

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Bayesian Meta-Prior Learning Using Empirical Bayes

Sareh Nabi

Foster School of Business, University of Washington, Seattle, Washington 98195, snabi@uw.edu

Houssam Nassif, Joseph Hong

Amazon, Seattle, Washington 98109, houssamn@amazon.com, josehong@amazon.com

Hamed Mamani

Foster School of Business, University of Washington, Seattle, Washington 98195, hmamani@uw.edu

Guido Imbens

Graduate School of Business, Stanford University, California 94305, gimbens@amazon.com

Adding domain knowledge to a learning system is known to improve results. In multi-parameter Bayesian frameworks, such knowledge is incorporated as a prior. On the other hand, the various model parameters can have different learning rates in real-world problems, especially with skewed data. Two often-faced challenges in Operation Management and Management Science applications are the absence of informative priors, and the inability to control parameter learning rates. In this study, we propose a hierarchical Empirical Bayes approach that addresses both challenges, and that can generalize to any Bayesian framework. Our method learns empirical meta-priors from the data itself, and uses them to decouple the learning rates of first-order and second-order features (or any other given feature grouping) in a Generalized Linear Model. As the first-order features are likely to have a more pronounced effect on the outcome, focusing on learning first-order weights first is likely to improve performance and convergence time. Our Empirical Bayes method clamps features in each group together and uses the deployed model's observed data to empirically compute a hierarchical prior in hindsight. We report theoretical results for the unbiasedness, strong consistency, and optimal frequentist cumulative regret properties of our meta-prior variance estimator. We apply our method to a standard supervised learning optimization problem, as well as an online combinatorial optimization problem in a contextual bandit setting implemented in an Amazon production system. Both during simulations and live experiments, our method shows marked improvements, especially in cases of small traffic. Our findings are promising, as optimizing over sparse data is often a challenge.

Key words: Informative Prior, Meta-Prior, Empirical Bayes, Bayesian Bandit, Generalized Linear Models, Thompson Sampling, Feature Grouping, Learning Rate

1. Introduction

We start by introducing and motivating the problem at hand.

1.1. Sequential Decisions and Prior Knowledge

Many scenarios arise in real-world applications where a decision-making agent must interact with an unknown environment in a sequential fashion, in which the agent receives a reward signal after each selected action (Chakraborty 2013, Sutton and Barto 2018, Nabi et al. 2018, Biswas et al. 2019, Sunar et al. 2019, Geng et al. 2020). The agent's objective is to find the best action-selection strategy that maximizes its long-run cumulative reward while navigating the existent uncertainty.

One application of such scenarios is dynamic pricing, where prices are dynamically adjusted in response to real-time demand and supply information. Nowadays dynamic pricing is practiced in e-commerce, travel industries, ride-hailing apps, smart retail electricity pricing, to name a few (Kakade et al. 2013, Ban and Keskin 2021, Ferreira et al. 2018, Lei et al. 2018, Keskin et al. 2020). Firms often face uncertainties such as unknown demand and customer preferences when making pricing decisions. Their challenge is to learn these uncertainties to better optimize revenue.

Online advertising is another example of sequential decision making under uncertainty (Nassif et al. 2016, Cheung and Simchi-Levi 2017, Sawant et al. 2018, Rafeian 2019, Rafeian and Yoganarasimhan 2021). According to a recent Interactive Advertising Bureau (IAB) report, the digital ad revenue in the US hit a record of \$49.5 billions in the first half of 2018, an increase of 23% compared to the same period last year (PwC and IAB 2018). In search advertising, an agent (like a search engine) makes decisions sequentially on what ads to show when a user renders a webpage. As the search engine's revenue depends on whether the user clicks on the ad, the agent's objective is to maximize revenue by accurately predicting the click-through-rate. Obtaining such predictions is a challenge, specially in cold start settings where the search engine has little or no knowledge of the user.

A main characteristic that arises in these aforementioned scenarios is the trade-off between the exploration involved in learning the unknown environment and the exploitation which occurs due to reward maximization. One way of solving this explore/exploit trade-off is by using the Multi-Armed Bandit (MAB) approach, originally proposed by Robbins (1952). Since then many studies have worked on MAB modeling framework and proposed algorithms to address the explore/exploit trade-off.

This work focuses on Bayesian MABs, which allow the encoding of prior knowledge into the learning process and thereby increase its efficiency (Ghavamzadeh et al. 2015, Lattimore and Szepesvári 2020). Most such algorithms are based on Thompson Sampling (TS) (Thompson 1933). The TS procedure selects an action proportionally to its probability of being optimal, conditioned on previous observations. This algorithm has demonstrated promising empirical results and is widely applied in industry and academia.

The choice of an appropriate informative prior in a Bayesian framework is often challenging, especially for reward models with multiple unknown parameters as in contextual MABs. This is why Bayesian MABs

often assume a non-informative prior. Recent literature demonstrates that the design of a prior can have significant impact on the performance of Bayesian bandits (Bastani et al. 2021). Can we construct a general framework that computes an empirical prior from early randomized data? Given a categorization of features in a model, can we decouple the learning rates of the parameters in each category? To the best of our knowledge, there is no work in the literature that addresses these questions broadly. This study aims at providing a general solution in estimating empirical priors for a wide range of applications that are modeled as Bayesian bandits or involve Bayesian learning.

1.2. Research Agenda and Our Approach

Our objective in this study is to utilize early randomized data to construct an experiment-specific hierarchical informative prior that can be applied to any individual problem instance modeled as a Bayesian framework. Such hierarchical informative priors can decouple the parameter learning rates, improve MAB optimization, and potentially lead to higher cumulative reward and shorter convergence time.

We are motivated by our production setting in Amazon, where we aim at optimizing a webpage layout with multiple components (Hill et al. 2017). In the simplest case, we have one feature per component value (e.g. “image2”), and one interaction feature between each pair of components (e.g. “image2 AND title1”). We refer to the first type of features as “first-order features”, and to the second type as “second-order features”. These two features types are functionally distinct. Based on analysis of fully-converged past experiments, we observe that the second-order feature weights are clustered more closely around the mean, while the first-order feature weights have larger absolute values, with a more pronounced effect on the outcome. Mathematically, this translates into a wider variance for the estimated first-order weights, when compared to the second-order ones. Can we leverage such meta-variance to learn the first-order order weights first, then move to the second-order weights as more data becomes available?

In our use-case, this question is particularly prescient, as the first and second order features have different incidence rates. Let n be the number of values per component (e.g., $n = 4$ possible images). As time progresses, we observe the first-order feature values at a rate of $1/n$, while we observe the second-order feature values at a slower rate of $1/n^2$. Although our proposed method does not require a different incidence rate, let alone a higher rate for the higher importance category, this research was partly motivated by this observation.

We leverage Empirical Bayes (EB) techniques to extract better priors from early randomized data. Empirical Bayes, also known as maximum marginal likelihood, allows for hyper-parameter estimation at the highest level of the hierarchical Bayes models. These point estimates can be obtained using either parametric or non-parametric approaches. In essence, EB is a statistical inference procedure where the prior distribution is estimated empirically (and frequentistically) from the data. It exploits the finding that large datasets of parallel situations carry within them their own Bayesian information (Efron and Hastie 2016).

Our approach applies EB to Bayesian bandits to compute an informative prior in hindsight, and then use this prior to improve cumulative reward and convergence time. We use the early random (or pseudo-random) MAB traffic to compute the empirical prior. We then rewind the bandit, re-training it on the same traffic, augmented with the empirical prior. By grouping the first-order and second-order features together, imposing a hierarchical prior, and estimating the priors by applying the EB framework, we are able to decouple the learning rates of these two categories of features (or any other arbitrary feature grouping). Note that although we have a bandit use-case, our method can also be applied to Bayesian learning in a standard classification setting, as we also show.

1.3. Contributions and Findings

Our main contribution in this study are three-folded:

First, we derive a general formula to compute a hierarchical meta-prior variance for any Bayesian framework using empirical Bayes. We then provide an estimator to estimate the EB prior variance empirically using early randomized data. Our meta-prior estimator derivation is general and applicable to a wide range of applications that are modeled as Bayesian learners. Our framework enables decoupling parameter learning rates by categorizing features arbitrarily, imposing a hierarchical prior per each category, and estimating the meta-priors using our EB approach.

Second, we report theoretical results for the unbiasedness, consistency, and regret properties of our meta-prior variance estimator. We provide a closed form for the expectation of our estimator and articulate the amount of bias in terms of the variances of feature estimates. We further show that our estimator is unbiased in our use-cases, which involve a Generalized Linear Model (GLM) used as a supervised classifier and as a bandit. We prove the strong consistency of our meta-prior variance estimator in the regime where the feature estimates are independent. Furthermore, we show that an EB TS algorithm retains an optimal $\tilde{O}(d^{3/2}\sqrt{T})$ regret bound.

Third, we validate our meta-prior variance estimator and test its generalization by performing simulations on the publicly available Adult UCI Machine Learning Repository dataset (Blake and Merz 1998). The task is to predict whether one's income exceeds a certain amount per year based on census data. We study the effects of first-order features, small batch data, prior reset time, and prior variance. We further examine the performance of empirical Bayes prior on our Amazon live production system, using a TS generalized linear multi-armed bandit with a probit link function (Hill et al. 2017). Specifically, we study the selection of a web page layout with the objective of maximizing the likelihood of a service purchase. Both during simulations and live experiments, we observe marked improvements when empirical Bayes prior is used instead of a non-informative prior. We observe a more pronounced improvement in small traffic situations.

The rest of this paper is structured as follows. In Section 2, we review the related literature. In Section 3, we derive our empirical Bayes prior estimation and provide statistical properties and regret bounds in Section 4. Pre-processing steps and simulation experiments using a real-world supervised dataset are explained

in Section 5. Section 6 reports results on bandit live experiments in an Amazon production setting. We discuss our findings and conclude in Section 7.

2. Related Literature

In applications of sequential decision making under uncertainty, we face an inherent trade-off between the exploration involved in learning unknown parameters, and the exploitation which occurs due to reward optimization. One way of modeling this trade-off is by using a Multi-Armed Bandit (MAB), originally proposed by Robbins (1952). Many algorithms have been proposed to solve stochastic MAB problems and generalized linear bandits (Gittins 1989, Auer et al. 2002, Garivier and Cappé 2011, Bubeck and Cesa-Bianchi 2012, Bubeck et al. 2012, Agrawal and Goyal 2013b, Cesa-Bianchi et al. 2017, Riquelme et al. 2018, Biswas et al. 2019, Ding et al. 2021, Jun et al. 2021).

Our focus in this work is on Bayesian MABs that often assume non-informative priors (Kaufmann et al. 2012a, Agrawal and Goyal 2013a, Sawant et al. 2018). Most such algorithms are based on Thompson Sampling, with promising empirical results across industry (Scott 2010, Chapelle and Li 2011, Tang et al. 2013, Teo et al. 2016). Multiple researchers investigated its theoretical guarantees. Kaufmann et al. (2012b) and Agrawal and Goyal (2013a) derived the asymptotic optimality of TS using non-informative priors and Bernoulli reward models. Korda et al. (2013) extended the work to one-dimensional exponential family bandits and derived an asymptotic optimality with Jeffreys prior. Refer to Russo et al. (2018) for a tutorial on TS algorithm.

Recent research reveals that incorporating prior knowledge in Bayesian MABs improves bandit performance and increase its efficiency (Ghavamzadeh et al. 2015, Lattimore and Szepesvári 2020, Bastani et al. 2020). Raina et al. (2006) propose a transfer learning algorithm that constructs a multivariate Gaussian prior for a supervised learning task. Honda and Takemura (2014) showed that TS is asymptotically optimal using a uniform prior when the reward model is Gaussian with unknown means and variances. But they prove that TS with the choice of Jeffreys prior and reference prior does not lead to asymptotic optimality. They report that choice of priors can play a crucial role in attaining the asymptotic optimality when the reward model has more than one unknown parameter.

Note that our prior in question is the Bayesian prior, a probability distribution that represents one's beliefs before data is encountered. One can also encode previous domain knowledge using first-order logic rules, as in Inductive Logic Programming and Statistical Relational Learning (Nassif et al. 2012, 2013, Getoor and Taskar 2007). Such a prior knowledge also improves bandit and classifier performance (Dutra et al. 2011, Kuusisto et al. 2013, Kakadiya et al. 2021), but is outside the scope of this research.

Bayesian (or expected) regret can be used as a performance measure in Bayesian MABs. By assuming an arbitrary prior distribution on reward models, Bayesian regret represents the average frequentist regret where the expectation is taken with respect to the prior on reward models. Russo and Van Roy (2014, 2016),

and Bubeck and Liu (2013) quantified the Bayesian regret of TS algorithm with an informative prior. Liu and Li (2016) also studied the sensitivity of TS performance over the choice of priors with the focus on the frequentist regret. They obtained regret bounds that depend on the probability mass a given prior places on the true unknown reward model.

Empirical Bayes (EB) is a known methodology to construct informative priors from early randomized data (Robbins 1955, Efron and Morris 1972, Morris 1983, Berger 1985). For comprehensive reviews on EB methods, refer to Carlin and Louis (2010), Efron (2012), and Maritz (2018). We implement our prior construction framework on a TS-based GLM bandit model with probit link function and achieve optimal asymptotic frequentist regret (Abeille et al. 2017, Hamidi and Bayati 2020).

The work of Bastani et al. (2021) is relevant to our study, as they propose a transfer learning algorithm that learns the meta-prior across experiments for similar products in a dynamic pricing setting. This study assumes the existence of a sequence of related experiments, where one leverages earlier bandit experiments to compute a meta-prior to be applied to later bandit experiments. Bastani et al. (2021) demonstrate empirically that leveraging the prior knowledge in Thompson Sampling, instead of following a prior-independent approach, improves performance. The approach of Bastani et al. (2021) only applies in transfer learning cases, where one is performing a sequence of related experiments. For example, in a dynamic pricing setting with unknown demand, if similar products are being sold, the transfer learning technique can help expedite the convergence to the optimal price for the later products in the experiments.

On the other hand, we do not attempt any transfer learning between experiments but rather bootstrap the early phase of an experiment to compute an empirical meta-prior for that experiment. To continue the example above, our framework aims at speeding up convergence to the optimal price using only the experiment's data, and hence can be applied in the absence of a proper sequence of similar experiments. To adopt our approach for the pricing experiments in Bastani et al. (2021), one has to run each pricing experiment separately, collect all the data across experiments, and apply our prior formulation on the aggregated data to estimate the prior. One could potentially combine both approaches in transfer learning settings, where we compute a transfer learning prior across the sequence of similar experiments and use it to seed the empirical prior computation for each individual experiment. We leave such an extension for future work.

There exists a rich literature in computer science on meta-learning, where one solves a novel task efficiently by using prior experience on similar tasks. Often assuming a hierarchical data structure, the objective in meta-learning problems is to train a model that generalizes across tasks (Schmidhuber 1987, Bengio et al. 1990, Naik and Mammone 1992, Thrun and Pratt 2012). Meta-learning problems are often formulated using a hierarchical Bayes model (Good 1980, Berger 2013) or an empirical Bayes model (Robbins 1992, Grant et al. 2018, Hu et al. 2020). Bayesian meta-learning helps with fast and robust adaptation to new tasks due to its probabilistic nature in measuring uncertainties (Ravi and Beatson 2019, Yoon et al. 2018, Nichol

et al. 2018, Gordon et al. 2019). More recent work shows gradient-based meta-learning approaches to be effective in few-shot learning (Finn et al. 2017, 2018, Rajeswaran et al. 2019, Zou and Lu 2020).

Our meta-prior framework can potentially be applied to dynamic Bayesian learning settings with incomplete information. One such application is Bayesian dynamic pricing (e.g., Harrison et al. (2012), Chen and Wang (2018)). Birge et al. (2021) and Kao et al. (2020) study this problem in the context of spread betting markets and service subscriptions, respectively. Keskin et al. (2020) study dynamic pricing problems with demand learning in the presence of customers with heterogeneous valuations for quality. Another avenue of potential application is optimal stopping problems in a regime of dynamic learning with incomplete information. Harrison and Sunar (2015) study this problem in the context of investment strategies. Sunar et al. (2021) study a sequential investment game in a duopoly setting in the presence of competition and dynamic learning opportunities.

3. Empirical Bayes Prior

We now present our model assumptions, and derive an empirical Bayes prior estimator.

3.1. Empirical Prior Derivation

Our problem of interest is one where features can be grouped into two or more groupings. For example, in a recommendation setting, one can distinguish between item and user features. In a personalization setting, one can distinguish between non-interaction features (e.g., “gender”) and interaction features (e.g., “female user likes action movies”).

By grouping the features into non-overlapping categories, we impose a Bayesian hierarchical model. We assume that each category C_k has a distinct hyperparameter meta-prior distribution with mean ν_k and variance τ_k^2 . The meta-prior distribution can be determined using experts’ knowledge for each specific application. In our experiments, we set the meta-prior distribution to be Gaussian as explained in section 3.3. We assume that each feature’s true effect μ_i is drawn i.i.d (independent and identically-distributed) from that feature’s category meta-prior, thus:

$$\mathbb{E}[\mu_i] = \nu_k, \quad \mathbb{V}[\mu_i] = \tau_k^2, \quad \forall i \in C_k. \quad (1)$$

$\mathbb{E}[\cdot]$ and $\mathbb{V}[\cdot]$ refer to the expectation and variance operators over random draws $i \in C_k$.

Let $\tilde{\mu}_i$ and $\tilde{\sigma}_i^2$ respectively denote the estimated (observed) feature effect and variance for feature i . Our approach assumes the existence of a model (like a Bayesian regression) that estimates $\tilde{\mu}_i$ and $\tilde{\sigma}_i^2$. We assume $\tilde{\mu}_i | \mu_i$ is drawn from a distribution with a mean equal to the true effect μ_i , and variance equal to its estimated variance $\tilde{\sigma}_i^2$. Based on our assumptions, the following holds for any feature $i \in C_k$:

$$\mathbb{E}[\tilde{\mu}_i | \mu_i] = \mu_i, \quad \mathbb{V}[\tilde{\mu}_i | \mu_i] = \tilde{\sigma}_i^2, \quad \mathbb{E}[\tilde{\mu}_i] = \nu_k, \quad \forall i \in C_k. \quad (2)$$

For each feature in category C_k , we perform variance decomposition:

$$\begin{aligned}\mathbb{V}[\tilde{\mu}_i] &= \mathbb{E}[\mathbb{V}[\tilde{\mu}_i|\mu_i]] + \mathbb{V}[\mathbb{E}[\tilde{\mu}_i|\mu_i]] \\ &= \mathbb{E}[\tilde{\sigma}_i^2] + \tau_k^2, \quad \forall i \in C_k.\end{aligned}\quad (3)$$

The second equality follows from our assumptions provided in Equations (1) and (2). Note that the choice of meta-prior distribution is irrelevant in deriving Equation (3).

We can now solve for the meta-prior variance τ_k^2 by:

$$\tau_k^2 = \mathbb{V}[\tilde{\mu}_i] - \mathbb{E}[\tilde{\sigma}_i^2]. \quad (4)$$

We can estimate τ_k^2 for each category C_k using Equation (4) as long as we have a procedure to estimate the feature effect $\tilde{\mu}_i$ and its variance $\tilde{\sigma}_i^2$ for any feature i in category C_k . We can obtain estimates of $\mathbb{E}[\tilde{\sigma}_i^2]$ and $\mathbb{V}[\tilde{\mu}_i]$ by their corresponding sample mean and sample variance, respectively. In the next two sections, we provide details of our model and framework for estimating $\tilde{\mu}_i$ and $\tilde{\sigma}_i^2$ and propose an estimator for the meta-prior variance, τ_k^2 .

3.2. Model Description

Our application of interest is a Bayesian generalized linear bandit, which interacts with the environment and collects data, while updating its estimates of $\tilde{\mu}_i$ and $\tilde{\sigma}_i^2$ at each time t . In general, these bandits model each feature effect using an underlying Gaussian distribution, starting with a standard normal non-informative prior (Filippi et al. 2010, Chapelle and Li 2011). Let $\mathcal{N}(\tilde{\mu}_{i,t}, \tilde{\sigma}_{i,t}^2)$ be the model weight distribution associated with feature i at time t . As time progresses, the bandit learns from its interactions with the environment, updating the features' weight distributions accordingly. In a stochastic setting, as $t \rightarrow \infty$, we have $\tilde{\sigma}_{i,t}^2 \rightarrow 0$ and $\tilde{\mu}_{i,t} \rightarrow \mu_i$, where μ_i is the true feature effect (Bubeck and Cesa-Bianchi 2012).

In our implementation, we use the Bayesian probit regression of Graepel et al. (2010), referred to as the **Bayesian LInear Probit (BLIP)** model. Its probability distribution is given by:

$$P(y|x, \tilde{\mu}) = \Phi\left(y \cdot \frac{\tilde{\mu}^T x}{\beta}\right), \quad (5)$$

where $\Phi(\cdot)$ denotes the CDF of the standard Gaussian distribution and β scales its steepness. For simplicity, we set $\beta = 1$ and drop it in the remainder of the paper. Here, $y \in \{-1, 1\}$ is the response variable. Moreover, x denotes the feature vector and $\tilde{\mu}$ refers to its associated weight vector.

We learn the feature weights $\tilde{\mu}$ in a Bayesian fashion. The weights are modeled as mutually independent random variables (hence the diagonal covariance matrix) following a Gaussian posterior distributions $\mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)$. The model uses a conjugate standard Gaussian prior $\mathcal{N}(0, 1)$ for the weights. The weight variances are upper bounded by the initial prior configuration, as every subsequent weight update is guaranteed to not increase the variance (Graepel et al. 2010).

We refer to this GLM classifier as BLIP. BLIP can be used as the Bayesian underlying model for a Thompson Sampling bandit, as we do in Section 6.1. We refer to such a bandit as a BLIP-based TS bandit. Next, we discuss how to apply our empirical Bayes method to BLIP, generating an EB BLIP classifier and an EB BLIP TS bandit.

3.3. Empirical Prior Estimation

For our use-case, we specify the meta-prior distribution as a Gaussian, the conjugate prior of the feature weight distribution. Each feature's true effect μ_i is drawn from that feature's category meta-prior:

$$\mu_i \sim \mathcal{N}(\nu_k, \tau_k^2), \quad \forall i \in C_k. \quad (6)$$

Let $\hat{\tau}_{k,t}^2$ denotes an estimator for the meta-prior variance τ_k^2 at time t . By applying Equation (4), we have

$$\hat{\tau}_{k,t}^2 = \widehat{\mathbb{V}[\tilde{\mu}_{i,t}]} - \widehat{\mathbb{E}[\tilde{\sigma}_{i,t}^2]} = \frac{\sum_{i \in C_k} (\tilde{\mu}_{i,t} - \hat{\nu}_{k,t})^2}{N_k - 1} - \frac{\sum_{i \in C_k} \tilde{\sigma}_{i,t}^2}{N_k}, \quad (7)$$

where N_k refers to the number of features in category C_k and superscript $\hat{\cdot}$ notation represents an estimator. We estimate $\widehat{\mathbb{E}[\tilde{\sigma}_{i,t}^2]}$ using the empirical mean of the observed variances. We use the basic sample variance formula for $\widehat{\mathbb{V}[\tilde{\mu}_{i,t}]}$. $\hat{\nu}_{k,t}$ is the empirical mean estimation of the meta-prior mean ν_k computed at time t for category C_k :

$$\hat{\nu}_{k,t} = \frac{\sum_{i \in C_k} \tilde{\mu}_{i,t}}{N_k}, \quad \forall C_k. \quad (8)$$

Equation (7) provides a formula for estimating the meta-prior variance τ_k^2 at time t per category C_k . $\sum_i (\tilde{\mu}_{i,t} - \hat{\nu}_{k,t})^2 / (N_k - 1)$ is the unbiased estimator of the true parameter τ_k^2 with estimation noise $\sum_i \tilde{\sigma}_{i,t}^2 / N_k$.

In our experiments, we set $\nu_k = 0$ to ensure the model is invariant to input feature sign changes. As long as one includes a bias (intercept) term in the generalized linear model, setting ν_k to any value has little effect as its value will be absorbed into the bias term. We indeed confirmed this hypothesis in preliminary experiments, where we compared setting $\nu_k = 0$ to setting it using Equation (8). Setting $\nu_k = 0$ gains one degree of freedom, ensuring the sample variance denominator is N_k and not $N_k - 1$. Hence, Equation (7) simplifies to:

$$\hat{\tau}_{k,t}^2 = \frac{\sum_{i \in C_k} [\tilde{\mu}_{i,t}^2 - \tilde{\sigma}_{i,t}^2]}{N_k}, \quad \forall C_k. \quad (9)$$

To ensure a non-degenerate τ_k^2 , one can enforce a minimum value threshold (see Section 5.1.1 for additional approaches). Our method is a parametric (zero-mean Gaussian) g-modeling EB approach (Efron and Hastie 2016) where we aim at estimating the variance τ_k^2 .

In practice, we start the model with a non-informative prior. At some small t (even at the end of batch $t = 1$, where all the data is random), we compute the empirical Bayes prior $\mathcal{N}(0, \hat{\tau}_{k,t}^2)$ for each feature grouping k using Equation (9). We then restart the model using the new informative prior and re-train it

using the data from the elapsed t batches. One can repeat these steps at multiple time-points t , each time re-computing a new τ_k^2 in an expectation-maximization fashion, but the approximation after one round is likely sufficient.

The computational complexity overhead of this method is negligible, as the meta-prior is computed and applied only once. Computing $\hat{\tau}_k^2$ in Equation (9) is $O(N)$, linear in the total number of features N . Once $\hat{\tau}_k^2$ is computed, one needs to re-train with data up to time t . The time complexity of such batch update is model dependent, and can be as small as $O(N)$ if a closed-form formula exists for already-processed batch updates. If one has to actually re-process the data, the training time complexity would be $O(tN)$, where t is a small constant $t \ll T$. In the latter case, the batch re-training can take place offline to further reduce the time overhead. In fact, most production systems have regular batch updates at pre-determined intervals, and this update can simply tag along the regular schedule.

Note that our derivation of $\hat{\tau}_{k,t}^2$ in this section is general and applicable to any Bayesian learning framework, and not tied to our specific bandit user-case. One can estimate the meta-prior variance using Equation (7) as long as estimates of the feature effect and its variance are available.

4. Theoretical Results

In this section, we first provide statistical properties of our meta-prior variance estimator. We then report an upper bound on the regret of our TS-based GLM bandit model.

4.1. Empirical Prior: Unbiasedness Property

In this section, we prove that the meta-prior variance estimator, $\hat{\tau}_k^2$, given in Equation (7), is unbiased as long as the estimates of $\tilde{\mu}_i$ are uncorrelated for all $i \in C_k$. The meta-prior variance can be computed at any time t and the results in this section hold for any time selected. Therefore, for clarity of notation, we drop the subscript t in our analysis.

PROPOSITION 1. *Let the meta-prior variance estimator, $\hat{\tau}_k^2$, for each category C_k , be*

$$\hat{\tau}_k^2 = \widehat{\mathbb{V}}[\tilde{\mu}_i] - \widehat{\mathbb{E}}[\tilde{\sigma}_i^2] = \frac{\sum_{i \in C_k} (\tilde{\mu}_i - \hat{\nu}_k)^2}{N_k - 1} - \frac{\sum_{i \in C_k} \tilde{\sigma}_i^2}{N_k}, \quad (10)$$

where $\tilde{\mu}_i$ and $\tilde{\sigma}_i^2$ respectively denote the estimated feature effect, and estimated variance, for feature i . The $\mathbb{E}[\cdot]$ and $\mathbb{V}[\cdot]$ operators refer to the expectation and variance over random draws $i \in C_k$. Moreover, $\hat{\nu}_k = \frac{\sum_{i \in C_k} \tilde{\mu}_i}{N_k}$ represents the empirical mean estimation of the meta-prior mean ν_k , and N_k refers to the number of features in category C_k . We have:

$$\mathbb{E}[\hat{\tau}_k^2] = \tau_k^2 - \frac{1}{N_k(N_k - 1)} \sum_{\substack{i,j \in C_k \\ i \neq j}} \text{Cov}[\tilde{\mu}_i, \tilde{\mu}_j]. \quad (11)$$

Proof One can re-write $\widehat{\mathbb{V}}[\widehat{\tilde{\mu}}_i]$ as:

$$\widehat{\mathbb{V}}[\widehat{\tilde{\mu}}_i] = \frac{1}{N_k - 1} \sum_{i \in C_k} (\tilde{\mu}_i - \hat{\nu}_k)^2 = \frac{1}{N_k - 1} \left[\sum_{i \in C_k} \tilde{\mu}_i^2 - N_k \hat{\nu}_k^2 \right]. \quad (12)$$

Taking the expectation on both sides and applying the linearity of expectation gives:

$$\begin{aligned} \mathbb{E}[\widehat{\mathbb{V}}[\widehat{\tilde{\mu}}_i]] &= \frac{1}{N_k - 1} \left[\sum_{i \in C_k} \mathbb{E}[\tilde{\mu}_i^2] - N_k \mathbb{E}[\hat{\nu}_k^2] \right] \\ &= \frac{1}{N_k - 1} \left[\sum_{i \in C_k} \left(\mathbb{E}[\tilde{\mu}_i]^2 + \mathbb{V}[\tilde{\mu}_i] \right) - N_k \left(\mathbb{E}[\hat{\nu}_k]^2 + \mathbb{V}[\hat{\nu}_k] \right) \right] \\ &= \frac{1}{N_k - 1} \left[\sum_{i \in C_k} \left(\nu_k^2 + \mathbb{V}[\tilde{\mu}_i] \right) - N_k \left(\nu_k^2 + \frac{1}{N_k^2} \left(\sum_{i \in C_k} \mathbb{V}[\tilde{\mu}_i] + \sum_{\substack{i, j \in C_k \\ i \neq j}} \text{Cov}[\tilde{\mu}_i, \tilde{\mu}_j] \right) \right) \right] \\ &= \frac{1}{N_k} \sum_{i \in C_k} \mathbb{V}[\tilde{\mu}_i] - \frac{1}{N_k(N_k - 1)} \sum_{\substack{i, j \in C_k \\ i \neq j}} \text{Cov}[\tilde{\mu}_i, \tilde{\mu}_j] \\ &= \frac{1}{N_k} \sum_{i \in C_k} (\tau_k^2 + \mathbb{E}[\tilde{\sigma}_i^2]) - \frac{1}{N_k(N_k - 1)} \sum_{\substack{i, j \in C_k \\ i \neq j}} \text{Cov}[\tilde{\mu}_i, \tilde{\mu}_j] \\ &= \tau_k^2 + \mathbb{E} \left[\frac{\sum_{i \in C_k} \tilde{\sigma}_i^2}{N_k} \right] - \frac{1}{N_k(N_k - 1)} \sum_{\substack{i, j \in C_k \\ i \neq j}} \text{Cov}[\tilde{\mu}_i, \tilde{\mu}_j]. \end{aligned} \quad (13)$$

The second equality follows from applying the variance formula ($\mathbb{V}[\mu] = \mathbb{E}[\mu^2] - \mathbb{E}[\mu]^2$) to the random variables $\tilde{\mu}_i$ and $\hat{\nu}_k$. The third equality follows from one of our assumptions provided in Equation (2), $\mathbb{E}[\tilde{\mu}_i] = \nu_k \forall i \in C_k$, and the definition of $\hat{\nu}_k = \frac{\sum_{i \in C_k} \tilde{\mu}_i}{N_k}$. In the penultimate equality, we substitute the value of $\mathbb{V}[\tilde{\mu}_i]$ using Equation (4).

Taking the expectation of $\hat{\tau}_k^2$ in Equation (10) and substituting $\mathbb{E}[\widehat{\mathbb{V}}[\widehat{\tilde{\mu}}_i]]$ from Equation (13) gives:

$$\mathbb{E}[\hat{\tau}_k^2] = \mathbb{E} \left[\frac{\sum_{i \in C_k} (\tilde{\mu}_i - \hat{\nu}_k)^2}{N_k - 1} \right] - \mathbb{E} \left[\frac{\sum_{i \in C_k} \tilde{\sigma}_i^2}{N_k} \right] = \tau_k^2 - \frac{1}{N_k(N_k - 1)} \sum_{\substack{i, j \in C_k \\ i \neq j}} \text{Cov}[\tilde{\mu}_i, \tilde{\mu}_j], \quad (14)$$

which concludes our proof. \square

COROLLARY 1. *The meta-prior variance estimator $\hat{\tau}_k^2$ for each category C_k , given in Equation (10), is unbiased when the estimates of $\tilde{\mu}_i$ and $\tilde{\mu}_j$, for all $i, j \in C_k$ where $i \neq j$, are uncorrelated.*

Proof This corollary simply follows from the results of Proposition 1. Since $\tilde{\mu}_i$ and $\tilde{\mu}_j$ are uncorrelated, $\text{Cov}[\tilde{\mu}_i, \tilde{\mu}_j] = 0$ for all $i, j \in C_k$ where $i \neq j$. Applying Equation (11) gives $\mathbb{E}[\hat{\tau}_k^2] = \tau_k^2$ which proves the unbiasedness of the meta-prior variance estimator $\hat{\tau}_k^2$. \square

COROLLARY 2. *The meta-prior variance estimator $\hat{\tau}_k^2$ for each category C_k , given in Equation (10), is unbiased when applied to our Bayesian GLM bandit model presented in Equation (5).*

Proof This result holds due to the modeling assumptions of Section 3.2. We assume that feature weights are mutually independent which results in a diagonal covariance matrix, making the estimates of $\tilde{\mu}_i$ uncorrelated. Therefore by Proposition 1 and Corollary 1, our meta-prior variance estimator in our Bayesian GLM model is unbiased. \square

4.2. Empirical Prior: Strong Consistency

In this section, we show strong consistency of our proposed meta-prior variance estimator in the regime where the estimates of $\tilde{\mu}_i$ and $\tilde{\sigma}_i^2$ are independent for all $i \in C_k$, and random variables $\tilde{\mu}_i$, $\tilde{\mu}_i^2$, and $\tilde{\sigma}_i^2$ have finite expectations and variances.

PROPOSITION 2. *The meta-prior variance estimator $\hat{\tau}_k^2$, given in Equation (10), is strongly consistent when the feature effect estimates $\tilde{\mu}_i$ and their observed variances $\tilde{\sigma}_i^2$ are independent $\forall i \in C_k$, and the expectation and variance exist for random variables $\tilde{\mu}_i$, $\tilde{\mu}_i^2$, and $\tilde{\sigma}_i^2$.*

Proof The strong consistency of our meta-prior variance estimator follows from the Kolmogorov Strong Law of Large Numbers; refer to Theorem 2.3.10 in Sen and Singer (1994) for details and proof.

THEOREM 1 (Kolmogorov Strong Law of Large Numbers). *(Thm 2.3.10 in Sen and Singer (1994)) Let X_i , $i \geq 1$, be independent random variables such that $\mathbb{E}[X_i] = \mu_i$ and $\mathbb{V}[X_i] = \sigma_i^2$ exists for every $i \geq 1$. Also let, $\bar{X}_n = \frac{\sum_{i=1}^n X_i}{n}$ and $\bar{\mu}_n = \frac{\sum_{i=1}^n \mu_i}{n}$ for $n \geq 1$. Then,*

$$\sum_{k \geq 1} k^{-2} \sigma_k^2 < \infty \implies \bar{X}_n \xrightarrow{a.s.} \bar{\mu}_n.$$

To prove strong consistency of meta-prior variance estimator $\hat{\tau}_k^2$, we first rewrite the sample variance estimator $\widehat{\mathbb{V}}[\tilde{\mu}_i]$ and apply the definition of $\hat{\nu}_k$:

$$\widehat{\mathbb{V}}[\tilde{\mu}_i] = \frac{\sum_{i \in C_k} (\tilde{\mu}_i - \hat{\nu}_k)^2}{N_k - 1} = \frac{\sum_{i \in C_k} \tilde{\mu}_i^2}{N_k - 1} - \frac{N_k}{N_k - 1} \hat{\nu}_k^2 = \frac{N_k}{N_k - 1} \left[\frac{\sum_{i \in C_k} \tilde{\mu}_i^2}{N_k} - \left(\frac{1}{N_k} \sum_{i \in C_k} \tilde{\mu}_i \right)^2 \right].$$

Hence the meta-prior variance estimator $\hat{\tau}_k^2$, given in Equation (10), becomes:

$$\hat{\tau}_k^2 = \frac{N_k}{N_k - 1} \left[\frac{\sum_{i \in C_k} \tilde{\mu}_i^2}{N_k} - \left(\frac{1}{N_k} \sum_{i \in C_k} \tilde{\mu}_i \right)^2 \right] - \frac{\sum_{i \in C_k} \tilde{\sigma}_i^2}{N_k}. \quad (15)$$

In the following, we provide details on how to apply Kolmogorov Strong Law of Large Numbers on the summation terms with random variables $\tilde{\mu}_i$, $\tilde{\mu}_i^2$, and $\tilde{\sigma}_i^2$ and prove strong consistency of $\hat{\tau}_k^2$.

Given our assumptions, estimates $\tilde{\mu}_i \forall i \in C_k$ are independent and the expectation of $\tilde{\sigma}_i^2$ exists. The latter implies $\exists \epsilon$ s.t. $\mathbb{E}[\tilde{\sigma}_i^2] \leq \epsilon, \forall i \in C_k$. We further have $\mathbb{E}[\tilde{\mu}_i] = \nu_k$ per our assumption in Equation (2), and Equation (3) results in $\mathbb{V}[\tilde{\mu}_i] = \tau_k^2 + \mathbb{E}[\tilde{\sigma}_i^2] \leq \tau_k^2 + \epsilon$. Note that $\sum_{i \geq 1}^{\infty} \frac{1}{i^2} = \frac{\pi^2}{6}$. We can now apply the Kolmogorov strong law of large numbers on $\tilde{\mu}_i$:

$$\sum_{i \geq 1} i^{-2} \mathbb{V}[\tilde{\mu}_i] = \sum_{i \geq 1} \frac{\tau_k^2 + \mathbb{E}[\tilde{\sigma}_i^2]}{i^2} \leq \frac{\pi^2}{6} (\tau_k^2 + \epsilon) < \infty \implies \frac{1}{N_k} \sum_{i \in C_k} \tilde{\mu}_i \xrightarrow{a.s.} \frac{1}{N_k} \sum_{i \in C_k} \mathbb{E}[\tilde{\mu}_i] = \nu_k. \quad (16)$$

We apply similar arguments to the independent random variables $\tilde{\mu}_i^2, \forall i \in C_k$. We have

$$\mathbb{E}[\tilde{\mu}_i^2] = \mathbb{V}[\tilde{\mu}_i] + \mathbb{E}[\tilde{\mu}_i]^2 = \tau_k^2 + \mathbb{E}[\tilde{\sigma}_i^2] + \nu_k^2,$$

using the variance formula, our assumption in Equation (2), and Equation (3). Moreover, we assume that $\mathbb{V}[\tilde{\mu}_i^2]$ exists which implies that $\exists \epsilon'$ s.t. $\mathbb{V}[\tilde{\mu}_i^2] \leq \epsilon', \forall i \in C_k$. Therefore, $\sum_{i \geq 1} i^{-2} \mathbb{V}[\tilde{\mu}_i^2] = O(1)$. Hence, by the Kolmogorov strong law of large numbers, as $N_k \rightarrow \infty$, almost surely

$$\frac{1}{N_k} \sum_{i \in C_k} \tilde{\mu}_i^2 \xrightarrow{a.s.} \frac{1}{N_k} \sum_{i \in C_k} \mathbb{E}[\tilde{\mu}_i^2] = \tau_k^2 + \nu_k^2 + \frac{1}{N_k} \sum_i \mathbb{E}[\tilde{\sigma}_i^2]. \quad (17)$$

By applying the results in Equations (16) and (17), we obtain

$$\begin{aligned} \lim_{N_k \rightarrow \infty} \widehat{\mathbb{V}}[\widehat{\tilde{\mu}}_i] &= \lim_{N_k \rightarrow \infty} \frac{N_k}{N_k - 1} \left[\lim_{N_k \rightarrow \infty} \frac{\sum_{i \in C_k} \tilde{\mu}_i^2}{N_k} - \left(\lim_{N_k \rightarrow \infty} \frac{1}{N_k} \sum_{i \in C_k} \tilde{\mu}_i \right)^2 \right] \\ &= \tau_k^2 + \nu_k^2 + \frac{1}{N_k} \sum_i \mathbb{E}[\tilde{\sigma}_i^2] - \nu_k^2 \\ &= \tau_k^2 + \frac{1}{N_k} \sum_i \mathbb{E}[\tilde{\sigma}_i^2]. \end{aligned} \quad (18)$$

Finally, by taking the limit of the meta-prior variance estimator $\hat{\tau}_k^2$, given in Equation (10), we obtain

$$\begin{aligned} \lim_{N_k \rightarrow \infty} \hat{\tau}_k^2 &= \lim_{N_k \rightarrow \infty} \widehat{\mathbb{V}}[\tilde{\mu}_i] - \lim_{N_k \rightarrow \infty} \widehat{\mathbb{E}}[\tilde{\sigma}_i^2] \\ &= \tau_k^2 + \frac{1}{N_k} \sum_i \mathbb{E}[\tilde{\sigma}_i^2] - \lim_{N_k \rightarrow \infty} \frac{\sum_{i \in C_k} \tilde{\sigma}_i^2}{N_k} \\ &= \tau_k^2 + \frac{1}{N_k} \sum_i \mathbb{E}[\tilde{\sigma}_i^2] - \frac{1}{N_k} \sum_i \mathbb{E}[\tilde{\sigma}_i^2] \\ &= \tau_k^2, \end{aligned}$$

where the second equality applies the results in (18). The third equality follows from the Kolmogorov strong law of large numbers applied to the independent random variables $\tilde{\sigma}_i^2 \forall i \in C_k$ that have finite expectation and variance due to our assumptions. This concludes the strong consistency proof of our proposed meta-prior variance estimator. \square

4.3. Bandit Cumulative Regret Upper Bound

In this section, we provide an upper bound on the regret of our EB BLIP-based bandit method, described in Section 3.2. In the stochastic GLM bandit setting, at every round t , we sample $\dot{\mu} \sim \tilde{\mu}_t$, and play arm:

$$x_t(\dot{\mu}) = \arg \max_{x \in X} \Phi(\dot{\mu}^T x). \quad (19)$$

Let us denote the optimal weight vector as μ^* and the optimal arm as $x^* = x(\mu^*)$. The instantaneous regret at time step t is the difference between the expected reward of the optimal arm x^* and the selected arm x_t , written as $\Phi(\mu^{*T} x^*) - \Phi(\mu^{*T} x_t)$. Over a time horizon of length T , the cumulative regret becomes:

$$R^{GLM}(T) = \sum_{t=1}^T \left[\Phi(\mu^{*T} x^*) - \Phi(\mu^{*T} x_t) \right]. \quad (20)$$

Following Lemma 4 of Abeille et al. (2017), and under a set of assumptions, the cumulative regret $R^{GLM}(T)$ is upper bounded. We restate the Lemma for completeness and report on our regret bounds in the two corollaries that follow afterward.

LEMMA 1. (Lemma 4 of Abeille et al. (2017))

Assume that:

1. The arm set X is a bounded closed subset of \mathbb{R}^d such that $\|x\| \leq 1, \forall x \in X$.
2. There exists a known $S \in \mathbb{R}^+$ such that $\|\mu^*\| \leq S$.
3. The noise process is a martingale difference sequence given a filtration, and is conditionally R -subgaussian for some constant $R \geq 0$.
4. The link function Φ is continuously differentiable, Lipschitz with constant k_Φ , and its derivative is lower bounded by $c_\Phi = \inf_{\mu \in \mathbb{R}^d, x \in X} \Phi'(\mu^T x) > 0$.

Then, with probability $1 - \delta$, the cumulative regret $R^{GLM}(T)$ of the GLM TS bandit is upper bounded by:

$$R^{GLM}(T) \leq \frac{k_\Phi}{c_\Phi} (\beta_T(\delta') + \gamma_T(\delta')(1 + 2/p)) \sqrt{2Td \log \left(1 + \frac{T}{\lambda}\right)} + \frac{2k_\Phi \gamma_T(\delta')}{pc_\Phi} \sqrt{\frac{8T}{\lambda} \log \frac{4}{\delta}}. \quad (21)$$

Here $\delta' = \frac{\delta}{4T}$, λ denotes the regularization parameter, p is an anti-concentration probability bound,

$$\beta_t(\delta) = R \sqrt{2 \log \frac{(\lambda + t)^{d/2} \lambda^{-d/2}}{\delta}} + \sqrt{\lambda} S, \quad \gamma_t(\delta) = \beta_t(\delta') \sqrt{cd \log(c'd/\delta)}, \quad (22)$$

and c, c' are positive concentration constants.

COROLLARY 3. The cumulative regret of the EB BLIP-based TS bandit is upper bounded by Equation (21) and is thus of order of $\tilde{O}(d^{3/2} \sqrt{T})$.

Proof Our approach ascribes to the Lemma 1 assumptions, and is hence bounded. Assumption 1 is satisfied by using a 1-hot encoding for categorical features, and centering and scaling numerical features. In practice, one can take a large S to satisfy Assumption 2. As our reward is binary, Assumption 3 is satisfied, with $R = 1/2$. For Assumption 4, our $\Phi(\cdot)$ is the CDF of the standard Gaussian distribution:

$$\Phi(x) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{x}{\sqrt{2}} \right) \right], \quad \text{where} \quad \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt. \quad (23)$$

By taking the derivative of $\Phi(x)$ w.r.t x and substituting for $\frac{d}{dx} \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} e^{-x^2}$, we see that it is continuously differentiable: $\frac{d\Phi(x)}{dx} = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$. By the mean-value theorem, any differentiable function with bounded derivative is Lipschitz. $\Phi(\cdot)$ is Lipschitz, since its first derivative is bounded by $0 < \frac{d\Phi(x)}{dx} < \frac{1}{\sqrt{2\pi}}$, giving us a suitable value for $k_\Phi = \frac{1}{\sqrt{2\pi}}$. Note that $c_\Phi \rightarrow 0^+$ as $\mu^T x \rightarrow \pm\infty$. As both x and μ are bounded by assumptions 1 and 2, this guarantees $c_\Phi > 0$.

Now that all four assumptions are satisfied, Lemma 1 holds, and the cumulative regret of our EB BLIP bandit is bounded according to Equation (21). This bound is of order of $\tilde{O}(d^{3/2} \sqrt{T})$. \square

COROLLARY 4. *The cumulative regret of the BLIP-based TS bandit is of order $\tilde{O}(d^{3/2}\sqrt{T})$, whether an empirical Bayes prior $\mathcal{N}(0, \tau_k^2)$ per feature category C_k is applied, or a non-informative prior $\mathcal{N}(0, 1)$ is used.*

Proof Following the results in Lemma 1 and Corollary 3, we obtain a frequentist regret bound of order $\tilde{O}(d^{3/2}\sqrt{T})$. This bound is optimal for GLM TS bandits (Hamidi and Bayati 2020). This bound applies to our bandit, regardless if we use the empirical Bayes prior $\mathcal{N}(0, \tau_k^2)$, or the non-informative prior $\mathcal{N}(0, 1)$. The difference between these two alternatives resides in different values of the anti-concentration (p) and concentration (c, c') constants introduced in Abeille et al. (2017), with no significant effect on the bound.

Starting with a $\mathcal{N}(0, 1)$ prior, we get anti-concentration probability $p = \frac{e^{-\frac{1}{2}}}{4\sqrt{\pi}}$ and concentration constants $c = c' = 2$, as per Appendix A in Abeille et al. (2017). We repeat this exercise starting from a $\mathcal{N}(0, \tau_k^2)$ prior per feature category C_k . Following the same strategy as in the standard Gaussian prior case, and applying the Gaussian tail bound provided in Chang et al. (2011), we get anti-concentration probability $p = \frac{e^{\left(\frac{1}{2} - \frac{1}{\tau_{\min}^2}\right)}}{4\sqrt{\pi}}$ and concentration constants $c = 2\tau_{\max}^2$ and $c' = 2$. Here $\tau_{\min} = \min_{k \in C_k} \tau_k$, and $\tau_{\max} = \max_{k \in C_k} \tau_k$. For the sake of completeness, we provide details on how to compute these values in the appendix. \square

Let *regretBound* refer to the r.h.s of Equation (21). We note that $\text{regretBound} \propto \frac{\gamma}{p} \propto \frac{\sqrt{c}}{p} \propto \tau_{\max} e^{\frac{1}{\tau_{\min}^2}}$. As $\tau_{\min} \rightarrow 0$, $\text{regretBound} \rightarrow \infty$. This finding is concordant with previous findings (Honda and Takemura 2014, Liu and Li 2016). As τ_k shrinks, exploration is severely hampered, the bandit fails to identify the best arm, and the the cumulative regret grows linearly. As our approach decouples the learning rates of different categories of features, $\tau_k \approx 0$ reflects a category with no variance, and one can use this information to prune or cluster such features. We discuss guardrails in Section 7.

The regret bound and asymptotic optimal regret analysis we provide is very specific to the finite horizon bandit problem (i.e., a Markov Decision Process (MDP) with a single state). In general, it is unclear that these methods would generalize to arbitrary Bayesian decision problems. In particular cases such as tabular MDPs, finite-time regret bounds have been worked out in a frequentist sense (Simchowitz and Jamieson 2019).

As for an analysis of exact Bayesian optimality and regret, we refer the reader to Chapter 35 of Lattimore and Szepesvári (2020). As a quick summary, in infinite horizon discounted Bayesian k-armed bandits, Bayesian optimal policy takes the form of a Gittins index when the objective is to maximize total expected discounted reward (Theorem 35.9). For finite horizon bandits, the Bayesian optimal policy is often not an index policy. Computing Bayesian optimal policy for large finite horizon settings is intractable. For small horizon and number of arms, the dynamic programming formulation for achieving the Bayesian optimal policy has a large number of states. But the benefits of obtaining exact optimality may be worth the extensive computation effort in these cases.

Feature	#Values	Feature	#Values
occupation	14	workclass	7
education	16	education-num	10
relationship	6	marital status	7
gender	2	race	5
hours-per-week	12	native-country	41
capital gain	3	capital loss	2
age	20		

Table 1 Adult dataset features and number of possible values.

5. Supervised Learning Dataset and Experiments

In order to validate our method and test its generalization, we first report simulations on a public optimization dataset from a different domain.

5.1. Supervised Learning Dataset

We pick the Adult dataset available from the UCI Machine Learning Repository (Blake and Merz 1998). The target is to predict whether one’s income exceeds \$50,000 per year based on census data. The train and test datasets have 30,162 and 15,060 observations respectively, after removing rows with empty feature values. The dataset has 13 categorical features (see Table 1).

The original 13 features constitute our first-order features. We mimic our use-case of interest by generating second-order features through pairwise combinations of first-order features. We add all these second-order interaction terms (total of 78) to the Adult dataset, forming our raw feature vector. The first-order and second-order groupings form a natural categorization for the $\hat{\tau}_k^2$ computation given in Equation (9).

5.1.1. Bootstrapping for Small Traffic To simulate batch updates, we divide the training set equally into 6 batches of 5027 examples each. We train the probit linear model on the first batch, starting with a $\mathcal{N}(0, 1)$ prior. We then apply our EB approach described in Section 3 to compute empirical priors $\hat{\tau}_1^2$ over the first-order features, and $\hat{\tau}_2^2$ over the second-order features. To our initial surprise, this resulted in degenerative negative $\hat{\tau}_k^2$ values.

This degenerative result stems from an inappropriate initial prior, to our point. As we had only observed a small number of data points, the default $\mathcal{N}(0, 1)$ initial prior still dominates the first-batch posterior used to compute the empirical Bayes prior. The value of $\tilde{\mu}_{i,1}$ was still close to 0 and the value of $\sigma_{i,1}^2$ close to 1. As $\tilde{\mu}_{i,1} < \tilde{\sigma}_{i,1}$, Equation (9) returns $\hat{\tau}_k^2 < 0$. This observation suggests that our first batch is too small for empirical prior estimation.

The necessary traffic required to ensure a non-degenerative $\hat{\tau}_k^2$ depends on the update mechanism of the inference model used, the initial prior, and the magnitude of the true meta-prior variance (see Equation 7). To pursue a model-free discussion, we define large traffic as one where the first update batch is enough to

generate viable $\hat{\tau}_k^2 > 0, \forall k$. Medium traffic is one where $\hat{\tau}_k^2 > 0$ within a user-specified number of batches (in our case, one week if one training batch used per day), and small traffic requires waiting even longer.

One solution is to simply consume additional batches before computing the empirical prior. This ensures that $\tilde{\mu}_i$ and $\tilde{\sigma}_i^2$ deviate enough from their initial values to produce a viable $\hat{\tau}_k^2$. This solution works well in medium traffic cases, where waiting for the next batch can resolve the issue. In small traffic situations, one may not be able to wait long enough.

Here lies the essence of our challenge. Starting from an informative prior helps most in small traffic cases and small traffic hinders the proper computation of the empirical prior. Observe though that the value of $\hat{\tau}_k^2$ increases with the number of samples due to the reduction of the posterior variance $\tilde{\sigma}_{i,1}^2$. One can use techniques like bootstrapping, ensembles, or epoch training to artificially boost the number of samples enough to produce a proper $\hat{\tau}_k^2$.

For our simulation dataset, and given the first batch of 5K instances, we bootstrap it (sampling with replacement) repeatedly creating multiple sets of 5K instances each. We treat each bootstrapped set as a training epoch. After each training epoch, we compute $\hat{\tau}_k^2$. We stop training as soon as we attain positive $\hat{\tau}_k^2$ values, which was 8 epochs in our case. We also tried training epochs without bootstrapping, with similar results. We keep $\mathcal{N}(0, 1)$ as the initial prior, since this is the default non-informative prior used in many applications. Once we compute the empirical prior $\hat{\tau}_k^2$, we restart the model using the new informative prior and re-train it with the initial non-bootstrapped first batch.

5.1.2. Data Pre-Processing One may suspect that not all features in a model are relevant. To that effect, and to balance the overfitting aspect of small data bootstrapping, we prune the bootstrapped model by applying adaptive lasso. We pick adaptive lasso due to its oracle properties, as it can identify the right subset of features to retain as if the true underlying model was given in advance (Zou 2006).

The objective function for adaptive lasso is

$$\widehat{W} := \arg \min_W \|y - XW\|^2 + \lambda \sum_i \zeta_i |w_i|, \quad (24)$$

where λ is the shrinkage parameter estimated using cross validation, W is the model weight vector where each w_i is the weight coefficient of feature i , y is the response vector, and X the design matrix. ζ is the adaptive weight vector and is defined as $\zeta_i = \frac{1}{|\hat{w}_i|^\gamma}$, where \hat{w}_i is an initial estimate of coefficient w_i obtained by performing ridge regression in our simulations. Moreover, γ is a positive constant that adjusts the adaptive lasso weight vector and is set to $\gamma = 1$ in our experiments. Figure 1 shows an example of adaptive lasso feature selection using the glmnet package (Friedman et al. 2010).

In this section, we use the terms “batch” and “day” interchangeably, as it is more intuitive to think in terms of a temporal framework. Note that one could use any arbitrary time period, “day” is simply an example. Although we focus on batch updates, our method works equally well with online updates.




Figure 1 Mean squared error and coefficients selected by adaptive lasso. The left plot shows the mean squared error where the dashed line denotes the optimal value of the shrinkage parameter λ obtained through cross validation. The right plot shows the adaptive lasso’s pruning order where the top values denote the number of non-zero weight coefficients as the shrinkage parameter increases.

5.2. Scenarios Description

We consider three scenarios that only differ on what happens at the end of a pre-specified time step t . At the start of the experiment, we initialize the three models using a standard normal prior $\mathcal{N}(0, 1)$. At the end of each day, the models are trained in batch with the day’s observed data. At the end of day t , the following three scenarios are applied:

BLIP: Our base model, which stands for Bayesian Linear Probit as mentioned earlier. We update the model in batch with day t data. This is the Bayesian regression method that our model builds and improves upon, and the basis of our production GLM TS bandit in Section 6. One can use any Bayesian GLM model, we opt for BLIP as a fast and scalable state-of-the-art production system (He et al. 2014, Ding et al. 2019).

BLIPBayes: We reset the model. We bootstrap all the data observed until day t (as per Section 5.1.1) and train the model on the bootstrapped data. We then use the Empirical Bayes computation of Section 3 to compute separate informative priors $\hat{\tau}_k^2$ for the first and second order features. We then restart the model with a $\mathcal{N}(0, \hat{\tau}_1^2)$ prior for the first-order features and a $\mathcal{N}(0, \hat{\tau}_2^2)$ prior for the second-order features. We update the new EB model with the original data observed up until day t . This is our proposed model.

BLIPTwice: We update the model twice, first with the same bootstrapped data as BLIPBayes and second with all the data observed until day t . The rationale is that BLIPTwice uses the same amount of data as used by BLIPBayes. BLIPTwice is not a proper optimization algorithm, we only include it when studying the effect of data reuse.

We train the three models on the same batches. At the end of day t , adaptive lasso prunes the feature space identically for the three scenarios. We start evaluating the models after the day t update. At the end

of each batch, we evaluate the models on the holdout testing set using binary log loss (cross-entropy). Let n be the number of observed data points, $y_j \in \{0, 1\}$ the true binary label of instance j , and P_j the model's predicted probability of $y_j = 1$, then:

$$\text{Log Loss} = -\frac{1}{n} \sum_{j=1}^n [y_j \ln P_j + (1 - y_j) \ln(1 - P_j)]. \quad (25)$$

Note that the output variable in BLIP model given in Equation (5) takes values in $\{-1, 1\}$ set. We adjust our labels before computing the log loss given in Equation (25).

5.3. Simulation Results

We now provide the results of our experiments using the Adult dataset. Specifically, we study the effect of first-order features and batch size across our three scenarios, BLIP, BLIPBayes, and BLIPTwice. We also study the effect of prior reset time and prior variance on BLIPBayes, and compare them against BLIP. All these experiments used the same total traffic amount for training: 30,162 records from the Adult dataset.

5.3.1. First-Order Feature Effect: Based on the findings of Section 5.1.2, and after day $t = 1$, we compute the hierarchical empirical prior over a bootstrapped data pruned using adaptive lasso. Adaptive lasso retained 7 first-order and 11 second-order features. As we surmise that the first-order features may hold more predictive power during the first batches, we test retaining all 13 first-order features alongside the 11 pruned second-order features. Keeping all first-order features results in $\hat{\tau}_1^2 = 0.852$ and $\hat{\tau}_2^2 = 0.241$. Keeping only selected first-order features returns $\hat{\tau}_1^2 = 0.714$ and $\hat{\tau}_2^2 = 0.460$.

Figure 2 plots the log loss of our scenarios. We observe that BLIPBayes outperforms BLIP and BLIPTwice in both cases. We suspect that the bad performance of BLIPTwice is due to overfitting to the first t batches. Retaining all first-order features improves the three methods' prediction accuracies. We also note that keeping all first-order features results in a markedly better performance for BLIPBayes. This may be due to a better estimate of $\hat{\tau}_1^2$. In our subsequent experiments, we retain all first-order features, and use adaptive lasso to prune the second-order features.

5.3.2. Effect of Prior Reset Time: In this experiment, we reset the empirical prior at the end of day $t = 3$, after observing 15,000 samples. Adaptive lasso pruning retains 11 second-order features. EB results in $\hat{\tau}_1^2 = 0.862$ and $\hat{\tau}_2^2 = 0.414$.

Figure 3 compares the log loss for BLIP, BLIPBayes with $t = 1$ reset, and BLIPBayes with $t = 3$ reset. BLIPBayes gives lower log loss values and therefore higher prediction accuracy when using more data for constructing the prior. As we observed when comparing all versus selected first-order features, more data improves performance.




Figure 2 Log loss with prior reset after day 1. Figure (a) keeps all 13 first-order features and only the 11 second-order features selected by adaptive lasso. Figure (b) only keeps the 7 first-order and 11 second-order features selected by adaptive lasso.




Figure 3 Comparing log loss of no prior reset (BLIP), prior reset after day 1 (BLIPBayes-1d), and after day 3 (BLIPBayes-3d). Using more data to compute the empirical prior improves performance.

5.3.3. Small Batch Dataset: In this experiment, we divide the train dataset into thirty batches, each with 1000 data points. Our objective is to observe the performance of EB when training occurs on a longer period with smaller data batches. We reset the prior at $t = 1$, but this time we bootstrap and train over 12,000 instances. Adaptive lasso retains 21 second-order features. EB results in $\hat{\tau}_1^2 = 0.799$ and $\hat{\tau}_2^2 = 0.132$.

Figure 4 plots the log loss over the thirty days. BLIPBayes outperforms both other methods. What is remarkable is that the BLIPBayes advantage persists over the whole range, indicating that such a method can be especially valuable for small batch training.




Figure 4 Log loss over smaller batch dataset. The total traffic amount is unchanged.

5.3.4. Effect of Prior Variance τ_k^2 : It is possible that our improvements simply stem from the fact that our empirical prior variance is below 1, its non-informative counterpart. This would not adhere to the finding uncovered in Section 4.3, where *regretBound* increases as τ_k shrinks. To shed additional light on this issue, we simulate the worse-case scenario of $\hat{\tau}_1 = \hat{\tau}_2$, and vary their values from $\approx 0^+$ onward. Recall from our first experiment setting that EB returns $\hat{\tau}_1^2 = 0.852$ and $\hat{\tau}_2^2 = 0.241$ by following the formula given in Equation (9). We experiment with additional $\hat{\tau}_k^2$ settings, namely $\hat{\tau}_1^2 = \hat{\tau}_2^2 = 5$, $\hat{\tau}_1^2 = \hat{\tau}_2^2 = 0.1$, and $\hat{\tau}_1^2 = \hat{\tau}_2^2 = 0.01$.

Figure 5 plots the log loss for the aforementioned scenarios, alongside the EB values (“optimal”) and BLIP. We observe that the optimal BLIPBayes consistently outperforms all other variations. We also note that BLIP (with its $\hat{\tau}_1^2 = \hat{\tau}_2^2 = 1$ priors) outperforms the non-optimal BLIPBayes versions (with a negligible overlap with $\hat{\tau}_k^2 = 0.1$ after day 5). This suggests that one needs to set the hierarchical priors in a principled manner, and that $\hat{\tau}_k^2 < 1$ is not necessarily better. We also note that $\hat{\tau}_k^2 = 0.01$ achieves the worst results by far, concordant with our Section 4.3 finding, and that $\hat{\tau}_k^2 = 5$ under-performs $\hat{\tau}_k^2 = 0.1$. This suggests that erring towards a large prior variance is more easily overcome by data than erring towards a small prior variance (Bastani et al. 2020). In fact, in a Bayesian setting, the relative effect of the observed data on the posterior increases with a larger prior variance.

6. MAB Live Experiments

We now examine the performance of EB on the live production system described in Hill et al. (2017).

6.1. Experimental Settings

The problem we address is the selection of a layout A of a web page with the objective of maximizing the expected value of a binary reward $y \in \{-1, 1\}$. Each layout A is generated from a common template that




Figure 5 Comparing the effect of τ_k^2 on performance. We varied $\hat{\tau}_k^2$ away from its optimal EB values, $\hat{\tau}_1^2 = 0.852$ and $\hat{\tau}_2^2 = 0.241$.

contains D widgets representing the contents of the page. The i^{th} widget has n_i alternative variations for its content; however, for simplicity of notation, we assume equal number of variations per widget denoted by n . A web page has thus n^D possible layouts. We represent a layout as $A \in \{1, 2, \dots, n\}^D$, a D -dimensional vector where $A[i]$ denotes the content chosen for the i^{th} widget.

We model this scenario as a stochastic contextual multi-armed bandit with linear reward. Our model is a TS generalized linear MAB with a probit link function and $\mathcal{N}(0, 1)$ prior (Kveton et al. 2020). Its core is the same model used for classification in Section 3.2. At each time step t , we select the layout to present to a user using Thompson Sampling. We sample the model parameters from their posterior and pick the layout that maximizes the reward, following Equation (19).

Let B_A be the featured representation of a given layout. Let $R = +1$ indicate the user took the desired action which refers to the service purchase in our live experiment and $R = -1$ indicate otherwise. The reward probability is modeled as:

$$p_r(R|A) = \Phi(R * B_A^\top W), \quad (26)$$

where W presents the unknown parameter vector and Φ refers to the CDF of the standard Normal distribution. Note that this is the same model as in Equation (5). Although the original formulations take user context into account, this work investigates the case where the features only reflect the layout content. Fur-

thermore, we only consider pair-wise interactions and assume that the effects of higher order interaction terms are negligible. Therefore, our linear model becomes:

$$B_A^\top W = W^0 + \sum_{i=1}^D W_i^1(A) + \sum_{j=1}^D \sum_{k=j+1}^D W_{j,k}^2(A), \quad (27)$$

where W^0 is a common bias weight, $W_i^1(A)$ is a weight associated with the content in the i^{th} widget, and $W_{j,k}^2(A)$ is a weight for the interaction between the contents of the j^{th} and k^{th} widgets. For EB computations, W^1 are the weights of the first-order features, and W^2 the weights of the second-order features.

At each time step t , we select the layout to present to a user using Thompson Sampling. A layout is shown proportional to its probability of being optimal. We sample the model parameters from their posterior and pick the layout that maximizes the reward. Let \tilde{W}_t be the sampled weights. We present the layout that solves the following optimization problem:

$$A_t = \arg \max_A B_A^\top \tilde{W}_t. \quad (28)$$

As we have N^D layouts, a standard MAB results in N^D arms. For small N and D values, we solve this optimization using exhaustive search. This is the case in our live experiments in Section 6.2. In settings with large N and D values, we recommend deploying a greedy hill climbing optimization procedure with random restarts (Casella and Berger (2002)).

In our live experiments, we aim at optimizing a message that promotes the purchase of an Amazon service. The message had $D = 4$ widgets with $n_i = \{2, 3\}$ possible options per widget, for a total of $d = 24$ distinct combinatorial layouts. The target is binary, whether the customer purchased the service or not. The messages were shown to the selected customers during a browsing session on Amazon.com.

We performed A/B tests with three treatments, a production baseline algorithm, the standard probit bandit, and EB applied to the probit bandit. We randomly diverted a constant subset of our traffic to this experiment, with the standard and EB bandits receiving equal shares. The baseline algorithm has a different pool of messages, and dynamically adjusts to seasonal shifts in an adversarial manner. In order to study the effect of the empirical prior reset, we disable seasonality adjustment for both standard and EB bandits, deploying them as stochastic MABs.

6.2. Live Results

We start both bandits with a random phase, where the bandits allocate traffic equally between their 24 layouts. We exclude this random phase from our plots and analysis. We then compute the empirical prior and re-train the EB MAB on the random data. Unlike the simulation experiments, we do not prune features nor do we bootstrap the data for EB computation. As we do not know the ground truth, we can not compute regret nor log loss. Instead, we compute the cumulative success rates of each bandit.

6.2.1. Traffic Effect In the first experiment, we diverted a constant traffic percentage to our MABs. We set the random phase to 3 time units, after which we ran the experiment for 12 time units during a non-seasonal time frame. EB resulted in $\hat{\tau}_1^2 = 0.613$ and $\hat{\tau}_2^2 = 0.195$. Figure 6 plots the cumulative success rate of each bandit, relative to the baseline's final cumulative success rate.

The EB MAB clearly dominates the standard MAB in cumulative rewards. We notice that EB MAB stabilizes after five time units, while the standard MAB needs eight to plateau. This is an indication that EB MAB converged faster. After $t = 8$, both bandits attain the same performance, maintaining similar success rates. From $t = 1$, the empirical prior focused the exploration of the EB bandit, while the standard bandit had to fully explore the space of suboptimal layouts. This resulted in improved cumulative gain, as reflected by the gap in the relative cumulative success rates.

The dip at $t = 5$ is due to a production update delay that affected all algorithms, including the baseline. The standard bandit suffers the most, as it was still exploring and deployed a suboptimal policy. EB proves robust to such disruptions, a crucial quality as data issues are commonly observed in any production system.

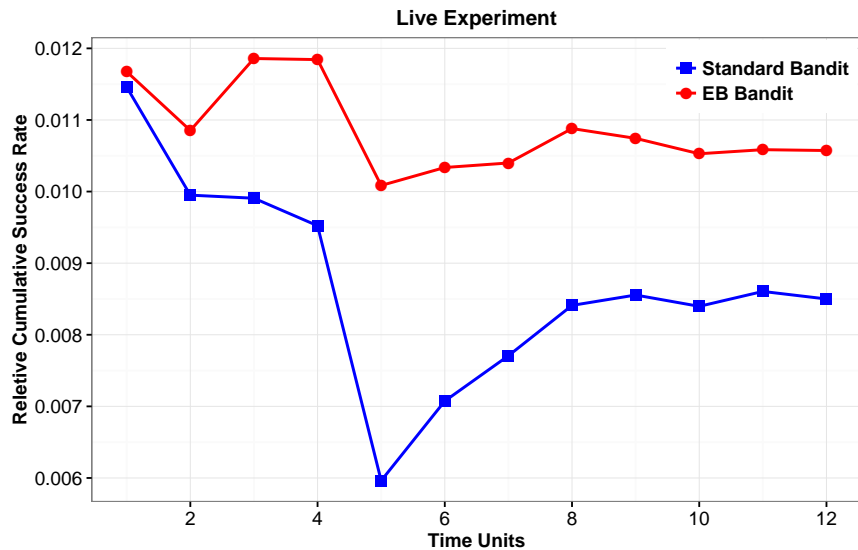


Figure 6 Cumulative success rate in first live experiment, relative to baseline final cumulative success rate.

We followed this experiment with a shorter one (4 time units total) during a fixed season, where we reduced the traffic by a quarter, and reduced the random phase to 1 time unit. EB resulted in $\hat{\tau}_1^2 = 0.489$ and $\hat{\tau}_2^2 = 0.087$. At the end of the experiment, we compared the cumulative performance of both MABs against the baseline. Using two-tailed proportion z-test with pooled variance, EB bandit had a 0.05 p-value, and standard bandit a 0.11 p-value. EB MAB significantly outperformed the baseline, and effectively converged, while the standard MAB did not. Another indication that EB can be most valuable for smaller traffic cases. We report significance against the baseline to measure convergence. Testing EB MAB versus standard MAB is not significant and is irrelevant for us, as both converge to the same solution given enough data.

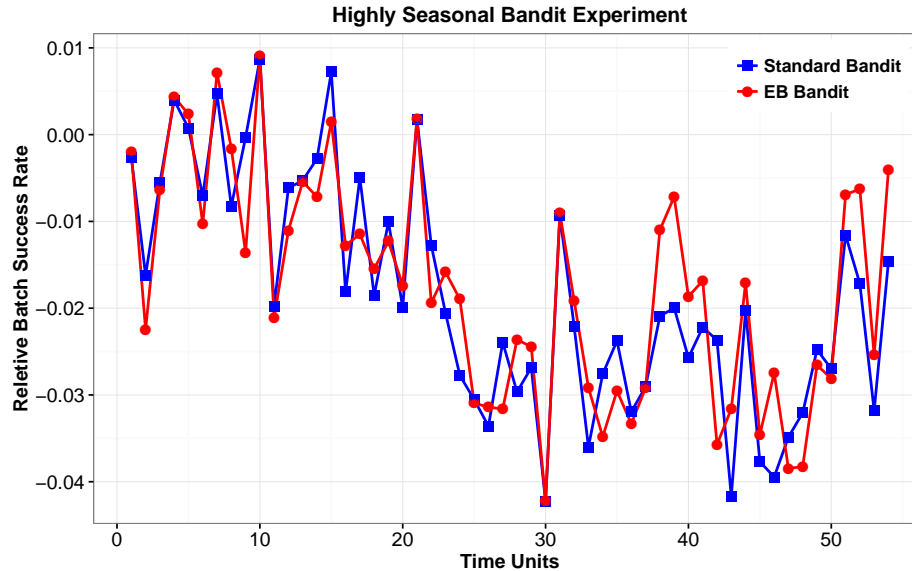


Figure 7 Batch success rate in highly-seasonal third live experiment, relative to the baseline success rate.

6.2.2. Seasonality Effect To test the impact of seasonality, we ran our final experiment over a holiday season, with significant pre and post seasonal changes. The experiment extended over 54 time units, following a random phase of 2 time units, and did not suffer from major production system delays. We dialed up the traffic assigned to each MAB, doubling the traffic on average, while tripling it in the first two weeks. EB resulted in $\hat{\tau}_1^2 = 0.448$ and $\hat{\tau}_2^2 = 0.133$.

At this level of high traffic, both bandits behaved indistinguishably. The empirical prior had little to no effect, as expected in higher traffic situations. As the two bandits' cumulative success rates were almost identical, Figure 7 plots their batched (per unit time) success rates. One can see how they fluctuate similarly. We also note that seasonality had an adverse effect on the bandits, as both lagged behind the seasonality-aware baseline. Recall that we had disabled seasonality adjustments for both bandits, but not for the baseline. This explains their negative relative success rates compared to the baseline.

7. Discussion and Conclusion

We conclude by analyzing and discussing our findings.

7.1. Discussion

It is interesting that, empirically, most $\hat{\tau}_k^2$ priors computed during our experiments had values below 1. Recall from Section 4.3 that τ_k^2 affects the bandit regret bound through constants p and c . Parametrizing p, c using τ_k^2 in Equations (21) and (22), we notice that $regretBound(\tau_k^2 < 1) > regretBound(\tau_k^2 = 1)$. Here $regretBound$ refers to the r.h.s of Equation (21). As a small τ_k^2 reduces exploration, this could explain the higher regret bound (Liu and Li 2016). But in our experiments, we observe that EB constantly outperforms the standard bandit. This suggests that imposing an EB hierarchical prior to decouple the learning rates of the first and second-order features could have a tangible effect on the regret.

In a supervised learning setting, the unbiasedness and consistency results of Section 4, coupled with our experimental results, suggest that our EB computation method is likely to yield a lower log loss. In a bandit setting, we do observe an empirical improvement, even when *regretBound* may be interpreted as suggesting otherwise for $\tau_k^2 < 1$. This seems to indicate a trade-off between loosing on exploration potential, and gaining by introducing a meta-prior that decouples learning rates. In cases where there is a difference between the effects of different categories' features, i.e. when true $\tau_1 \neq \tau_2 \neq 1$, EB would find a better prior that is likely to improve on the regret. Since the choice of the hierarchical prior variance has no effect on the asymptotic regret $\tilde{O}(d^{3/2}\sqrt{T})$, one may loose little by applying the EB method, even if true $\tau_1 = \tau_2 = 1$.

In all of our simulation and live experiments, we always had $\hat{\tau}_1^2 > \hat{\tau}_2^2$, reflecting a difference in the effects of the first and second order features. This confirms our initial conjecture, that the second-order features are likely to be less important. This result is likely to generalize to many applications.

By grouping the first-order and second-order features together and imposing a hierarchical prior, we effectively clamped each category's weights together. Since the relative effect of the observed data on the posterior increases with a larger prior variance, our method is putting more weight on the first-order features, and is shrinking the second-order effect. The model thus focuses on learning the first-order effects first. This may explain the increased stability and convergence speed of the EB model.

Of interest is how the improvement is correlated with the amount of available data. Our findings suggest that the EB improvement is most marked in cases of low to medium traffic, and is lost at high traffic. This is promising, as low traffic cases are the hardest to optimize. At very low traffic, direct computation of empirical prior variances may fail, with $\hat{\tau}_k^2 < 0$. Even if $\hat{\tau}_k^2 \approx 0^+$, one needs guardrails to preserve a sublinear regret bound, as per Section 4.3. One may either wait longer before computing the prior (see Section 5.3.2), bootstrap from available data (see Section 5.1.1), prune the feature space (see Section 5.1.2), add a minimal τ_k threshold, or perform transfer learning. We leave investigating the last two suggestions for future work.

In order to study the effect of the empirical prior reset, we disabled seasonality adjustment for both standard and EB bandits in Section 6. For short-period experiments, it is a reasonable approximation to assume stationarity. We do not recommend disabling seasonality adjustment for long-period experiments, as indicated in Section 6.2.2. Since we compute $\hat{\tau}_k^2$ close to the start of the experiment, the effects of seasonality and algorithmic adjustments have little bearing on its value. We confirmed this finding in subsequent experiments. For long-period experiments where seasonality adjustment is paramount, the choice of the prior becomes less relevant with additional data.

As with most Bayesian techniques, we model the meta-prior using the conjugate distribution of the posterior. As we used a probit regression, the feature weights are modeled as Gaussians, with a Gaussian conjugate prior. This work seeks to improve the starting prior of an existing model, without questioning the validity of the prior's distribution family. If one has no valid distribution shape for the prior, techniques like

the ones explored in Luedtke et al. (2020) could help. Even if the prior distribution is unknown or misspecified, one could invoke the Central Limit Theorem to justify the choice of a Gaussian variance representation for Equation (9).

7.2. Future Work

Our findings raise multiple questions and opportunities for future work. In this paper, we decouple learning rates by imposing a hierarchical meta-prior, and estimate its variance prior using empirical Bayes. One could investigate other methods to estimate the priors, and compare techniques to ensure non-degenerative empirical priors. This work specifically focuses on first and second-order features, with different incidence rates. One could further investigate the effect of incidence rates on the meta-prior computation.

We used EB to effectively bootstrap an existing model using its own data. It could be possible to apply EB to different but related use-cases as a transfer learning technique (Bastani et al. 2021). This may be valuable if the related use-case has a large data volume, while the target scenario is highly sparse. Our technique requires pre-specified feature categories, a harder challenge is to automatically infer the feature categories.

Separately, decoupling learning rates is a valuable by-product of our method. It can be attributed to: (i) the features are grouped in different categories and share different priors, and (ii) the EB approach provides more accurate per-category prior estimation. One can further study the dynamics of decoupling learning rates using EB, as well as comparing it to other methods.

Another potential research question is whether our meta-prior framework can be extended to dynamic Bayesian learning in general settings. One such application is dynamic learning with optimal stopping, which occurs in launch/innovation and investment strategies. As one example, optimal policies in Harrison and Sunar (2015) are characterized based on optimal lower and upper critical beliefs that depend upon posterior estimations. A more informative prior leads to more accurate estimations of these parameters at any given time. This, in turn, can increase or decrease the length of the learning continuation region described in Harrison and Sunar (2015). Data deficiency issues may also arise in optimal stopping problems due to the potential shorter learning horizon. Leveraging existing structure given the application domain and problem formulation may benefit meta-prior learning and help overcome potential data deficiencies.

7.3. Conclusion

This study proposes an Empirical Bayes method that leverages early observed data in a given experiment to compute category-specific informative meta-priors for that experiment. Such informative meta-priors have the added effect of decoupling learning rates of different categories of features. We show that our GLM-based estimator is unbiased, and prove its strong consistency in the regime where the feature estimates are independent. Furthermore, we show that our EB BLIP-based TS bandit algorithm retains an optimal $\tilde{O}(d^{3/2}\sqrt{T})$ regret bound. We apply this method to decouple the learning rates of first and second order features in a supervised learning and a bandit setting. Our empirical results reveal that our Empirical Bayes

technique leads to higher cumulative rewards and lower convergence time for bandits and improves prediction accuracy for classifiers. Of special note are the observed improvements in cases of smaller traffic, leading us to believe that Empirical Bayes may offer an adequate solution for the challenges of sparse-data optimization.

Acknowledgments

The authors thank Gregory Duncan, Sham Kakade, Milos Curcic, Lalit Jain, Omid Rafeian, Devin Didericksen, Yi Liu, Karthik Mohan, Emily Ikeda-Flowers, Calvin Kwok, Zachary Austin, and Jason Yang for their support and helpful discussions.

References

- Abeille M, Lazaric A, et al. (2017) Linear thompson sampling revisited. *Electronic Journal of Statistics* 11(2):5165–5197.
- Agrawal S, Goyal N (2013a) Further optimal regret bounds for thompson sampling. *AISTATS*, 99–107.
- Agrawal S, Goyal N (2013b) Thompson sampling for contextual bandits with linear payoffs. *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 127–135.
- Auer P, Cesa-Bianchi N, Fischer P (2002) Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.
- Ban GY, Keskin NB (2021) Personalized dynamic pricing with machine learning. *Management Science* .
- Bastani H, Bayati M, Khosravi K (2020) Mostly exploration-free algorithms for contextual bandits. *Management Science* 67(3):1329–1349.
- Bastani H, Simchi-Levi D, Zhu R (2021) Meta dynamic pricing: Transfer learning across experiments. *Management Science* .
- Bengio Y, Bengio S, Cloutier J (1990) *Learning a synaptic learning rule* (Citeseer).
- Berger JO (1985) *Statistical decision theory and Bayesian analysis*. Springer Series in Statistics (New York: Springer), 2nd edition.
- Berger JO (2013) *Statistical decision theory and Bayesian analysis* (Springer Science & Business Media).
- Birge JR, Feng Y, Keskin NB, Schultz A (2021) Dynamic learning and market making in spread betting markets with informed bettors. *Operations Research* .
- Biswas A, Pham TT, Vogelsong M, Snyder B, Nassif H (2019) Seeker: Real-time interactive search. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2867–2875.
- Blake CL, Merz CJ (1998) UCI repository of machine learning databases.
- Bubeck S, Cesa-Bianchi N (2012) Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5(1):1–122.

- Bubeck S, Cesa-Bianchi N, Kakade S (2012) Towards minimax policies for online linear optimization with bandit feedback. *Annual Conference on Learning Theory (COLT)*, volume 23, 1–14.
- Bubeck S, Liu CY (2013) Prior-free and prior-dependent regret bounds for thompson sampling. *Advances in Neural Information Processing Systems (NIPS)*, 638–646.
- Carlin BP, Louis TA (2010) *Bayes and empirical Bayes methods for data analysis* (Chapman and Hall/CRC).
- Casella G, Berger RL (2002) *Statistical inference*, volume 2 (Duxbury Pacific Grove, CA).
- Cesa-Bianchi N, Gentile C, Lugosi G, Neu G (2017) Boltzmann exploration done right. *Advances in Neural Information Processing Systems (NeurIPS)*, 6284–6293.
- Chakraborty B (2013) *Statistical methods for dynamic treatment regimes* (Springer).
- Chang SH, Cosman PC, Milstein LB (2011) Chernoff-type bounds for the gaussian error function. *IEEE Transactions on Communications* 59(11):2939–2944.
- Chapelle O, Li L (2011) An empirical evaluation of thompson sampling. *Advances in neural information processing systems (NeurIPS)*, 2249–2257.
- Chen X, Wang Z (2018) Bayesian dynamic learning and pricing with strategic customers, available at SSRN.
- Cheung WC, Simchi-Levi D (2017) Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models, available at SSRN 3075658.
- Ding Q, Hsieh CJ, Sharpnack J (2021) An efficient algorithm for generalized linear bandit: Online stochastic gradient descent and thompson sampling. *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1585–1593.
- Ding W, Govindaraj D, Vishwanathan S (2019) Whole page optimization with global constraints. *Knowledge Discovery in Databases Conference (KDD)*, 3153–3161.
- Dutra I, Nassif H, Page D, Shavlik J, Strigel RM, Wu Y, Elezaby ME, Burnside E (2011) Integrating machine learning and physician knowledge to improve the accuracy of breast biopsy. *American Medical Informatics Association Symposium (AMIA)*, 349–355.
- Efron B (2012) *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1 (Cambridge University Press).
- Efron B, Hastie T (2016) *Computer age statistical inference*, volume 5 (Cambridge University Press).
- Efron B, Morris C (1972) Limiting the risk of bayes and empirical bayes estimators—Part II: The empirical bayes case. *Journal of the American Statistical Association* 67(337):130–139.
- Ferreira KJ, Simchi-Levi D, Wang H (2018) Online network revenue management using thompson sampling. *Operations research* 66(6):1586–1602.
- Filippi S, Cappé O, Garivier A, Szepesvári C (2010) Parametric bandits: The generalized linear case. *Advances in Neural Information Processing Systems 23 (NIPS)*, 586–594.

- Finn C, Abbeel P, Levine S (2017) Model-agnostic meta-learning for fast adaptation of deep networks. *Proceedings of the International Conference on Machine Learning (ICML)*, 1126–1135.
- Finn C, Xu K, Levine S (2018) Probabilistic model-agnostic meta-learning. *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*.
- Friedman J, Hastie T, Tibshirani R (2010) Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software* 33(1):1–22.
- Garivier A, Cappé O (2011) The KL-UCB algorithm for bounded stochastic bandits and beyond. *COLT*, 359–376.
- Geng S, Nassif H, Manzanares C, Reppen M, Sircar R (2020) Deep PQR: Solving inverse reinforcement learning using anchor actions. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 3431–3441.
- Getoor L, Taskar B (2007) *Introduction to statistical relational learning* (The MIT Press).
- Ghavamzadeh M, Mannor S, Pineau J, Tamar A, et al. (2015) Bayesian reinforcement learning: A survey. *Foundations and Trends® in Machine Learning* 8(5-6):359–483.
- Gittins JC (1989) *Multi-armed Bandit Allocation Indices* (Wiley, Chichester, NY).
- Good IJ (1980) Some history of the hierarchical bayesian methodology. *Trabajos de estadística y de investigación operativa* 31(1):489–519.
- Gordon J, Bronskill J, Bauer M, Nowozin S, Turner RE (2019) Meta-learning probabilistic inference for prediction. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Graepel T, Candela JQ, Borchert T, Herbrich R (2010) Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. *Proceedings of the 27th International Conference on Machine Learning (ICML)*, 13–20.
- Grant E, Finn C, Levine S, Darrell T, Griffiths T (2018) Recasting gradient-based meta-learning as hierarchical bayes. *International Conference on Learning Representations (ICLR)*.
- Hamidi N, Bayati M (2020) On worst-case regret of linear thompson sampling, arXiv preprints.
- Harrison JM, Keskin NB, Zeevi A (2012) Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science* 58(3):570–586.
- Harrison JM, Sunar N (2015) Investment timing with incomplete information and multiple means of learning. *Operations Research* 63(2):442–457.
- He X, Pan J, Jin O, Xu T, Liu B, Xu T, Shi Y, Atallah A, Herbrich R, Bowers S, Candela JQn (2014) Practical lessons from predicting clicks on ads at facebook. *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, 1–9.
- Hill DN, Nassif H, Liu Y, Iyer A, Vishwanathan S (2017) An efficient bandit algorithm for realtime multivariate optimization. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 1813–1821.

- Honda J, Takemura A (2014) Optimality of thompson sampling for gaussian bandits depends on priors. *Artificial Intelligence and Statistics (AISTATS)*, 375–383.
- Hu SX, Moreno PG, Xiao Y, Shen X, Obozinski G, Lawrence ND, Damianou A (2020) Empirical bayes transductive meta-learning with synthetic gradients. *International Conference on Learning Representations (ICLR)*.
- Jun KS, Jain L, Mason B, Nassif H (2021) Improved confidence bounds for the linear logistic model and applications to linear bandits. *Proceedings of the International Conference on Machine Learning (ICML)*.
- Kakade SM, Lobel I, Nazerzadeh H (2013) Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research* 61(4):837–854.
- Kakadiya A, Natarajan S, Ravindran B (2021) Relational boosted bandits. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Kao YM, Keskin NB, Shang K (2020) Bayesian dynamic pricing and subscription period selection with unknown customer utility, available at SSRN 3722376.
- Kaufmann E, Cappé O, Garivier A (2012a) On bayesian upper confidence bounds for bandit problems. *AISTATS*, 592–600.
- Kaufmann E, Korda N, Munos R (2012b) Thompson sampling: An asymptotically optimal finite-time analysis. *International Conference on Algorithmic Learning Theory*, 199–213 (Springer).
- Keskin NB, Li Y, Sunar N (2020) Data-driven clustering and feature-based retail electricity pricing with smart meters, available at SSRN.
- Korda N, Kaufmann E, Munos R (2013) Thompson sampling for 1-dimensional exponential family bandits. *Advances in Neural Information Processing Systems (NeurIPS)*, 1448–1456.
- Kuusisto F, Dutra I, Nassif H, Wu Y, Klein ME, Neuman HB, Shavlik J, Burnside ES (2013) Using machine learning to identify benign cases with non-definitive biopsy. *International Conference on E-Health Networking, Application & Services (HealthCom)*, 283–285.
- Kveton B, Zaheer M, Szepesvari C, Li L, Ghavamzadeh M, Boutilier C (2020) Randomized exploration in generalized linear bandits. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2066–2076.
- Lattimore T, Szepesvári C (2020) *Bandit algorithms* (Cambridge University Press).
- Lei Y, Jasin S, Sinha A (2018) Joint dynamic pricing and order fulfillment for e-commerce retailers. *Manufacturing & Service Operations Management* 20(2):269–284.
- Liu CY, Li L (2016) On the prior sensitivity of thompson sampling. *Proceedings of the 27th International Conference on Algorithmic Learning Theory (ALT)*, 321–336.
- Luedtke A, Chung I, Sofrygin O (2020) Adversarial monte carlo meta-learning of optimal prediction procedures, arXiv preprints.
- Maritz JS (2018) *Empirical Bayes Methods with Applications* (Chapman and Hall/CRC).

- Morris CN (1983) Parametric empirical bayes inference: Theory and applications. *Journal of the American Statistical Association* 78(381):47–55.
- Nabi R, Kanki P, Shpitser I (2018) Estimation of personalized effects associated with causal pathways. *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 239.
- Naik DK, Mammone RJ (1992) Meta-neural networks that learn by learning. *International Joint Conference on Neural Networks (IJCNN)*, volume 1, 437–442.
- Nassif H, Cansizlar KO, Goodman M, Vishwanathan SVN (2016) Diversifying music recommendations. *Proceedings of the International Conference on Machine Learning (ICML) Workshops*.
- Nassif H, Kuusisto F, Burnside ES, Shavlik J (2013) Uplift modeling with ROC: An SRL case study. *International Conference on Inductive Logic Programming (ILP)*, pp. 40–45.
- Nassif H, Santos Costa V, Burnside ES, Page D (2012) Relational differential prediction. *European Conference on Machine Learning (ECML)*, 617–632.
- Nichol A, Achiam J, Schulman J (2018) On first-order meta-learning algorithms, arXiv preprint arXiv:1803.02999.
- PwC, IAB (2018) IAB internet advertising revenue report. Technical report, Interactive Advertising Bureau.
- Rafeian O (2019) Optimizing user engagement through adaptive ad sequencing. Technical report, University of Washington.
- Rafeian O, Yoganasimhan H (2021) Targeting and privacy in mobile advertising. *Marketing Science* 40(2):193–218.
- Raina R, Ng AY, Koller D (2006) Constructing informative priors using transfer learning. *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, 713–720.
- Rajeswaran A, Finn C, Kakade S, Levine S (2019) Meta-learning with implicit gradients. *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*.
- Ravi S, Beatson A (2019) Amortized bayesian meta-learning. *International Conference on Learning Representations (ICLR)*.
- Riquelme C, Tucker G, Snoek J (2018) Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. *International Conference on Learning Representations (ICLR)*.
- Robbins H (1952) Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* 58(5):527–535.
- Robbins H (1955) *An empirical Bayes approach to statistics* (Office of Scientific Research, US Air Force).
- Robbins HE (1992) An empirical bayes approach to statistics. *Breakthroughs in statistics*, 388–394 (Springer).
- Russo D, Van Roy B (2014) Learning to optimize via posterior sampling. *Mathematics of Operations Research* 39(4):1221–1243.
- Russo D, Van Roy B (2016) An information-theoretic analysis of thompson sampling. *The Journal of Machine Learning Research* 17(1):2442–2471.

- Russo DJ, Van Roy B, Kazerouni A, Osband I, Wen Z, et al. (2018) A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning* 11(1):1–96.
- Sawant N, Namballa CB, Sadagopan N, Nassif H (2018) Contextual multi-armed bandits for causal marketing. *Proceedings of the International Conference on Machine Learning (ICML'18) Workshops*.
- Schmidhuber J (1987) *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. Ph.D. thesis, Technische Universität München.
- Scott SL (2010) A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry* 26(6):639–658.
- Sen PK, Singer JM (1994) *Large sample methods in statistics: an introduction with applications*, volume 25 (CRC press).
- Simchowitz M, Jamieson KG (2019) Non-asymptotic gap-dependent regret bounds for tabular mdps. *Advances in Neural Information Processing Systems (NeurIPS)*, 1151–1160.
- Sunar N, Birge JR, Vitavasiri S (2019) Optimal dynamic product development and launch for a network of customers. *Operations Research* 67(3):770–790.
- Sunar N, Yu S, Kulkarni VG (2021) Competitive investment with bayesian learning: Choice of business size and timing. *Operations Research* .
- Sutton RS, Barto AG (2018) *Reinforcement learning: An introduction* (MIT press).
- Tang L, Rosales R, Singh A, Agarwal D (2013) Automatic ad format selection via contextual bandits. *Proceedings of the 22nd ACM international conference on Information & Knowledge Management (CIKM)*, 1587–1594.
- Teo CH, Nassif H, Hill D, Srinivasan S, Goodman M, Mohan V, Vishwanathan S (2016) Adaptive, personalized diversity for visual discovery. *Proceedings of the ACM Conference on Recommender Systems (RecSys)*, 35–38.
- Thompson WR (1933) On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25(3/4):285–294.
- Thrun S, Pratt L (2012) *Learning to learn* (Springer Science & Business Media).
- Yoon J, Kim T, Dia O, Kim S, Bengio Y, Ahn S (2018) Bayesian model-agnostic meta-learning. *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS)*, 7343–7353.
- Zou H (2006) The adaptive lasso and its oracle properties. *Journal of the American statistical association* 101(476):1418–1429.
- Zou Y, Lu X (2020) Gradient-EM bayesian meta-learning. *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, 20865–20875.

Appendix. Computation of Bandit Cumulative Regret Constants

Here, we provide details on how to compute anti-concentration probability p and constant values c and c' in the regret upper bound given in Equation (21) and Corollary 4.

Abeille et al. (2017) show that TS can be defined as a randomized algorithm that is built upon a Regularized Least Square (RLS) estimate, given past history of arms pulled and rewards observed. They show that TS samples from a perturbed parameter that depends on this RLS estimate and a multivariate distribution that follows anti-concentration and concentration properties. We first restate these properties and compute their corresponding parameter values for our EB BLIP TS model.

DEFINITION 1. (As provided in Abeille et al. 2017). D^{TS} is a multivariate distribution on \mathbb{R}^d absolutely continuous with respect to Lebesgue measure which satisfies the following properties:

1. *Anti-concentration property:* There exists a strictly positive probability p such that for any $u \in \mathbb{R}^d$ with $\|u\| = 1$,

$$P_{\eta \sim D^{TS}} \left(u^T \eta \geq 1 \right) \geq p. \quad (29)$$

2. *Concentration property:* There exists c, c' positive constants such that $\forall \delta \in (0, 1)$,

$$P_{\eta \sim D^{TS}} \left(\|\eta\| \leq \sqrt{cd \log \frac{c'd}{\delta}} \right) \geq 1 - \delta. \quad (30)$$

In our EB BLIP, $\eta \sim \mathcal{N}(0, \mathbf{I}_{d,\tau})$ where $\mathbf{I}_{d,\tau}$ denotes the diagonal covariance matrix of size $d * d$ with diagonal values τ_i^2 , $\forall i = 1, \dots, d$. Let η_{τ_i} denote the i -th component of η . Hence, $\eta_{\tau_i} \sim \mathcal{N}(0, \tau_i^2)$, $\forall i = 1, \dots, d$. We have

$$P \left(u^T \eta \geq 1 \right) \geq P \left(\eta_{\tau_{\min}} \geq 1 \right) = \frac{1}{2} \operatorname{erfc} \left(\frac{1}{\tau_{\min} \sqrt{2}} \right) \geq \frac{e^{\frac{1}{2} - \frac{1}{\tau_{\min}^2}}}{4\sqrt{\pi}}, \quad (31)$$

where $\tau_{\min} = \min_k \tau_k$ and $\eta_{\tau_{\min}} \sim \mathcal{N}(0, \tau_{\min}^2)$. The first inequality holds since the anti-concentration property must be satisfied for any arbitrary $u \in \mathbb{R}^d$. The second equality and third inequality are due to the definition of the Gaussian Q-function and its lower bound, given in Chang et al. (2011). Let $Q(x)$ denote the Gaussian Q-function for $X \sim \mathcal{N}(0, 1)$,

$$Q(x) := P(X > x) = \frac{1}{2} \operatorname{erfc} \left(\frac{x}{\sqrt{2}} \right) \quad \text{where} \quad \operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt, x > 0. \quad (32)$$

Chang et al. (2011) prove that $f(x) = \alpha e^{-\beta x^2}$ is a lower bound on $Q(x)$ for values of $\beta > 1$ and $0 < \alpha \leq \sqrt{\frac{2e}{\pi}} \frac{\sqrt{\beta-1}}{\beta}$. We set $\beta = 2$ and $\alpha = \frac{1}{2} \sqrt{\frac{2e}{\pi}}$ to obtain the third inequality in (31). Hence, by definition 1, we have anti-concentration probability $p = \frac{e^{\frac{1}{2} - \frac{1}{\tau_{\min}^2}}}{4\sqrt{\pi}}$.

We now provide details on how to compute concentration constants c and c' in Equation (30). For any $\gamma > 0$, and using the complement rule and union bound, we have:

$$P \left(\|\eta\| \leq \gamma \sqrt{d} \right) \geq P \left(\forall 1 \leq i \leq d, |\eta_{\tau_i}| \leq \gamma \right) = 1 - P \left(\exists i \text{ s.t. } |\eta_{\tau_i}| > \gamma \right) \geq 1 - \sum_{i=1}^d P \left(|\eta_{\tau_i}| > \gamma \right). \quad (33)$$

Let $\tau_{\max} = \max_k \tau_k$ and $\eta_{\tau_{\max}} \sim \mathcal{N}(0, \tau_{\max}^2)$. We have,

$$P \left(\|\eta\| \leq \gamma \sqrt{d} \right) \geq 1 - dP \left(|\eta_{\tau_{\max}}| > \gamma \right) \geq 1 - 2de^{-\frac{\gamma^2}{2\tau_{\max}^2}}, \quad (34)$$

where the last inequality follows from applying the Chernoff bound (standard Gaussian concentration inequality). Using Equation (30), we substitute for $\delta = 2de^{-\frac{\gamma^2}{2\tau_{\max}^2}}$, which results in $\gamma = \sqrt{2\tau_{\max}^2 \log \frac{2d}{\delta}}$, $c = 2\tau_{\max}^2$, and $c' = 2$.