Industrial and
Systems Engineering

Information Collection Optimization in Designing Marketing Campaigns for Market Entry

Somayeh Moazeni[1], Boris Defourny[2], and Monika J. Wilczak[3]

[1]School of Business, Stevens Institute of Technology, USA

[2]Department of Industrial and Systems Engineering, Lehigh University, USA

[3]Accenture Digital, Accenture, USA

ISE Technical Report 17T-001

LEHIGH UNIVERSITY

Information Collection Optimization in Designing Marketing Campaigns for Market Entry

Somayeh Moazeni
School of Business, Stevens Institute of Technology, 1 Castle Point Terrace, Hoboken, NJ 07030, USA
smoazeni@stevens.edu

Boris Defourny
Department of Industrial and Systems Engineering, Lehigh University, 200 W Packer Ave, Bethlehem, PA 18015, USA
defourny@lehigh.edu

Monika J. Wilczak
Accenture, 1345 Avenue of the Americas, New York, New York 10105, USA
monika.j.wilczak@accenture.com

Developing marketing strategies for a new product or a new target population is challenging, due to the scarcity of relevant historical data. Building on dynamic Bayesian learning, a sequential information collection optimization assists in creating new data points, within a finite number of learning phases. This procedure identifies effective advertisement design elements as well as customer segments that maximize the expected outcome of the final marketing campaign. In this paper, the marketing campaign performance is modeled by a multiplicative advertising exposure model with Poisson jumps. The intensity of the Poisson process is a function of the marketing campaign features. A forward-looking measurement policy is formulated to maximize the expected improvement in the value of information in each learning phase. Solving this information collection optimization is reduced to a mixed-integer second-order cone programming problem. A computationally efficient approach is proposed that consists in solving a sequence of mixed-integer linear optimization problems. The performance of the optimal learning policy over commonly used benchmark policies is evaluated using examples from the property and casualty insurance industry.

Key words: adaptive experiments, optimal learning, dynamic optimization, marketing analytics

History: February 17, 2017

1. Introduction

One of the challenging marketing decisions is to design strategies for market entry, in which either the product or business line to be advertised is new or the target population differs from those in the previous marketing campaigns for the same product. In such marketing environments, the marketer faces significant uncertainty about the effectiveness of potentially considered marketing variables. Hence, in order to collect relevant information before committing to a final marketing campaign, the marketer may want to devote some time to try a few marketing strategies on a selected group of customers and to measure their effectiveness.

Since the created data points depend on the marketer's actions and the number of learning phases is finite, she faces a strategic choice in the selection of the sequence of marketing campaigns to be tried. As the number of considered marketing variables increases, there will be numerous variations of the marketing campaign to choose from. For example, after each measurement and updating the current state of knowledge, the marketer can select a marketing campaign which seems the most appealing, given her knowledge to date. This myopic learning does not take into account the effect of new information on future choices and the final estimated impact of marketing variables. Alternatively, the marketer can explore further by trying currently less-than-ideal marketing campaigns, in order to make a more informed decision in the future. Our goal in this paper is to investigate a principled process to guide forward-looking marketers through the design of a marketing campaigns sequence, in order to infer design features that exert significant impact on the outcome of a marketing campaign.

Marketing strategies will be ineffective when a "one size fits all" or even "one size fits many" approach is deployed in designing the campaign. We thus express a marketing campaign by features related to both the advertisement design and customer segments. Design inputs include advertisement content and channels or marketing mix variables defined by the marketer. Examples of customer segmentation features are the education level or past relationships of target customers with the company or industry. Various interactions among the advertisement design elements and customer segmentation attributes can be captured by constraints on the binary variables modeling these marketing features. A predictive model of the marketing campaign performance links these input features to a measurable output of concern, such as purchasing the product of interest, referring the industry to a friend, customer awareness levels, product perceptions, or levels of sales or profits (Lilien et al. 2012). We describe the marketing campaign outcome by a multiplicative advertising exposure model. Jumps here represent the exposure of the customers to the marketing campaign (Farris et al. 2015); for example, the customers' visits of a web advertisement, the number of passersby viewing a billboard, or the volume of customers who answer a telemarketing call. We let the arrival rate of this Poisson process depend on the marketing campaign features. The importance of models where the advertising exposure is characterized by the marketing campaign attributes are discussed in Barajas et al. (2016). The jump size explains the effect of input variables on the marketing campaign outcome per exposure. We define it by a noisy linear function of marketing campaign features.

Bayesian statistics can be employed to model evolving beliefs about the coefficients with uncertain expectations in the marketing campaign performance model. Bayesian models enable us to incorporate prior knowledge from similar marketing campaigns or expert opinions (Rossi and Allenby 2003). For the multiplicative advertising exposure model with Poisson jumps and under relatively general correlated belief assumptions, we establish closed-form equations for the Bayesian updating procedure.

Given a finite number of learning phases and a predictive model of the marketing campaign outcome, the sequential marketing campaign design problem falls into the class of finite-horizon discrete-time Markov Decision Processes (MDP), e.g., see (Puterman 1994). We formulate the information collection optimization as a dynamic programming problem. When the number of features and possible alternatives increases, the curse of dimensionality (Powell 2011, Bertsekas 2012) imposes challenges on exact computation of the value function and an optimal policy for this problem. In fact, the computational complexity of forward-looking learning has been one of the main reasons for the popularity of pure exploration and myopic learning approaches for market entry among marketing practitioners.

An approximation of this dynamic optimization maximizes the expected improvement in the value of information in each learning phase, where the value is quantified according to the one-step lookahead utility function. It was first introduced by (Gupta and Miescke 1996), and further studied as the expected improvement policy in (Chick 2006, Chick et al. 2010). This approximate dynamic programming policy has been revisited and analyzed as the knowledge-gradient (KG) policy by (Frazier et al. 2008) and later in the book (Ryzhov and Powell 2012b). The optimality and statistical consistency of this policy for linear predictor functions were investigated in (Frazier et al. 2009, Negoescu et al. 2011, Ryzhov et al. 2012, Ryzhov and Powell 2012a, Han et al. 2015), under different statistical assumptions for the coefficients. We develop this approximate policy for the multiplicative advertising exposure model of the marketing campaign outcome.

Enumeration of all feasible alternatives is a typical approach in the optimal learning literature to derive the KG decision. Efficient approaches for computing the KG policy are discussed. This step leverages large-scale convex optimization reformulations and solution techniques in mixed-integer second-order cone optimization. There are often numerous input features that a marketer would like to consider when designing marketing campaigns. This increases the computational time of solving the mixed-integer second-order cone optimization problem. We thus propose an iterative approach which consists in solving a number of mixed-integer linear programming problems. This algorithm often requires very few iterations and is significantly faster, as the number of features grows.

The information collection optimization methodology for market entry is illustrated by an application from property and casualty (P&C) insurance market. Relevant input features for P&C insurance products and their requirements are constructed based on the report from (McGranahan et al. 2013, Whelan and O'Neill 2014) and the industry experience of one of the authors. We elaborate on the formulation of various types of marketing variables and their constraints. We then compare the performance of the KG policy, adaptive myopic policy, Thompson sampling, and a number of other popular design policies (Bechhofer et al. 1995, Ryzhov and Powell 2012b). In a recent study for display advertising (Schwartz et al. 2016), the Thompson sampling policy outperforms a large number of other considered policies. Our simulations show that designing marketing campaigns using the KG policy consistently improves the marketing campaign outcome in comparison to other benchmark policies. This improvement is more prominent when the number of learning phases is small. Learning the effective marketing campaign features within fewer time steps and faster is a desirable property for market entry and new product development.

In summary, this paper contributes to the literature on marketing analytics and optimal learning. The contributions to the marketing literature include framing a multiplicative advertising exposure model for marketing campaign outcome in terms of both advertisement design features and customer segmentation variables, formalizing the sequential information collection optimization for market entry into a dynamic programming problem, and developing the Knowledge-Gradient policy to identify effective marketing campaign features. The present research also contributes to the optimal learning literature by extending the methodology for the predictive model with Poisson jumps whose arrival rate is non-constant and is a function of the input features, and proposing an efficient approach based on mixed-integer linear optimization to compute the KG policy.

This paper is organized as follows. Related literature is reviewed in Section 2. The model to predict the marketing campaign performance with the advertising exposure process is discussed in Section 3. The Bayesian inference method for model parameters is developed in Section 4. The sequential information collection optimization and its MDP formulation are described in Section 5. The concept of the expected value of information and the KG policy are presented in Section 6. Algorithms to compute the KG Policy are developed in Section 7. Illustrative examples from marketing campaigns in property and casualty insurance and computational results are explained in Sections 8 and 9. Conclusions are given in Section 10.

2. Related Literature

Sequential adaptive experiments have been practiced in interactive marketing and online advertising. However, the typical industry practice adopts marketing campaigns currently thought to be most promising and updates this belief as the outcome is observed. There are a number of efforts in this literature to develop alternative heuristic policies, such as Thompson sampling-based policies (e.g., see Agrawal and Goyal (2012), Chapelle and Li (2012), Russo and Roy (2014)) and Gittins index (Gittins et al. 2011) in online display advertising. Bertsimas and Mersereau (2007) study the problem of optimizing message selection from M given messages assigned to N customer encounters in a segment at each phase with the overall objective of maximizing total rewards over a finite time horizon. This assumes that customer encounters are statistically homogenous within a segment, the parameters to be learned are independent, and message effectiveness can be learnt independently across segments. Schwartz et al. (2016) study the advertiser's resource allocation optimization over time across many ad creatives and websites by sequentially learning about ad performance in order to maximize customer acquisition rates. They frame the problem of testing and learning for the online advertising problem as a multi-armed bandit, and investigate Thompson sampling as well as other multi-armed bandit policies by simulation and a field experiment. Both Bertsimas and Mersereau (2007) and Schwartz et al. (2016) select an ad from K pre designed ads. Our work differs in that we consider features to describe ad designs and target customer segments. This enables us to examine a greater variety of possible marketing campaigns and allows for dependence among advertisements and dependence among customer segments, which had been pointed out in Bertsimas and Mersereau (2007) as one of the challenges of their framework. In addition, as the main contribution of our paper, we investigate the KG policy for the marketing campaign design problem and develop an efficient approach to compute it even for a large number of possible alternatives. We then compare the performance of the KG policy with several other policies including the Thompson sampling policy.

The marketing literature includes studies of experiential learning dynamics from a customer's point of view, e.g., see (Ching et al. 2013) and the references therein, for a review on consumer learning models about product attributes over time. Forward-looking consumer experiential learning problems using Bayesian updating are investigated in (Lin et al. 2015), where a heuristic approach to balance exploitation and exploration is proposed. In contrast to this literature, the present paper aims to build on information collection techniques to guide a marketer through learning effective marketing campaign features.

Various marketing campaign performance metrics have been considered to approach strategic decision problems in marketing analytics, e.g., see (Hanssens et al. 2005, Bose and Chen 2009, Lilien et al. 2012, Farris et al. 2015). The outcome of a marketing action can be expressed by an explicit parametric function of the static variables of the campaign, such as the number of advertisements mailed. Examples include the growth curve model or the linear regression models in (Huxley 1980, Hill 1981, Bauer 1987, 1991, Basu et al. 1995, Fox et al. 1997), and the probabilistic response models in (Chun 2012, Chun and Jung 2015) for direct marketing. These models can be extended by including other advertisement design attributes, such as marketing channel features, e.g., see (Abe et al. 2004). Another class of models describes the marketing outcome only as a function of features related to the target customers, such as behavioral and demographic features; e.g., see (Coussement and Buckinx 2011, McCrary 2009, Moro et al. 2011, 2014) for telemarketing and direct marketing. These models are often used to identify an optimal customer segment for a given marketing communication. To accommodate the interdependency between marketing design features and customer segmentation features, similar to (Lu and Boutilier 2014), we consider both sets of features simultaneously. In addition, the multiplicative marketing campaign outcome model in the present paper explicitly treats advertising exposures through the Poisson random variable.

3. The Model

We express a marketing campaign by two classes of features: variables specifying customer segments such as demographic attributes or the state of previous promotions offered to a customer, and features related to the advertisement design including the advertisement content or the level of price reduction in a promotion[1]. These marketing variables are modeled by binary features $x_i \in \{0,1\}$.

Various requirements can be imposed on the features, either by definition or by the marketer's subjective opinion to address mutually exclusive features, conditional features, combined effect of multiple features, disjunctive features, and multiple-choice decisions. These constraints on binary variables can be formulated by linear constraints. For example, $x_j \leq x_i$ requires that feature j is conditional on feature i, the requirement $x_i + x_j \leq 1$ ensures that features i and j are mutually exclusive, and $x_i \geq x_j$ implies that feature i is a co-requisite for j. To exclude the simultaneous occurrence of the two possibilities i and j, e.g., two story types cannot be included simultaneously, we include the constraint $x_i + x_j = 1$. When the combined effect of two features x_i and x_j is different than the sum of their individual effects, the interaction between these attributes is captured by including the additional binary variable x_k and the respective requirement $x_k = x_i x_j$. This nonlinear constraint can be linearized using the three inequalities $x_k \leq x_i$, $x_k \leq x_j$, $x_i + x_j - 1 \leq x_k$. Hence, the set of feasible marketing campaigns, represented through m binary features, is explained by

$$\mathcal{X} \stackrel{\text{def}}{=} \{x \in \{0,1\}^m \text{ such that } Ax = h, \quad Bx \leq b\}, \tag{1}$$

for some matrices A and B, and vectors h and b. We assume that the set of constraints on the allowable inputs contains at least one equality constraint; for example, the linear system $Ax = h$ includes the trivial constraint $x_1 = 1$. The set $\mathcal{X}$ is always bounded; denote $K \stackrel{\text{def}}{=} |\mathcal{X}| < \infty$, which often depends exponentially on m. In marketing and advertisement design practices m is relatively large.

We distinguish between the features whose expected impacts on the marketing campaign outcome are uncertain to the marketer and those with known expected effects. Let the features $x_Z \in \{0,1\}^{r_0}$ refer to the marketing campaign attributes with known average impacts, and $x_B \in \{0,1\}^r$ denote

[1] The analyses in this paper can also accommodate the inclusion of other types of features, such as those related to the marketing campaign goal, e.g., to attract new customers or to encourage customers to be ambassadors of the brand.

those input features whose expected effects on the campaign outcome are uncertain. Note that $r + r_0 = m$.

For a marketing campaign characterized by the attributes $x = (x_Z, x_B) \in \mathcal{X}$, its performance is evaluated by the following model

$$\eta = \kappa J\left(\zeta, \beta, x_Z, x_B, \epsilon\right), \tag{2}$$

where η is a measurable marketing campaign outcome of interest. Here,

$\kappa \sim \text{Poisson}(\lambda_x)$: number of customer exposure events,

ζ: the effect per advertising exposure of features x_Z on the marketing campaign outcome, with the known average impacts $\mu_Z \stackrel{\text{def}}{=} \mathbb{E}[\zeta]$,

β: the effect per advertising exposure of features x_B on the marketing campaign outcome, with an uncertain expected effect $\mu_B \stackrel{\text{def}}{=} \mathbb{E}[\beta]$,

$\epsilon \sim \mathcal{N}\left(0, \frac{1}{\rho}\right)$: zero-mean noise term with $\text{var}[\epsilon] = \frac{1}{\rho}$, where the precision parameter ρ is uncertain,

and J denotes a deterministic function of the input features (x_Z, x_B), model parameters ζ and β, and the random term ϵ.

The marketing campaign outcome, $\eta \in \mathbb{R}$, can represent the customer spend or customer sales greater than a given level during the promotional period or observation time window, the proportion of advertisements that elicit purchases or other desired responses (McCrary 2009). The advertising exposure κ expresses the number of times that a target customer segment is exposed to the marketing campaign. For example, in direct mail marketing a sent mail may not even be received or opened by the customer; whence it elicits no response. In this case, κ is the number of times that the target population would actually receive and read the mails. Similarly, in online advertising campaigns, the exposure refers to the number of times that a target customer visits the website and views the advertisement. For further discussion about the concept of this marketing metric, the reader is referred to (Farris et al. 2015), where no specific modeling approach to capture this metric when quantifying the marketing campaign outcome is suggested.

We model the number of exposure events κ by a Poisson random variable with the arrival rate λ_x. We let the intensity of customers' advertising exposure during the time period of observation vary with the marketing campaign attributes:

$$\lambda_x = \lambda^{\mathscr{M}}(x), \tag{3}$$

for some deterministic function $\lambda^{\mathscr{M}}$. In practice, the level of exposure to the advertisement depends only on a few features, for example, only the advertisement channel. The Poisson random variable κ is assumed to be independent of ζ, β, and ϵ.

We note that the marketing campaign outcome can be recast in terms of the reward per impression or exposure, $\hat{\eta}$, where

$$\hat{\eta} = \begin{cases} J(\zeta, \beta, x_Z, x_B, \epsilon) & \text{with probability } p(x), \\ 0 & \text{with probability } 1 - p(x). \end{cases}$$

For η as in equation (2) and $\kappa \sim \text{Poisson}(\lambda_x)$, we have $\hat{\eta} = \frac{\eta}{\kappa}$ and $p(x) = \Pr(\kappa > 0) = 1 - e^{-\lambda_x}$. For the rest of this paper, we use the expression in model (2).

When enough relevant historical data about past marketing campaigns $x_1, \cdots, x_n$ and their corresponding outcomes $\eta_1, \cdots, \eta_n$ is available, the marketer can use this data set to estimate model parameters $\mu_Z \in \mathbb{R}^{r_0}$ and $\mu_B \in \mathbb{R}^r$. She would then choose marketing variables that maximize the expected campaign outcome. This involves solving the following constrained integer optimization problem with m binary decision variables, given the parameter μ_B:

$$\begin{aligned} U(\mu_B) \stackrel{\text{def}}{=} \max_{x \in \mathcal{X}} \mathbb{E}[\eta] &= \max_{x \in \mathcal{X}} \mathbb{E}\left[\kappa J(\zeta, \beta, x_Z, x_B, \epsilon) \,|\mu_B\right] \\ &= \max_{x \in \mathcal{X}} \lambda^{\mathscr{M}}(x) \mathbb{E}\left[J(\zeta, \beta, x_Z, x_B, \epsilon) \,|\mu_B\right]. \end{aligned} \tag{4}$$

This problem can be solved efficiently for a wide range of functions J. For example, when functions J and $\lambda^{\mathscr{M}}$ are linear in x, then problem (4) is reduced to a linearly constrained integer quadratic optimization.

In designing marketing campaigns for market entry, however, such a historical data set is either unavailable or insufficient. This imposes considerable uncertainty in the estimation of the model coefficients μ_B. In such circumstances, market research and testing alternatives can be employed, e.g., see Chapter 12 of Farris et al. (2015), to sequentially learn the coefficient values. This process is formally discussed next.

4. Performance Parameter Inference

An approach for modeling uncertainty in effectiveness of marketing campaign features is Bayesian learning. The Bayesian approach for making inference can accommodate multiple levels of randomness and correlation through prior distributions for model parameters. In addition, existing knowledge and managerial inputs can be incorporated in prior probability distributions. We adopt Bayesian modeling for expected impact of marketing campaign features and the variance of the noise term.

The precision ρ is expressed by a Gamma distribution, with the shape parameter a_0 and the rate parameter b_0:

$$\rho \sim Gamma(a_0, b_0). \tag{5}$$

The Bayesian priors on coefficients μ_B are assumed to follow a multivariate normal with the mean vector θ_0 and the covariance matrix $\frac{1}{\rho}\Sigma_0$:

$$\mu_B|\rho \sim \mathcal{N}\left(\theta_0, \frac{1}{\rho}\Sigma_0\right). \tag{6}$$

The multivariate model captures dependencies among unknown parameters. This implies that learning about one marketing campaign may inform the marketer about the effectiveness of several other campaigns. Capturing these statistical dependencies improves the learning speed.

We index learning phases by n, where the first decision is made at step $n = 0$. Using the distributional assumptions (5) and (6), the belief model after the n^{th} learning phase is described by

$$S_n \stackrel{\text{def}}{=} (\theta_n, \Sigma_n, a_n, b_n). \tag{7}$$

Denote the space of possible belief parameters by $\mathcal{S}$. The initial belief $S_0 = (\theta_0, \Sigma_0, a_0, b_0)$ is prespecified by the marketer. Subsequently, S_n and s refer to a random variable and a fixed point in the state space, respectively.

As new information is collected, the parameters of the belief distributions are updated, leading to improved estimations of the expected effect of features on the marketing campaign outcome. The updating procedure involves obtaining a posterior predictive distribution conditioned on S_{n+1} from the prior predictive distribution of μ_B and ρ, conditioned on S_n and the observation $(\kappa_{n+1}, \eta_{n+1})$. The graphical model (Wainwright et al. 2008) of the multiplicative advertising exposure model is illustrated in Figure 1.

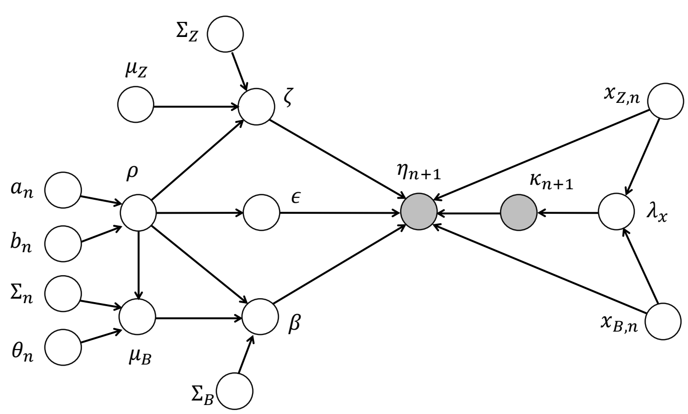

Figure 1 Graphical model for the multiplicative advertising exposure model

The following proposition shows that for the multiplicative advertising exposure model (2), the conditional distribution of (μ_B, ρ) given S_n remains multivariate normal-gamma for all n. While the result remains valid for broader classes of J for the marketing campaign outcome per exposure, here, we consider the linear model

$$J(\zeta, \beta, x_Z, x_B, \epsilon) = \zeta^\top x_Z + \beta^\top x_B + \epsilon. \tag{8}$$

The following proposition establishes probabilistic inference for the belief parameters.

PROPOSITION 1. *Consider the marketing campaign outcome model* (2)*, where*

$$\zeta \sim \mathcal{N}\left(\mu_Z, \frac{1}{\rho}\Sigma_Z\right), \qquad \beta \sim \mathcal{N}\left(\mu_B, \frac{1}{\rho}\Sigma_B\right).$$

Suppose that the conditional distribution of (μ_B, ρ) *given* S_n *is multivariate normal-gamma with parameters* $S_n = (\theta_n, \Sigma_n, a_n, b_n)$*. Then, the conditional distribution of* (μ_B, ρ) *given* $\{S_n, x_n, \kappa_{n+1}, \eta_{n+1}\}$ *with* $x_n = (x_{Z,n}, x_{B,n})$ *and* $\kappa_{n+1} > 0$ *is multivariate normal-gamma with parameters:*

$$\theta_{n+1} = \theta_n + \frac{\eta_{n+1} - \kappa_{n+1}(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n})}{\kappa_{n+1}(1 + x_{B,n}^\top \Sigma_n x_{B,n} + x_{Z,n}^\top \Sigma_Z x_{Z,n} + x_{B,n}^\top \Sigma_B x_{B,n})} \Sigma_n x_{B,n} \tag{9}$$

$$\Sigma_{n+1} = \Sigma_n - \frac{\Sigma_n x_{B,n} x_{B,n}{}^\top \Sigma_n}{1 + x_{B,n}{}^\top \Sigma_n x_{B,n} + x_{Z,n}^\top \Sigma_Z x_{Z,n} + x_{B,n}^\top \Sigma_B x_{B,n}}, \tag{10}$$

$$a_{n+1} = a_n + \frac{1}{2}, \tag{11}$$

$$b_{n+1} = b_n + \frac{\left(\eta_{n+1} - \kappa_{n+1}(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n})\right)^2}{2\kappa_{n+1}^2 \left(1 + x_{B,n}^\top \Sigma_n x_{B,n} + x_{Z,n}^\top \Sigma_Z x_{Z,n} + x_{B,n}^\top \Sigma_B x_{B,n}\right)}. \tag{12}$$

A proof of Proposition 1 is provided in Appendix A.

When $\kappa_{n+1} = 0$, model (2) yields $\eta_{n+1} = 0$. Hence, the conditional distribution of (μ_B, ρ) given $\{S_n, x_n, \kappa_{n+1}, \eta_{n+1}\}$ remains multivariate normal-gamma with parameters in S_n, i.e., the parameters $\theta_n, \Sigma_n, a_n, b_n$ will not be updated.

Equation (10) and the Courant-Fischer theorem (Theorem 8.1.5 in Golub and Loan (1996)) yield

$$\|\Sigma_{n+1}\|_2 = \lambda_{\max}\left(\Sigma_n - \frac{\Sigma_n x_{B,n} x_{B,n}{}^\top \Sigma_n}{1 + x_{B,n}{}^\top \Sigma_n x_{B,n} + x_{B,n}{}^\top \Sigma_B x_{B,n} + x_{Z,n}^\top \Sigma_Z x_{Z,n}}\right) \le \lambda_{\max}(\Sigma_n) = \|\Sigma_n\|_2,$$

where the inequality comes from $\lambda_{\min}(\Sigma_n x_{B,n} x_{B,n}{}^\top \Sigma_n) = \lambda_{\min}((\Sigma_n x_{B,n})(\Sigma_n x_{B,n})^\top) \ge 0$. Thus, the norm of the covariance matrix Σ_n decreases at each phase of the Bayesian inference. Minimizing the norm of the covariance matrix, i.e., the maximum eigenvalue of Σ_{n+1}, is the goal of E-optimal experiment design, see e.g. (Boyd 2004). Similarly, the trace of the covariance matrix $\mathrm{tr}(\Sigma_n)$ decreases by the updating procedure (10), i.e.,

$$\begin{aligned}\mathrm{tr}(\Sigma_{n+1}) &= \mathrm{tr}(\Sigma_n) - \frac{1}{1 + x_{B,n}{}^\top \Sigma_n x_{B,n} + x_{B,n}{}^\top \Sigma_B x_{B,n} + x_{Z,n}^\top \Sigma_Z x_{Z,n}} \mathrm{tr}(\Sigma_n x_{B,n} x_{B,n}{}^\top \Sigma_n) \\ &= \mathrm{tr}(\Sigma_n) - \frac{1}{1 + x_{B,n}{}^\top \Sigma_n x_{B,n} + x_{B,n}{}^\top \Sigma_B x_{B,n} + x_{Z,n}^\top \Sigma_Z x_{Z,n}} x_{B,n}{}^\top \Sigma_n^2 x_{B,n}. \end{aligned} \tag{13}$$

An A-optimal experiment design seeks to find a marketing campaign x_n which minimizes $\mathrm{tr}(\Sigma_{n+1})$.

Collected observations on the advertising exposure κ can also improve the estimation of the Poisson intensity rate $\lambda^{\mathscr{M}}(x)$. The following proposition presents the Bayesian updating procedure for a special class of point-wise Poisson intensity functions.

PROPOSITION 2. *Consider the point-wise intensity function $\lambda^{\mathscr{M}}(x_k) = \lambda_k$, for $k = 1, \cdots, K$. Suppose that for every $x_k \in \mathcal{X}$, the conditional distribution of $\lambda^{\mathscr{M}}(x_k)$ given S_n is Gamma with the shape parameter $c_n(x_k) \in \mathbb{R}_+$ and rate parameter $d_n(x_k) \in \mathbb{R}_+$. Then, by implementing the marketing campaign $x_n = (x_{Z,\bar{k}}, x_{B,\bar{k}})$, the conditional distribution of $\lambda_{\bar{k}}$ given $S_{n+1} = \{S_n, x_n, \kappa_{n+1}, \eta_{n+1}\}$ is Gamma with parameters:*

$$c_{n+1}(x_{\bar{k}}) = c_n(x_{\bar{k}}) + \kappa_{n+1}, \tag{14}$$

$$d_{n+1}(x_{\bar{k}}) = d_n(x_{\bar{k}}) + 1. \tag{15}$$

A proof of Proposition 2 is provided in Appendix A.

Bayesian inference for general Poisson intensity models is more involved and is not the focus of this paper. For further discussion, the reader is referred to (Kottas and Sanso 2007, Adams et al. 2009).

The following proposition characterizes the distribution of the marketing campaign outcome η_{n+1}.

PROPOSITION 3. *When η_{n+1} can be observed, i.e., $\kappa_{n+1} > 0$, the conditional distribution of η_{n+1} given S_n, $x_n = (x_{Z,n}, x_{B,n})$, and κ_{n+1} follows a univariate Student's t-distribution with $2a_n$ degrees of freedom, i.e.,*

$$\eta_{n+1}|S_n, x_n, \kappa_{n+1} = \kappa_{n+1}\left(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n}\right) + \kappa_{n+1}\sqrt{\frac{b_n}{a_n}\left(1 + x_{Z,n}{}^\top \Sigma_Z x_{Z,n} + x_{B,n}{}^\top (\Sigma_B + \Sigma_n) x_{B,n}\right)}\, T_{2a_n}, \tag{16}$$

where T_{2a_n} is a standard Student's t-distribution random variable with $2a_n$ degrees of freedom.

The proof of Proposition 3 is provided in Appendix A.

For $\kappa_{n+1} = 0$, the multiplicative advertising exposure model (2) yields $\Pr(\eta_{n+1} \mid S_n, x_n, \kappa_{n+1} = 0) = 1$ if $\eta_{n+1} = 0$, and $\Pr(\eta_{n+1} \mid S_n, x_n, \kappa_{n+1} = 0) = 0$ if $\eta_{n+1} > 0$.

Applying Proposition 3 in equation (9), the conditional distribution of θ_{n+1} can be expressed as

$$\theta_{n+1} = \theta_n + \sqrt{\frac{b_n}{a_n\left(1 + x_{Z,n}^\top \Sigma_Z x_{Z,n} + x_{B,n}{}^\top (\Sigma_B + \Sigma_n) x_{B,n}\right)}}\,\Sigma_n x_{B,n} T_{2a_n}. \tag{17}$$

Interestingly, the coefficient of the random variable T_{2a_n} is constant and does not depend on the actual value of κ_{n+1}. This representation of θ_{n+1} facilitates analyzing the information collection optimization problem and the computation of optimal policies, to be presented next.

5. The Sequential Information Collection Optimization

The Bayesian inference procedure, outlined in the previous section, assist the marketer to improve the estimation of the parameters in the predictive model for the marketing campaign outcome, before finalizing the best combination of features and committing to a large-scale marketing campaign. When the marketer is allotted to conduct only N sequential experimental marketing campaigns, where $N \ll K$, the marketer faces the strategic decision of selecting N measurements from the K alternatives.

Given the finite number of learning phases N, the sequential information collection optimization problem aims to design the sequence of N marketing campaigns in order to maximize the expected outcome of the final marketing campaign, where the expectation is computed based on the state of knowledge at the end of the N steps. More formally, the goal is to choose a policy π for the sequence of experimental campaigns $(x_0, \cdots, x_{N-1})$ and implementation campaign x_N that solves

$$\sup_{\pi \in \Pi} \mathbb{E}^{\pi} \left[U \left(\theta_N \right) \right], \tag{18}$$

where U is defined as in equation (4). Here, Π is the set of policies $\pi = (\pi_0(S_0), \cdots, \pi_{N-1}(S_{N-1}))$, where each decision rule π_k maps $S_k \in \mathcal{S}$ to an action $x_k \in \mathcal{X}$. The expectation $\mathbb{E}^{\pi}$ denotes the conditional expectation corresponding to the policy π.

Problem (18) can be framed as a discrete-time finite-horizon dynamic programming, in which the updating procedure induces the transition function $S^{\mathscr{M}}$ on $\mathcal{S}$, $S_{n+1} = S^{\mathscr{M}}(S_n, x_n, (\kappa_{n+1}, \eta_{n+1}))$. For $n = 0, \cdots, N$, the value of a policy $\pi \in \Pi$ at any state $s \in \mathcal{S}$ is defined as the reward-to-go from state s when $(N-n)$ experimental marketing campaigns are to be conducted under policy π. It is given by

$$V_n^{\pi}\left(S_n\right) \stackrel{\text{def}}{=} \mathbb{E}^{\pi}\left[U\left(\mathbb{E}[\mu_B | S_N]\right) | S_n\right] = \mathbb{E}^{\pi}\left[U\left(\theta_N\right) | S_n\right], \tag{19}$$

$$V_N^{\pi}\left(S_N\right) \stackrel{\text{def}}{=} U(\mathbb{E}[\mu_B | S_N]). \tag{20}$$

The value of the optimal policy π^*, denoted by $V_n(\cdot)$, is given by

$$V_n(S_n) = \max_{\pi \in \Pi} V_n^{\pi}(S_n), \qquad n = 0, \cdots, N.$$

The value function can be determined recursively by

$$V_n(S_n) = \max_{x_n \in \mathcal{X}} \mathbb{E}\left[V_{n+1}\left(S^{\mathscr{M}}\left(S_n, x_n, (\kappa_{n+1}, \eta_{n+1})\right)\right)\right]. \tag{21}$$

Define the Q-functions, $Q_n : \mathcal{S} \times \mathcal{X} \rightarrow \mathbb{R}$, as

$$Q_n(S_n, x_n) \stackrel{\text{def}}{=} \mathbb{E}\left[V_{n+1}\left(S^{\mathscr{M}}\left(S_n, x_n, (\kappa_{n+1}, \eta_{n+1})\right)\right)\right]. \tag{22}$$

Hence, the decisions on experimental marketing campaigns suggested by the optimal policy π^* satisfy

$$x_n^* \stackrel{\text{def}}{=} \pi_n^*(S_n) \in \arg\max_{x_n \in \mathcal{X}} \; Q_n(S_n, x_n). \tag{23}$$

From any knowledge state, exploring one more marketing campaign always improves the expected performance of the final (implementation) marketing campaign, i.e., the value of measuring the performance of the marketing campaign x at phase n, $Q_n(S_n, x)$, is no less than the value of making no measurement, $V_{n+1}(S_n)$. This is summarized in the following proposition. A similar result was first proved formally in Frazier et al. (2008) for independent normal populations and an identical known variance.

Proposition 4. *Consider the multiplicative advertising exposure model* (2)*, where J is as in* (8) *and $\lambda^{\mathscr{M}}(x)$ is a function of input features. For $n = 0, \cdots, N-1$ and for any state $S_n \in \mathcal{S}$, we have $Q_n(S_n, x) \geq V_{n+1}(S_n)$ for every $x \in \mathcal{X}$. In particular, $V_n(S_n) \geq V_{n+1}(S_n)$ for all possible states $S_n \in \mathcal{S}$, i.e., the value of the optimal policy decreases as the number of remaining learning phases reduces.*

A proof of Proposition 4 is provided in Appendix A.

The dynamic optimization principle in equation (21) characterizes an optimal policy. However, the value functions are often difficult to compute due to the high dimensionality of the state space. The curse of dimensionality of such decision making problems in marketing analytics with moderate or large numbers of marketing variables has often prevented marketers from adopting a systematic approach. In the following section, we investigate an approximate dynamic programming policy for the information collection problem (18).

6. Knowledge Gradient Policy

One of the approximate policies for the information collection optimization relies on the concept of expected improvement criterion (Jones et al. 1998) and the value of information measured by the expected single period reward $\mathbb{E}_n[V_N(S_{n+1}) - V_N(S_n)]$. This quantity is referred to as *Knowledge-Gradient (KG)* by (Frazier et al. 2008). For the marketing campaign design problem (18), it takes the following form: given $x \in \mathcal{X}$ and knowledge state S_n, at the n^{th} learning phase,

$$v_x^{\text{KG},n} \stackrel{\text{def}}{=} \mathbb{E}\left[V_N\left(S^{\mathscr{M}}(S_n, x, (\eta_{n+1}, \kappa_{n+1}))\right) - V_N(S_n) \mid S_n, x\right]. \tag{24}$$

The KG quantity $v_x^{\text{KG},n}$ represents the marginal value of a single measurement x. The KG policy, denoted by π^{KG}, is defined as the policy that maximizes the KG quantity. The decision rules of this policy, $x_n^{\text{KG}} : \mathcal{S} \to \mathcal{X}$, chooses the design of the $(n+1)^{th}$ marketing campaign to be

$$x_n^{\text{KG}} \stackrel{\text{def}}{=} \arg\max_{x \in \mathcal{X}} \; v_x^{\text{KG},n}. \tag{25}$$

If there are ties in the argmax in (25), the alternative with the smallest $\|x\|_1$ is chosen.

The KG policy can be viewed as a single step (Bayesian) look-ahead policy, e.g., see Powell (2011). For discussions on optimality of the KG policy, the reader is referred to (Frazier et al. 2008, 2009, Ryzhov and Powell 2012b). In particular, the KG policy would be optimal to implement, if this were the last experiment with no additional chances to collect information. It is also asymptotically optimal as the number of measurements N grows large.

The following proposition characterizes the KG quantity for the multiplicative advertising exposure model (2).

PROPOSITION 5. *Let the marketing campaign outcome η be as in* (2). *Then for any $x \in \mathcal{X}$ the KG quantity is given by*

$$v_x^{KG,n} = \left(1 - e^{-\lambda^{\mathscr{M}}(x)}\right)\left(\mathbb{E}\left[\max_{y\in\mathcal{X}} \ \left(p_y^n + q_y^n(x)T_{2a_n}\right) \ | \ S_n\right] - \max_{y\in\mathcal{X}} \ p_y^n\right), \tag{26}$$

where

$$p_y^n \stackrel{\text{def}}{=} \lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_n^\top y_B\right), \tag{27}$$

$$q_y^n(x) \stackrel{\text{def}}{=} \lambda^{\mathscr{M}}(y)\sqrt{\frac{b_n}{a_n\left(1 + x_Z^\top \Sigma_Z x_Z + x_B^\top(\Sigma_B + \Sigma_n)x_B\right)}}{x_B}^\top \Sigma_n y_B. \tag{28}$$

The proof of Proposition 5 can be found in Appendix A.

It follows from equation (26) that $v_x^{\text{KG},n}$ is always nonnegative. When the exposure rate to the marketing campaign is approximately zero, i.e., $\lambda^{\mathscr{M}}(x) \approx 0$, Proposition 5 indicates that the expected value of information gained from the marketing campaign x will be close to zero, i.e., $v_x^{\text{KG},n} \approx 0$.

In another special case, when the arrival rate does not depend on the measurement, i.e., $\lambda^{\mathscr{M}}(x) = \lambda$ for all $x \in \mathcal{X}$, the KG measurement can be computed by

$$\begin{aligned} x_n^{\text{KG}} &= \arg\max_{x\in\mathcal{X}} \ \left(1 - e^{-\lambda}\right)\left(\mathbb{E}\left[\max_{y\in\mathcal{X}}\left(p_y^n + q_y^n(x)T_{2a_n}\right) \ | \ S_n\right] - \max_{y\in\mathcal{X}} p_y^n\right) \\ &= \arg\max_{x\in\mathcal{X}} \ \mathbb{E}\left[\max_{y\in\mathcal{X}}\left(p_y^n + q_y^n(x)T_{2a_n}\right) \ | \ S_n\right]. \end{aligned} \tag{29}$$

Here, the second equality follows from the fact that the term $\max_{y\in\mathcal{X}} p_y^n$ does not depend on x.

7. Computation of the KG Policy

Given x and y, p_y^n and $q_y^n(x)$ are deterministic and T_{2a_n} is a one-dimensional random variable. Under this assumption Frazier et al. (2009) proposed an algorithm based on enumeration of elements in x to compute the KG quantity (26): Given S_n and x, sort the set $\{p_y^n, q_y^n(x)\}_{x\in\mathcal{X}}$ with respect to the values $q_y^n(x)$ in ascending order, and label the sorted set by $\{p_k^n, q_k^n\}_{k=1}^K$, i.e., $q_1^n \le q_2^n \le \cdots \le q_K^n$.

Let $\{c_k^n\}_{k=1}^{\bar{K}}$ be the non-decreasing sequence of breakpoints where the lines $p_k^n + q_k^n t$ intersect. Here, $\bar{K} \leq K$. Frazier et al. (2009) then prove that

$$\mathbb{E}\left[\max_{y\in\mathcal{X}}\left(p_y^n + q_y^n(x)T_{2a_n}\right) \mid S_n\right] - \max_{y\in\mathcal{X}} p_y^n = \sum_{k=1}^{\bar{K}}(q_{k+1}^n - q_k^n)\mathbb{E}\left[\max\{0, T_{2a_n} - |c_k^n|\}\right]. \qquad (30)$$

The expectation $\mathbb{E}[\max\{0, T_{2a_n} - |c_i^n|\}]$ can be expressed in terms of the pdf, $f_{T_{2a_n}}(\cdot)$, and cdf, $F_{T_{2a_n}}(\cdot)$, of the standard Student's t-distribution with $2a_n$ degrees of freedom, e.g., see Section 5.3 of (Ryzhov and Powell 2012b),

$$\mathbb{E}\left[\max\{0, T_{2a_n} - |c_k^n|\}\right] = \left(\frac{2a_n + (c_k^n)^2}{2a_n - 1} f_{T_{2a_n}}(|c_k^n|) - |c_k^n|(1 - F_{T_{2a_n}}(|c_k^n|))\right). \qquad (31)$$

Applying equations (30) and (31) in (26), we arrive at

$$v_x^{\mathrm{KG},n} = \left(1 - e^{-\lambda\mathscr{M}(x)}\right)\left(\sum_{k=1}^{\bar{K}}(q_{k+1}^n - q_k^n)\left(\frac{2a_n + (c_k^n)^2}{2a_n - 1} f_{T_{2a_n}}(|c_k^n|) - |c_k^n|(1 - F_{T_{2a_n}}(|c_k^n|))\right)\right). \qquad (32)$$

This expression indicates that when there is at least one $q_k^n < q_{k+1}^n$, we have $v_x^{\mathrm{KG},n} > 0$.

A KG decision x_n^{KG} can then be obtained by evaluating the exact KG quantity $v_x^{\mathrm{KG},n}$ at every feasible alternative $x \in \mathcal{X}$ using equation (32) within the optimization problem (25). This approach is typical method to compute the KG measurement in the optimal learning literature.

The above procedure involves identifying feasible alternatives in $\mathcal{X}$, computing p_y^n and $q_y^n(x)$ for every admissible measurement $y, x \in \mathcal{X}$, sorting the slopes, computing the breakpoints $\{c_k\}_k$, and finally selecting the best KG quantity $v_x^{\mathrm{KG},n}$. Thus, it remains computationally tractable only when identifying feasible alternatives in $\mathcal{X}$ is easy and $|\mathcal{X}| = K$ is relatively small. In marketing analytics, the number of advertisement design features are often large. As K increases, enumerating elements and assessing the associated $v_x^{KG,n}$ render computing the KG policy expensive. This intractability presents practical challenges in adopting KG policies in interactive marketing and online advertising.

Recently, Defourny et al. (2015) proposed to first apply an optimal quantization (e.g., see Pages and Wilbertz (2012), Graf and Luschgy (2000)) for the probability distribution of η, followed by developing a convex optimization formulation to maximize the KG quantity. This method was further investigated in Han et al. (2015). We adopt this algorithmic strategy for the multiplicative advertising exposure model (2).

Using the optimal quantization, the maximization problem in (25) is approximated by

$$\max_{x\in\mathcal{X}} \left(1 - e^{-\lambda\mathscr{M}(x)}\right)\left(\sum_{j=1}^{J} w_j \max_{y\in\mathcal{X}}\left(p_y^n + q_y^n(x)t_j\right) - \max_{y\in\mathcal{X}} p_y^n\right), \qquad (33)$$

where $\{t_j\}_{j=1}^J \subseteq \mathbb{R}$ is the sequence of points that minimizes the quadratic quantization error of the Voronoi quantizer for T_{2a_n}, $t_0 = -\infty$, $t_{J+1} = \infty$, and $w_j \stackrel{\text{def}}{=} F_{T_{2a_n}}\left(\frac{t_j+t_{j+1}}{2}\right) - F_{T_{2a_n}}\left(\frac{t_{j-1}+t_j}{2}\right)$, for $j = 1, \cdots, J$ with $\sum_{j=1}^J w_j = 1$. Next, we present a second-order cone reformulation of problem (33). In the subsequent discussion, we denote

$$P_n \stackrel{\text{def}}{=} \frac{a_n}{b_n}\left(\begin{matrix} A^\top A \\ h^\top h \end{matrix} + \begin{pmatrix} \Sigma_Z & 0 \\ 0 & \Sigma_n + \Sigma_B \end{pmatrix}\right), \tag{34}$$

$$\nu \stackrel{\text{def}}{=} \max_{y \in \mathcal{X}} \; p_y^n, \tag{35}$$

$$\mathcal{X}^+ \stackrel{\text{def}}{=} \{(\hat{x}, k) \in \{0, k\}^m \times \mathbb{R} \quad \text{such that} \quad k \geq 0, \; A\hat{x} = hk, \quad B\hat{x} \leq gk\}. \tag{36}$$

In addition, $\mathbf{1}_m$ denotes the vector of all ones of dimension $m \times 1$.

Proposition 6. *Problem (33), with p_y^n and $q_y^n(x)$ as in equations (27) and (28), is equivalent (in the sense of optimal objective value and projected solution) to the following mixed-integer optimization problem:*

$$\begin{aligned} \max_{\substack{(\hat{x},k)\in\mathcal{X}^+,\; k\leq M,\; x\in\mathcal{X} \\ y^j\in\mathcal{X},\quad j=1,\cdots,J}} \quad & \left\{\left(1 - e^{-\lambda^{\mathscr{M}}(x)}\right)\left(\sum_{j=1}^J w_j \lambda^{\mathscr{M}}\left(y^j\right)\left(\mu_Z^\top y_Z^j + \theta_n^\top y_B^j\right) + k - \nu\right)\right\} \\ \text{s.t.} \quad & k \cdot \boldsymbol{1}_m - M(\boldsymbol{1}_m - x) \leq \hat{x} \leq Mx, \\ & \|P_n^{1/2}\hat{x}\|_2 \leq \sum_{j=1}^J w_j t_j \lambda^{\mathscr{M}}(y^j) x_B^\top \Sigma_n y_B^j. \end{aligned} \tag{37}$$

The proof of Proposition 6 relies on the application of the lifting technique and the assumption $\|h\|_2 > 0$. For the details the reader is referred to (Defourny et al. 2015). In problem (37), M is a large constant. A choice for this parameter is

$$M = \frac{\lambda_{\max}(\Sigma_n)}{\sqrt{\lambda_{\min}(P_n)}}\left(\max_{y^j\in\mathcal{X}} \sum_{j=1}^J w_j \|t_j \lambda^{\mathscr{M}}(y^j) y_B^j\|_2\right).$$

Suppose $\{\lambda^{\mathscr{M}}(x)\}_{x\in\mathcal{X}} = \{\lambda_1, \cdots, \lambda_L\}$. Then problem (37) can be solved by $\max\{\mathcal{V}(\lambda_1), \cdots, \mathcal{V}(\lambda_L)\}$, where

$$\begin{aligned} \mathcal{V}(\lambda_\ell) \stackrel{\text{def}}{=} (1 - e^{-\lambda_\ell}) \max_{\substack{(\hat{x},k)\in\mathcal{X}^+,\; k\leq M,\; x\in\mathcal{X} \\ y^1\in\mathcal{X},\cdots,y^J\in\mathcal{X}}} \quad & \left\{\sum_{j=1}^J w_j \lambda^{\mathscr{M}}(y^j)(\mu_Z^\top y_Z^j + \theta_n^\top y_B^j) + k - \nu\right\} \\ \text{s.t.} \quad & \lambda^{\mathscr{M}}(x) = \lambda_\ell, \\ & k \cdot \mathbf{1}_m - M(\mathbf{1}_m - x) \leq \hat{x} \leq Mx, \\ & \|P_n^{1/2}\hat{x}\|_2 \leq \sum_{j=1}^J w_j t_j \lambda^{\mathscr{M}}(y^j) x_B^\top \Sigma_n y_B^j. \end{aligned} \tag{38}$$

The subproblems can be solved in parallel or, when L is not large, by maximizing $\sum_{\ell=1}^L l_\ell \mathcal{V}(\lambda_\ell)$, using binary decision variables $l_\ell \in \{0, 1\}$, where $\sum_{\ell=1}^L l_\ell = 1$.

When the advertising exposure rate $\lambda^{\mathscr{M}}(x)$ is a linear function of input features, the subproblems (38) are reduced to mixed-integer second-order cone problems. This is formally stated in the following proposition. Below, vec$(\cdot)$ is the vectorization operator, and $(\mu;\theta)$ denotes the column vector by stacking the vector μ on top of θ.

PROPOSITION 7. *Let the marketing campaign exposure rate be linear, i.e.,* $\lambda^{\mathscr{M}}(x) = \bar{\lambda}^\top x$ *for some* $\bar{\lambda} \in \mathbb{R}^m$*. Then, subproblems in* (38) *can be expressed as mixed-integer second-order cone optimization problems:*

$$
\begin{aligned}
\mathcal{V}(\lambda_\ell) = (1-e^{-\lambda_\ell})\ \max\ & \sum_{j=1}^{J} w_j vec\left((\mu_Z;\theta_n)\bar{\lambda}^\top\right)^\top vec(u^j) + k - \nu \qquad (39)\\
s.t.\ \ & \bar{\lambda}^\top x = \lambda_\ell,\\
& \|P_n^{1/2}\hat{x}\|_2 \le \sum_{j=1}^{J} w_j t_j vec\left(vec(\Sigma_n)\bar{\lambda}^\top\right)^\top vec(z^j)\\
& u^j_{i,i'} \le y^j_i, \quad j=1,\cdots,J,\ i,i'=1,\cdots,m,\\
& u^j_{i,i'} \le y^j_{i'}, \quad j=1,\cdots,J,\ i,i'=1,\cdots,m,\\
& u^j_{i,i'} \ge y^j_i + y^j_{i'} - 1, \quad j=1,\cdots,J,\ i,i'=1,\cdots,m,\\
& z^j_{(i+r(i'-1)),i''} \le (y^j_B)_i, \quad i,i'=1,\cdots,r,\ i''=1,\cdots,m,\\
& z^j_{(i+r(i'-1)),i''} \le (x_B)_{i'}, \quad i,i'=1,\cdots,r,\ i''=1,\cdots,m,\\
& z^j_{(i+r(i'-1)),i''} \le y^j_{i''}, \quad i,i'=1,\cdots,r,\ i''=1,\cdots,m,\\
& z^j_{(i+r(i'-1)),i''} \ge (y^j_B)_i + (x_B)_{i'} + (y^j_B)_{i''} - 2,\ i,i'=1,\cdots,r,\ i''=1,\cdots,m,\\
& k\cdot \boldsymbol{1}_m - M(\boldsymbol{1}_m - x) \le \hat{x} \le Mx, \quad \hat{x} \le k\cdot\boldsymbol{1}_m, \quad k \le M,\\
& (\hat{x},k)\in\mathcal{X}^+,\ x\in\mathcal{X},\ y^1,\cdots,y^J\in\mathcal{X},\quad u^j\in\{0,1\}^{m\times m},\ z^j\in\{0,1\}^{r^2\times m}.
\end{aligned}
$$

Here, ν *is given in* (35) *and can be computed by the following integer linear optimization problem,*

$$
\begin{aligned}
\nu = \max_{x\in\mathcal{X},\ u\in\{0,1\}^{m\times m}}\ & vec\left((\mu_Z;\theta_n)\bar{\lambda}^\top\right)^\top vec(u), \qquad (40)\\
s.t.\ \ & u_{i,i'} \le x_i, \quad i,i'=1,\cdots,m,\\
& u_{i,i'} \le x_{i'}, \quad i,i'=1,\cdots,m,\\
& u_{i,i'} \ge x_i + x_{i'} - 1, \quad i,i'=1,\cdots,m.
\end{aligned}
$$

The proof of Proposition 7 is presented in Appendix A.

Problem (39) can be solved using optimization solvers such as those by Gurobi Optimization (2015) or MOSEK ApS (2015). We note that when the strong duality holds for the integer optimization problem in (35), it may be computationally attractive to specify the value of ν using a set of constraints within problem (39) and avoid solving problem (40) separately.

The computational time of problem (39) grows rapidly as the number of input features m increases. We propose an iterative procedure to approximately solve this problem using a sequence of mixed-integer linear optimization problems. This algorithm is motivated in the next proposition. For simplicity in presentation, we assume that $\lambda^{\mathscr{M}}(x^j) = \lambda$. However, a similar treatment as in proposition 7 can be applied for more general functions, such as polynomial functions.

Proposition 8. *Suppose $\lambda^{\mathscr{M}}(x_j) = \lambda$. Consider the mixed-integer linear optimization problem $\mathcal{P}_\alpha$:*

$$(\mathcal{P}_\alpha) \qquad \max \left\{ \sum_{j=1}^{J} \lambda w_j \left(\mu_Z^\top y_Z^j + \theta_n^\top y_B^j \right) + \alpha \sum_{j=1}^{J} \lambda w_j t_j \, vec(\Sigma_n)^\top vec\left(z^j\right) - \nu \right\}, \tag{41}$$

$$s.t \quad vec(P_n)^\top vec(u) = 1, \tag{42}$$

$$0 \le u_{ii'} \le M x_i, \qquad i = 1, \cdots, m, \tag{43}$$

$$0 \le u_{ii'} \le M \hat{x}_{i'}, \qquad i' = 1, \cdots, m, \tag{44}$$

$$\hat{x}_{i'} - M(1 - x_i) \le u_{ii'} \le \hat{x}_{i'}, \qquad i, i' = 1, \cdots, m, \tag{45}$$

$$(z_B^j)_{i,i'} \le (y_B^j)_i, \qquad i = 1, \cdots, r, \tag{46}$$

$$(z_B^j)_{i,i'} \le (x_B)_{i'}, \qquad i' = 1, \cdots, r, \tag{47}$$

$$(z_B^j)_{i,i'} \ge (y_B^j)_i + (x_B)_{i'} - 1, \qquad i, i' = 1, \cdots, r, \tag{48}$$

$$k \cdot \boldsymbol{1}_m - M(\boldsymbol{1}_m - x) \le \hat{x} \le M x, \quad \hat{x} \le k \boldsymbol{1}_m, \quad k \le M, \tag{49}$$

$$(\hat{x}, k) \in \mathcal{X}^+, \; x \in \mathcal{X}, \; y^1, \cdots, y^J \in \mathcal{X}, \; u \in \mathbb{R}^{m \times m}, \; z^1, \cdots, z^J \in \mathbb{R}_+^{r \times r}.$$

Let k^* and x^* be part of an optimal solution of problem $\mathcal{P}_\alpha$. Then $x_n^{KG} = x^*$, when $k^* \approx \alpha^2$.

Proposition 8 is proved in Appendix A.

Problem $\mathcal{P}_\alpha$ is used in the algorithm described in Figure 2 to obtain a value for α, which leads to $k^* \approx \alpha^2$, and consequently an approximate solution for x_n^{KG}. In our computational investigation, we

Require: γ^{tol}, $N^{\max}$, α_0.
$N = 0$
$\alpha = \alpha_0$
$\gamma = 2\gamma^{tol}$
while $\gamma > \gamma^{tol}$ **and** $\bar{N} \le N^{\max}$ **do**
Solve problem $\mathcal{P}_\alpha$ **and compute** k^*
$\gamma = \lvert \alpha - k^* \rvert$
$\alpha = \sqrt{k^*}$
$\bar{N} = \bar{N} + 1$
end while
return x^*

Figure 2 An approximate algorithm using a sequence of mixed-integer linear programming problems for computing the KG policy

set $N^{\max} = 10$ and $\gamma^{tol} = 0.001$. However, as illustrated in Section 9, the iterative algorithm using MILPs tends to converge in very few (3 or 4) iterations. This procedure is significantly faster than the mixed-integer second-order optimization formulation for the large number of input features.

The computational methods in this section are employed in Section 9. We proceed by modeling marketing campaign features in property and casualty insurance and formulating typical requirements on these features.

8. Application to Property and Casualty Insurance

The insurance industry has been heavily engaged in advertising over the past few years, e.g., see (McGranahan et al. 2013). According to Whelan and O'Neill (2014), "*Insurance* was the most expensive keyword in Google AdWords in 2011. State Farm and Progressive each spent more than $43 million advertising on Google that year." Surveys on 26,500 U.S. property and casualty insurance consumers and 3,600 consumers of auto insurance are analyzed in Whelan and O'Neill (2014). From the empirical analysis of the survey data in this study, we inferred features that have been recognized to be effective. These features are summarized in Table 1. While their importance varies depending on the objective of the marketing campaign, e.g., customer acquisition, retention, cross-selling, or loyalty (Whelan and O'Neill 2014), we include them all as potential input features for insurance market entry, which is aimed primarily for acquisition.

These features fall into two categories: (i) customer segment features, which describe observable customer specifications such as demographic segmentation (e.g., age, education, gender, occupation, nationality, or household demographics such as marital status or home value), perceptual customer variables such as the customer's risk attitude or sensitivity to price, as well as customer-industry specifications including recorded attributes of the previous correspondence with the customer, relationship duration, shopping behavior, purchase history and frequencies. For further details on customer segmentation, see, e.g., (Wedel and Kamakura 2000, 2002, Webber 1998, Weinstein 2006, Foedermayr and Diamantopoulos 2008, Goyat 2011). (ii) advertisement design features, which constitute specific marketing strategies implemented in the campaign, such as the advertising channels, content, or variety of monetary and non-monetary premiums and incentives (Cobanoglu and Cobanoglu 2003). For general discussion on other marketing variables, key customer metrics, and their modeling challenges, see Gupta and Zeithaml (2006) and the references therein. For the benefits of advertisement design features related to personalized offers and messages in marketing, e.g., see (Schultz et al. 2009, Charski 2013, Nesamoney 2015).

In Table 1, the feature type *Insurance Needs* expresses the average number of insurance products held with main or other insurers. *Multi-channel* refers to agent and contact center, agent and digital, contact center and digital, agent and contact center and digital. *Service capabilities* emphasize

Customer Segment Features

Type	Categories
Age	18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older
Household Income	$<$ \$25K, \$25K$-$\$50K, $\cdots$, \$100K$-$\$150K, $>$ \$150
Insurance Needs	average number of insurance products: 1, 2, (3 or more)
Insurance Product	home, auto
Customer Mindset	peace-of-mind-focused, convenience-focused, price-focused
Time with an Insurer	< 1 year, $1-3$ years, $3-6$ years, > 6 years

Advertisement Design Features

Type	Categories
Advertisement Channel	digital, contact center, agent, multi-channel
Service Capabilities	emphasizing online options or easy service
Pricing Innovations	emphasizing accident forgiveness or telematics programs
Themes and Contents	informative, emotional

Table 1 Marketing campaign features in property and casualty insurance

options to purchase online without an agent, or easy service such as claim handling features, the existence of a mobile app, a video chat with an agent, video conferencing, etc. *Pricing innovations* accentuate price-related features such as first accident forgiveness, in-vehicle telematics programs, or route-based pricing. *Theme and Content* refer to different ways of composing stories and formulating the advertisement channel, or the written appeal in the advertisement. While this feature type is not explicitly considered in (Whelan and O'Neill 2014), we include it in our analysis motivated by other studies in the literature, e.g., see (Kunreuther et al. 2014), or (Klapdor et al. 2014) about the importance of content and words in marketing.

Categories of each feature type in Table 1 are modeled by a binary decision variable. For age, we introduce $x_1, \cdots, x_6$, where, $x_1 = 1$ if age is between 18 and 24, $x_2 = 1$ if age is between $25-34$, etc. The age-related features should satisfy the constraint $\sum_{i=1}^{6} x_i = 1$. Similarly, we define the income-related binary features $x_7, \cdots, x_{12}$, features related to the insurance needs x_{13}, x_{14}, x_{15}, decision variables x_{16} and x_{17} for Insurance Product, x_{18}, x_{19}, x_{20} for Customer Mindset, and $x_{21}, \cdots, x_{24}$ for Time with Insurer. Requirements on the customer segment features specify the following set

$$\mathcal{X}_{\mathrm{C}} = \left\{ x \in \{0,1\}^{37} \text{ s.t. } \sum_{i=1}^{6} x_i = 1, \sum_{i=7}^{12} x_i = 1, \sum_{i=13}^{15} x_i = 1, \sum_{i=16}^{17} x_i = 1, \sum_{i=18}^{20} x_i = 1, \sum_{i=21}^{24} x_i = 1 \right\}.$$

For Ad channels, we define variables x_{25} (agent), x_{26} (digital), x_{27} (contact center), x_{28} (agent and contact center), x_{29} (digital and contact center), x_{30} (digital and agent), x_{31} (agent and digital and contact center). For pricing innovation features, we introduce the binary variables x_{34} to specify that the advertisement emphasizes on first accident forgiveness, and x_{35}, if the advertisement emphasizes on in-vehicle telematics programs. The associated constraints define the following set:

$$\mathcal{X}_{\mathrm{A}} = \Big\{ x \in \{0,1\}^{37} \text{ s.t. } \; x_{28} = x_{25} \times x_{27},\; x_{29} = x_{26} \times x_{27},\; x_{30} = x_{25} \times x_{26},\; x_{31} = x_{25} \times x_{26} \times x_{27}, \\ \textstyle\sum_{i=25}^{27} x_i \geq 1,\; \sum_{i=32}^{33} x_i = 1,\; \sum_{i=34}^{35} x_i = 1,\; \sum_{i=36}^{37} x_i = 1. \Big\}.$$

The nonlinear constraints such as $x_{28} = x_{25} \times x_{27}$ in the definition of $\mathcal{X}_2$ above can be reformulated as linear constraints.

We also include a number of cross-type constraints to ensure that for price-focused customers the accident forgiveness pricing innovation property is emphasized, for customers who are with an insurer for more than three years the accident forgiveness is not mentioned, the easy service is emphasized for convenience-focused customers, while for customers older than 45 years options on online services and digital marketing channel are not highlighted, and finally the telematics programs are not emphasized for customers earning more than $100K$. These requirements imply the following set

$$\mathcal{X}_{\rm I} = \{x \in \{0,1\}^{37} \text{ s.t. } \quad x_{20} - x_{34} \le 0,\ x_\ell + x_{34} \le 1 \quad \ell = 23, 24,\ x_{19} - x_{33} \le 0,$$
$$x_i + x_{32} \le 1,\ x_i + x_{26} \le 1 \quad i = 4, 5, 6, \quad x_j + x_{35} \le 1 \quad j = 11, 12\}.$$

Hence, the set of feasible marketing campaigns $\mathcal{X} = \mathcal{X}_{\rm C} \cap \mathcal{X}_{\rm A} \cap \mathcal{X}_{\rm I}$ includes $K = |\mathcal{X}| = 34,560$ alternatives.

9. Computational Results

We use simulations to computationally study the performance of the KG policy and a number of other popular policies to identify effective features for property and casualty insurance marketing campaigns explained in Section 8.

The most popular policy among practitioners is the adaptive myopic policy, also called *pure exploitation* or *greedy heuristic*. It focuses on the alternatives that appear to be the best given our current estimates and implements $\text{argmax}_{x\in\mathcal{X}}\mathbb{E}[\eta|S_n]$ at phase n. The Thompson sampling policy, found outperforming in Schwartz et al. (2016), promotes more exploration by first drawing a single sample from n^{th} posterior distributions, $(\hat{\mu}_B, \hat{\rho}) \sim \text{Normal-Gamma}(\theta_n, \Sigma_n, a_n, b_n)$, and then, implementing the argmax of $\mathbb{E}[\eta|\hat{\mu}_B]$. Optimal design policies, also called *optimal design of experiments*, aim to choose $x \in \mathcal{X}$ which minimizes some scalarization of the error covariance matrix Σ_{n+1}, given the belief parameter Σ_n, e.g., see (Boyd 2004, Bechhofer et al. 1995). The *D-optimal design* minimizes the determinant of the error covariance matrix, i.e., $\log\det(\Sigma_{n+1})$, *E-optimal design* minimizes $\|\Sigma_{n+1}\|_2$, and *A-optimal design* minimizes the trace of the covariance matrix, i.e., $\min_{x\in\mathcal{X}} \text{tr}\,(\Sigma_{n+1})$, where $\text{tr}(\Sigma_{n+1})$ is computed as in equation (13). These problems can be cast as mixed integer second-order cone optimization, e.g., see Proposition 9 in Appendix A for E-optimal design policy. There are other strategies for balancing exploration and exploitation, such as Boltzmann exploration policy, uniform exploration policy, optimal computing budget allocation procedure, linear loss policy (e.g., see Frazier et al. (2008), Ryzhov and Powell (2012b), Chick and Frazier (2013)). Computing these policies often involves enumerating all alternatives; whence, the performance of these policies quickly degrades as K is large and the number of learning phases reduces.

For each policy, we simulate M replications, each of which involves generating a single *true* value of ρ from the prior distribution $Gamma(a_0, b_0)$, and the vector of *true* mean $\mu_B|\rho \sim \mathcal{N}(\theta_0, \frac{1}{\rho}\Sigma_0)$. These true values are not observed by any policy when making decisions and are used only for policy evaluations. We let $\mu_Z = 50\mathbf{1}_{r_0}$. The prior for ρ is set to be $a_0 = 1.5$ and $b_0 = 10$ in all experiments. In our simulations, we set the components of θ_0 by randomly drawing from a normal distribution. The prior (symmetric positive semidefinite) covariance is set to be $\Sigma_0 = (S + S^\top)^2$, for some matrix S with components $S_{ij} \sim \mathcal{N}(0,1)$. Similarly, we randomly generate covariance matrices $\bar{\Sigma}_0$ and $\bar{\Sigma}$ for Σ_Z and Σ_B. The time horizon in marketing is typically small. Hence, we focus on $N \leq 6$ in our experiments. The KG policy is computed with $J = 5$ in the quantization procedure. We assume that $\lambda^{\mathscr{M}}(x)$ only depends on the advertisement channel features:

$$\lambda^{\mathscr{M}}(x) = \bar{\lambda}_1 x_{25} + \bar{\lambda}_2 x_{26} + \bar{\lambda}_3 x_{27}, \tag{50}$$

where $\bar{\lambda}_1 = 10$, $\bar{\lambda}_2 = 15$, and $\bar{\lambda}_3 = 20$. Note that $m = 38$.

Each policy chooses a design x_n according to its decision rule and observes $(\kappa_{n+1}, \eta_{n+1})$. After N measurements, the performance of the policy π is evaluated by letting x^π be the solution that maximizes $\mathbb{E}[\eta|\theta_n]$, and computing the revenue R_N^π and the normalized regret $\varepsilon_N^{U,\pi}$, as follows

$$R_N^\pi \stackrel{\text{def}}{=} \lambda^{\mathscr{M}}(x^\pi)\left(\mu_Z^\top x_Z^\pi + \mu_B^\top x_B^\pi\right), \tag{51}$$

$$\varepsilon_N^\pi \stackrel{\text{def}}{=} \frac{1}{\bar{U} - \underline{U}}\left(\bar{U} - \lambda^{\mathscr{M}}(x^\pi)\left(\mu_Z^\top x_Z^\pi + \mu_B^\top x_B^\pi\right)\right), \tag{52}$$

where $\bar{U} \stackrel{\text{def}}{=} \max_{x\in\mathcal{X}} \lambda^{\mathscr{M}}(x)(\mu_Z^\top x_Z + \mu_B^\top x_B)$ and $\underline{U} \stackrel{\text{def}}{=} \min_{x\in\mathcal{X}} \lambda^{\mathscr{M}}(x)(\mu_Z^\top x_Z + \mu_B^\top x_B)$. Subsequently, we refer to $\bar{U}$, i.e., the optimal objective value given the true parameters μ_B and ρ, as the *ideal* policy.

Table 2 reports the average revenue of the policies for $N = 4, 5, 6$ and $M = 1000$ replications. In panel (a), where $r_0 = 13$ and $r = 20$, we have $K = 34,560$ marketing campaigns. In Tables 2(b) and 2(c), when $r_0 = 20$ and $r = 13$, we focus on a particular customer segment characterized by the first category of each feature type, i.e., $x_1 = 1$, $x_7 = 1$, $x_{13} = 1$, $x_{16} = 1$, $x_{18} = 1$, $x_{21} = 1$. In this scenario, the goal is to identify effective advertisement design elements for a fixed target customer segment. Hence, $K = 56$. In panel (c), $\Sigma_Z = \Sigma_B = 0$, i.e., the coefficients ζ and β are deterministic and equal $\zeta = \mu_Z$ and $\beta = \mu_B$. Table 2 indicates that all policies improve the expected outcome as N increases. The average revenue improvement (in percentage) panel in Table 2 shows that all policies significantly outperform the adaptive myopic policy. The left panel in Table 2 evinces that the average revenue of the *KG* policy is the most closest to that of the ideal policy in comparison with other policies including the Thompson sampling policy. This pattern holds in all three tables independent of the value of r_0 or covariance matrices. As Table 2(c) exhibits the revenue of the KG policy approaches to that of ideal when the variation in the performance model coefficients reduces.

	Average Revenue						**% Improvement**			
N	Ideal	Myopic	KG	Thompson Sampling	A-Design	E-Design	KG vs. Myopic	TS vs. Myopic	A-Design vs. Myopic	E-Design vs. Myopic
4	2628.21	1718.14	2202.84	1948.96	1943.59	2108.38	128.21	113.43	113.12	122.71
5	2628.21	1818.40	2248.63	2072.63	2065.92	2088.68	123.66	113.98	113.61	114.86
6	2628.21	1871.06	2316.47	2129.28	2182.05	2164.04	123.81	113.80	116.62	115.66

(a) $r_0 = 13$, $r = 20$, $\Sigma_Z = \bar{\Sigma}$, $\Sigma_B = \bar{\Sigma}_0$.

	Average Revenue						**% Improvement**			
N	Ideal	Myopic	KG	Thompson Sampling	A-Design	E-Design	KG vs. Myopic	TS vs. Myopic	A-Design vs. Myopic	E-Design vs. Myopic
4	2431.94	1475.76	1952.22	1704.67	1754.80	1894.59	132.29	115.51	118.91	128.38
5	2431.94	1529.45	2016.54	1784.38	1897.92	1934.81	131.85	116.67	124.09	126.50
6	2431.94	1561.59	2088.73	1868.37	1953.55	1958.28	133.76	119.65	125.10	125.40

(b) $r_0 = 20$, $r = 13$, $\Sigma_Z = \bar{\Sigma}$, $\Sigma_B = \bar{\Sigma}_0$.

	Average Revenue						**% Improvement**			
N	Ideal	Myopic	KG	Thompson Sampling	A-Design	E-Design	KG vs. Myopic	TS vs. Myopic	A-Design vs. Myopic	E-Design vs. Myopic
4	2536.41	1845.44	2359.34	2183.09	2100.19	2177.85	127.85	118.30	113.80	118.01
5	2536.41	1884.78	2478.83	2310.36	2280.58	2273.43	131.52	122.58	121.00	120.62
6	2536.41	1896.69	2501.83	2380.45	2306.52	2392.13	131.90	125.51	121.61	126.12

(c) $r_0 = 20$, $r = 13$, $\Sigma_Z = 0$, $\Sigma_B = 0$.

Table 2 Average revenue for different sequential learning policies

Figure 3 illustrates the average revenue R_N^π and the average regret ε_N^π as a function of learning phases for different marketing campaign design policies π. The solid horizontal lines in the left plots of Figure 3 indicates the average revenue of the ideal policy $\bar{U}$. These plots illustrate that, in all cases, the learning achieved from the KG policy occurs much more rapidly. This is an attractive property of a learning policy in marketing campaign design for market entry from a managerial viewpoint, since time is an important factor to maximize the value of marketing investment. Being capable to reduce the required number of learning phases to achieve the same level of learning makes the KG policy more favorable than alternative sequential learning policies. This confirms the value of adopting KG policy in designing marketing campaigns. By comparing plots 3(a) and 3(b) with those in 3(c), we observe that when Σ_Z and Σ_B is nonzero the A-Design and E-Design policies on average outperform the Thompson sampling policy.

Next, we investigate the performance of the approximate algorithm using a sequence of mixed-integer linear optimization problems in Figure 2 for computing the KG decision. For this experiment, we consider a feasible set $\mathcal{X} \subseteq \mathbb{R}^m$ with m decision variables, 5 equality constraints $Ax = h$

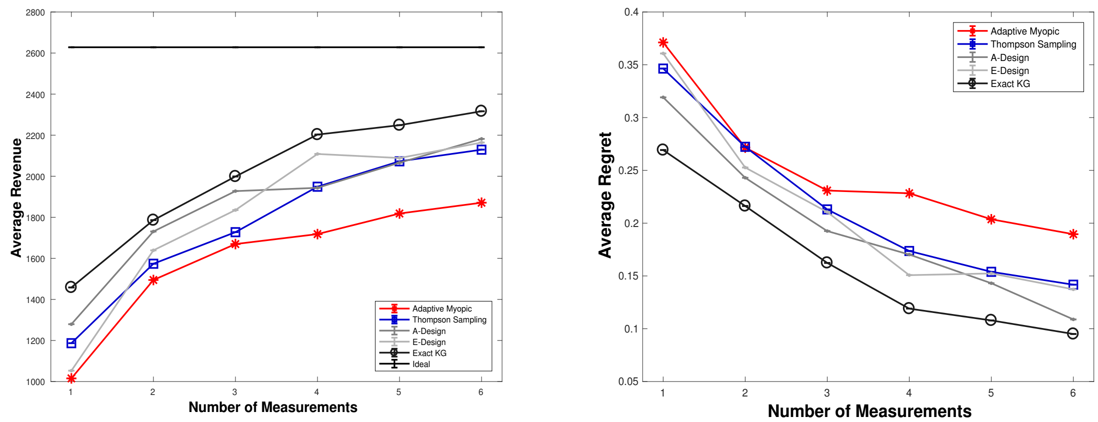

(a) $r_0 = 13$, $r = 20$, $\Sigma_Z = \bar{\Sigma}$, $\Sigma_B = \bar{\Sigma}_0$

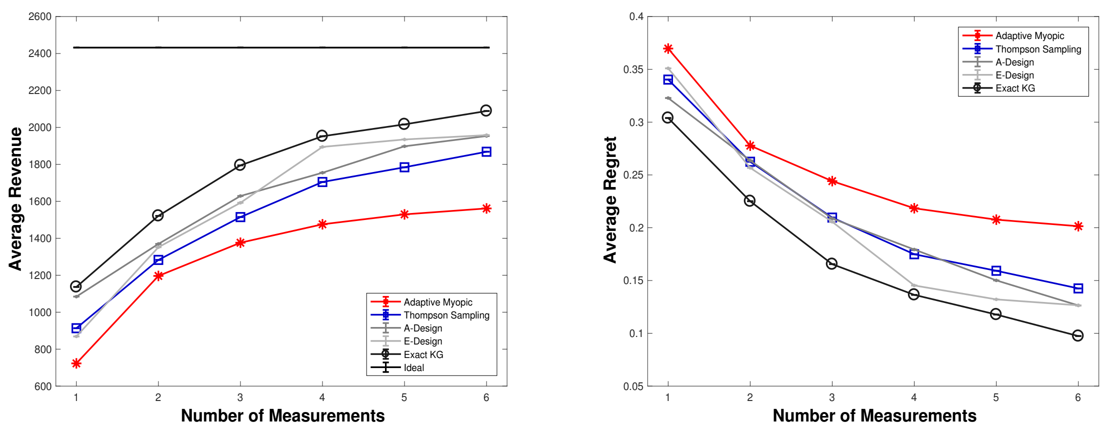

(b) $r_0 = 20$, $r = 13$, $\Sigma_Z = \bar{\Sigma}$, $\Sigma_B = \bar{\Sigma}_0$

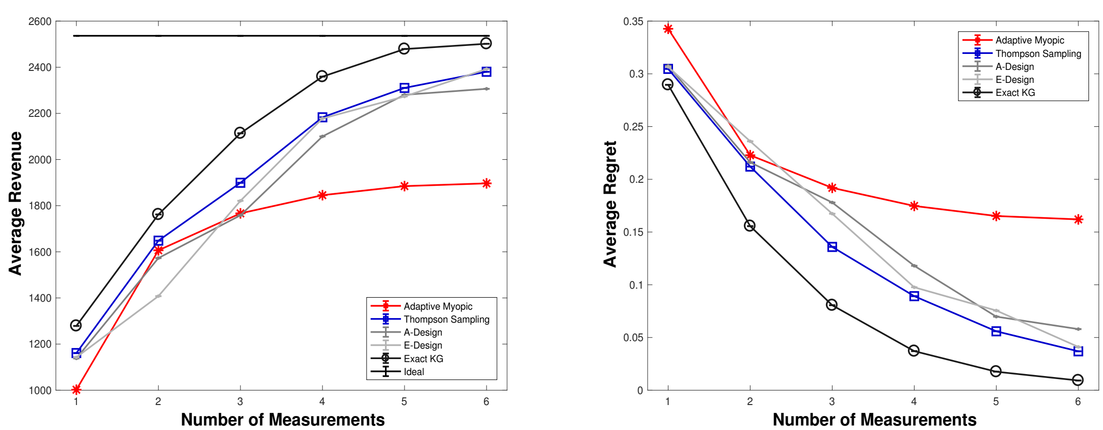

(c) $r_0 = 20$, $r = 13$, $\Sigma_Z = 0$, $\Sigma_B = 0$

Figure 3 **Average revenue and average regret of different sequential learning policies**

and no inequality constraint. In addition, we let $r = m$, $\Sigma_B = 0$, $\lambda = 1$, and $J = 10$. For a given value of m, we randomly generate 20 problem instances $(A, h, \theta_N, \Sigma_N)$. For each generated problem, we employ the iterative MILP algorithm in Figure 2 and the mixed-integer second-order cone optimization (MISOCO) problem (37) in Proposition 6. The average runtimes for each computational approach, the average number of iterations $\bar{N}$ in the MILP sequence algorithm, and average relative approximation error in the obtained optimal values are reported in Table 3. Here, the values are averaged over 20 problem instances. According to our simulations reported in Table 3, the algorithm in Figure 2 requires very few (around $\bar{N} \approx 3$) iterations on average to converge to a solution with less than 2% difference in the KG quantity. This holds for all values of m. The runtime of solving MISOCO formulation to compute the KG decision is lower for smaller values of m. As the number of input features m increases ($m \geq 24$), applying the algorithm in Figure 2 becomes much more attractive.

			Runtime	
m	Relative Difference in KG quantity	Average $\bar{N}$	Iterative MILP	MISOCO
12	0.59%	3.10	20.85	0.81
16	1.01%	3.05	53.54	4.09
20	1.81%	3.45	93.57	36.84
22	1.54%	3.05	106.33	91.58
24	1.47%	3.00	128.58	665.76
26	1.78%	2.95	144.66	1924.56

Table 3 **Runtime comparison between the MILP sequence algorithm in Figure 2 and MISOCO formulation**

10. Conclusions

This paper addresses modeling and computational challenges around adaptive experimentation in the context of marketing campaign design for market entry. We consider a multiplicative advertising exposure model which allows for characterizing a marketing campaign by detailed input features and capturing the advertising exposure effect by a Poisson process. A Bayesian approach to model parameter uncertainty is adopted to capture the interaction between design features and customer segmentation features. The model admits closed-form Bayesian updating equations and computationally efficient approaches to develop an approximate dynamic programming policy. We evaluate this policy and a set of alternative policies, including the practitioners' favorite adaptive myopic policy and the Thompson sampling policy, found often outperforming in (Schwartz et al. 2016), to design a property and casualty insurance marketing campaign.

A limiting factor of this study is that the computational investigation does not consider features related to the promotion period in the advertising schedule, which could describe the ordering of successive marketing campaigns that a same customer could be exposed to in order to induce an overall response. In (Tellis and Franses 2006) for instance it was shown that the manager can consider the promotion period as a decision variable, and let the optimization guide the next length of promotion period, taking into account the effectiveness period of advertising and the optimal data interval to estimate the advertising carryover.

Acknowledgement

This research was supported in part by a grant from Accenture.

Appendix A: Proofs

Sketch of Proof of Proposition 1 (see Section EC.1). From $\zeta \sim \mathcal{N}(\mu_Z, \Sigma_Z)$ and $\beta \sim \mathcal{N}(\mu_B, \Sigma_B)$ and the assumption that ζ, β, and ϵ are independent, $\zeta^\top x_Z + \beta^\top x_B + \epsilon \mid \mu_B, \rho$ is normally distributed, for every given $(x_{Z,n}, x_{B,n})$. Thus, given the multiplicative advertising exposure model (2), $\eta_{n+1}|\mu_B, \rho, x_n, \kappa_{n+1} > 0$ follows a normal distribution with mean $\kappa_{n+1}(\mu_Z^\top x_{Z,n} + \mu_B^\top x_{B,n})$ and variance $\frac{\kappa_{n+1}^2}{\rho}\hat{\sigma}_n$, where $\hat{\sigma}_n \overset{\text{def}}{=} 1 + x_{Z,n}^\top \Sigma_Z x_{Z,n} + x_{B,n}^\top \Sigma_B x_{B,n}$.

From the Bayes' rule, the probability density functions of $\eta_{n+1}|\mu_B, \rho, x_n, \kappa_{n+1} > 0$ and $\mu_B, \rho|S_n$ where $(\mu_B, \rho)|S_n \sim \text{Normal-Gamma}(\theta_n, \Sigma_n, a_n, b_n)$, we have

$$\begin{aligned}\Pr(\mu_B, \rho|S_n, x_n, \kappa_{n+1}, \eta_{n+1}) \propto \left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} |\Sigma_n|^{-1/2} e^{-\frac{\rho}{2}(\mu_B-\theta_n)^\top \Sigma_n^{-1}(\mu_B-\theta_n)} \times \frac{b_n^{a_n}}{\Gamma(a_n)} \rho^{a_n-1} e^{-b_n\rho} \\ \times \frac{\sqrt{\hat{\rho}}}{\sqrt{2\pi}\kappa_{n+1}} e^{-\frac{\hat{\rho}}{2}\left(\ell+(\theta_n-\mu_B)^\top x_{B,n}\right)^2},\end{aligned} \tag{53}$$

where $\ell \overset{\text{def}}{=} \frac{\eta_{n+1}}{\kappa_{n+1}} - (\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n})$. Given equation (9), Assumptions in Lemma EC.3 hold for $\gamma = 1$ and $\alpha = \hat{\sigma}^{-1}$. Therefore, by applying Lemma EC.3 and equation (12) in (53), and using equation (11), $a_{n+1} = a_n + \frac{1}{2}$, we have

$$\begin{aligned}&\left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} |\Sigma_n|^{-1/2} e^{-\frac{\rho}{2}(\mu_B-\theta_n)^\top \Sigma_n^{-1}(\mu_B-\theta_n)} \times \frac{b_n^{a_n}}{\Gamma(a_n)} \rho^{a_n-1} e^{-b_n\rho} \frac{\sqrt{c\rho}}{\sqrt{2\pi}\kappa_{n+1}} e^{-\frac{c\rho}{2}\left(\ell+(\theta_n-\mu_B)^\top x_{B,n}\right)^2} \\ &= c_0 \left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} |\Sigma_{n+1}|^{-\frac{1}{2}} e^{-\frac{\rho}{2}(\mu_B-\theta_{n+1})^\top \Sigma_{n+1}^{-1}(\mu_B-\theta_{n+1})} \times \frac{b_{n+1}^{a_{n+1}}}{\Gamma(a_{n+1})} \rho^{a_{n+1}-1} e^{-b_{n+1}\rho}.\end{aligned}$$

where c_0 does not depend on ρ and μ_B (see equation (EC.14)). Thus, $\mu_B, \rho|S_n, x_n, \kappa_{n+1}, \eta_{n+1} \sim \text{Normal-Gamma}(S_{n+1})$. □

Proof of Proposition 2. It follows from $\lambda^{\mathscr{M}}(x_{\bar{k}}) = \lambda_{\bar{k}} \sim \textit{Gamma}\left(c_n(x_{\bar{k}}), d_n(x_{\bar{k}})\right)$ that $\Pr\left(\lambda_{\bar{k}} \mid c_n(x_{\bar{k}}), d_n(x_{\bar{k}})\right) \propto \lambda_{\bar{k}}^{c_n(x_{\bar{k}})-1} e^{-d_n(x_{\bar{k}})\lambda_{\bar{k}}}$. Let κ_{n+1} be drawn from $\textit{Poisson}\left(\lambda^{\mathscr{M}}(x_n)\right)$. The conditional distribution of $\lambda^{\mathscr{M}}(x_n) = \lambda^{\mathscr{M}}(x_{\bar{k}}) = \lambda_{\bar{k}}$ given S_{n+1} is

$$\begin{aligned}\Pr\left(\lambda_{\bar{k}} \mid c_n(x_{\bar{k}}), d_n(x_{\bar{k}}), \kappa_{n+1}\right) &\propto \Pr\left(\kappa_{n+1}|\lambda_{\bar{k}}\right) \times \Pr\left(\lambda_{\bar{k}} \mid c_n(x_{\bar{k}}), d_n(x_{\bar{k}})\right) \\ &\propto \lambda_{\bar{k}}^{\kappa_{n+1}} e^{-\lambda_{\bar{k}}} \times \lambda_{\bar{k}}^{c_n(x_{\bar{k}})-1} e^{-d_n(x_{\bar{k}})\lambda_{\bar{k}}} = \lambda_{\bar{k}}^{c_n(x_{\bar{k}})+\kappa_{n+1}-1} e^{-\left(d_n(x_{\bar{k}})+1\right)\lambda_{\bar{k}}}.\end{aligned}$$

Hence, the posterior distribution remains a Gamma distribution with the shape parameter $c_n(x_{\bar{k}}) + \kappa_{n+1}$ and the rate parameter $d_n(x_{\bar{k}}) + 1$ as in equations (14) and (15). □

Sketch of Proof of Proposition 3 (see Section EC.2). The description of η_{n+1} in equation (2) and $\mu_B|S_n,\rho \sim \mathcal{N}(\theta_n, \frac{1}{\rho}\Sigma_n)$ imply that $\eta_{n+1} \mid S_n,\rho,x_n,\mu_B,\kappa_{n+1}$ follows a compound normal distribution whose mean $\kappa_{n+1}(\mu_Z^\top x_{Z,n} + \mu_B^\top x_{B,n})$ itself follows a Gaussian distribution. Therefore, $\eta_{n+1} \mid S_n,\rho,x_n,\kappa_{n+1}$ follows a normal distribution with mean $\kappa_{n+1}(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n})$ and variance $\frac{\kappa_{n+1}^2}{\rho}\hat{\sigma}_n + \frac{\kappa_{n+1}^2}{\rho} x_{B,n}{}^\top \Sigma_n x_{B,n}$, where $\hat{\sigma}_n = 1 + x_{Z,n}^\top \Sigma_Z x_{Z,n} + x_{B,n}^\top \Sigma_B x_{B,n}$. Define,

$$\hat{\eta}_{n+1} \stackrel{\text{def}}{=} \sqrt{\frac{a_n}{\kappa^2 b_n(\hat{\sigma}_n + x_{B,n}{}^\top \Sigma_n x_{B,n})}} \left(\eta_{n+1} - \kappa(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n})\right).$$

From the probability density of $\eta_{n+1} = \eta \mid S_n, x_n, \kappa_{n+1} = k > 0$, we observe that $\hat{\eta}_{n+1}|S_n, x_n, \kappa_{n+1} = k$ follows a standard Student's t-distribution with $\nu = 2a_n$ degrees of freedom, which yields (16). $\square$

Sketch of Proof of Proposition 4 (see Section EC.3). Denote the observed outcome of the marketing campaign characterized through the input feature vector y by (κ_y, η_y). The proof relies on backward induction on n, the distributional structure of η_{n+1} and $\mathbb{E}[T_{2a_n}] = 0$, and that the induced transition function $S^{\mathscr{M}}$ by the belief model satisfies $S^{\mathscr{M}}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right), x, (\kappa_x, \eta_x)\right) = S^{\mathscr{M}}\left(S^{\mathscr{M}}\left(S_n, x, (\kappa_x, \eta_x)\right), y, (\kappa_y, \eta_y)\right)$, for any $y, x \in \mathcal{X}$. At stage $n = N-1$, we have

$$Q_{N-1}(S_{N-1}, x_{N-1}) = \mathbb{E}\left[\max_{y\in\mathcal{X}} \lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_N^\top y_B\right)\right] \geq \max_{y\in\mathcal{X}} \lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_{N-1}^\top y_B\right) = V_N(S_{N-1}).$$

Assuming the statement for $n+1$, $Q_{n+1}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right), x_n\right) \geq V_{n+2}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right)\right)$ almost surely for all $y \in \mathcal{X}$, we have for any $x_n \in \mathcal{X}$,

$$\begin{aligned} Q_n(S_n, x_n) &\geq \max_{y\in\mathcal{X}} \ \mathbb{E}\left[\mathbb{E}\left[V_{n+2}\left(S^{\mathscr{M}}\left(S^{\mathscr{M}}(S_n, y, (\kappa_y, \eta_y)), x_n, (\kappa_{n+1}, \eta_{n+1})\right)\right)\right]\right] \\ &= \max_{y\in\mathcal{X}} \ \mathbb{E}\left[Q_{n+1}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right), x_n\right)\right] \geq \max_{y\in\mathcal{X}} \mathbb{E}\left[V_{n+2}(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right))\right] = V_{n+1}(S_n) \end{aligned}$$

This completes the proof of the first part. When the one more marketing campaign can be measured according to the optimal policy π^*, we have $Q_n(S_n, x_n^*) = V_n(S_n)$. This and the inequality $Q_n(S_n, x_n) \geq V_{n+1}(S_n)$ yields $V_n(S_n) \geq V_{n+1}(S_n)$. $\square$

Sketch of Proof of Proposition 5 (see Section EC.4). Equation (20) implies that $V_N^\pi(S_n) := U(\mathbb{E}[\mu_B|S_N = S_n]) = U(\theta_n)$. Therefore,

$$v_x^{\text{KG},n} = \mathbb{E}\left[U\left(\theta^{\mathscr{M}}\left(S_n, x, (\kappa_{n+1}, \eta_{n+1})\right)\right) \mid S_n, x\right] - U(\theta_n). \tag{54}$$

Here, $\theta^{\mathscr{M}}\left(S_n, x, (\eta_{n+1}, \kappa_{n+1})\right) = \mathbb{E}\left[\mu_B \mid S_N = S^{\mathscr{M}}\left(S_n, x, (\eta_{n+1}, \kappa_{n+1})\right)\right]$. The first term then equals

$$\begin{aligned} &\mathbb{E}_{\kappa_{n+1}}\left[\mathbb{E}_{\eta_{n+1}}\left[U\left(\theta^{\mathscr{M}}\left(S_n, x, (\kappa_{n+1}, \eta_{n+1})\right)\right) \mid S_n, x, \kappa_{n+1}\right]\right] \\ &= e^{-\lambda^{\mathscr{M}}(x)} U(\theta_n) + (1 - e^{-\lambda^{\mathscr{M}}(x)})\mathbb{E}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x, \kappa_{n+1} > 0\right]. \end{aligned} \tag{55}$$

Applying Proposition 3 and equation (17), we have

$$\begin{aligned} &\mathbb{E}\left[U\left(\theta^{\mathscr{M}}\left(S_n, x, (\kappa_{n+1}, \eta_{n+1})\right)\right) \mid S_n, x, \kappa_{n+1} > 0\right] \\ &= \mathbb{E}\left[\max_{y\in\mathcal{X}} \lambda(y)\left(\mu_Z^\top y_Z + \left(\theta_n + \sqrt{\frac{b_n}{a_n(\hat{\sigma}_n(x) + x_B^\top \Sigma_n x_B)}} \Sigma_n x_B T_{2a_n}\right)^\top y_B\right) \mid S_n, x, \kappa_{n+1} > 0\right] \\ &= \mathbb{E}\left[\max_{y\in\mathcal{X}} \left(p_y^n + q_y^n(x) T_{2a_n}\right) \mid S_n, x, \kappa_{n+1} > 0\right], \end{aligned} \tag{56}$$

where p_y^n and $q_y^n(x_n)$ are given in equations (27) and (28). Applying equation (56) in (55), and finally in equation (54) completes the derivation of (26). $\square$

Proof of Proposition 7. Under the linearity assumption for intensity $\lambda^{\mathscr{M}}(x) = \bar{\lambda}^\top x$, we have

$$\lambda^{\mathscr{M}}(y^j)(\mu_Z;\theta_n)^\top y^j = (y^j)^\top \bar{\lambda}(\mu_Z;\theta_n)^\top y^j = \mathrm{tr}\left(\bar{\lambda}(\mu_Z;\theta_n)^\top \cdot y^j (y^j)^\top\right) = \mathrm{vec}((\mu_Z;\theta_n)\bar{\lambda}^\top)^\top \mathrm{vec}(u^j), \quad \forall\, y^j \quad (57)$$

where u^j satisfies $u^j = y^j\,(y^j)^\top$. Similarly, by defining the decision variable $z^j = \mathrm{vec}\left(y_B^j x_B^\top\right)(y^j)^\top$,

$$\lambda^{\mathscr{M}}(y^j) x_B^\top \Sigma_n y_B^j = \bar{\lambda}^\top y^j \mathrm{tr}\left(x_B^\top \Sigma_n y_B^j\right) = \mathrm{tr}\left((y^j)^\top \bar{\lambda} \mathrm{tr}\big(\Sigma_n y_B^j x_B^\top\big)\right) = \mathrm{vec}\left(\mathrm{vec}(\Sigma_n)\bar{\lambda}^\top\right)^\top \mathrm{vec}(z^j). \quad (58)$$

The constraints $u^j = y^j\,(y^j)^\top$ and $z^j = \mathrm{vec}\left(y_B^j x_B^\top\right)(y^j)^\top$ can be expressed by linear inequalities. Applying (57) and (58) in the maximization (38) along with the linearized inequality constraints for u^j and z^j yields (39). Equation (40) comes from $\nu = \max\limits_{x\in\mathcal{X}} \mathrm{vec}\left((\mu_Z;\theta_n)\bar{\lambda}^\top\right)^\top \mathrm{vec}\left(xx^\top\right)$. $\square$

Proof of Proposition 8. The constraints (49) and $(\hat{x}, k) \in \mathcal{X}^+$ collectively imply that $kx = \hat{x}$. Constraints (43)-(45) yield $u_{ii'} = x_i \hat{x}_{i'}$. Hence, from constraint (42), $x^\top P \hat{x} = 1$ or equivalently $\hat{x}^\top P \hat{x} = k$. The constraints (46)-(48) yield $z^j = y^j x^\top$. Thus, $\mathrm{vec}(\Sigma_n)^\top \mathrm{vec}(z^j) = x_B^\top \Sigma_n z_B^j$, and

$$\begin{aligned}\alpha \textstyle\sum_{j=1}^J t_j w_j (z_B^j)^\top \Sigma_n x_B &= \frac{\alpha}{k}\sum_{j=1}^J t_j w_j (k x_B)^\top \Sigma_n z_B^j = \frac{\alpha}{k}\sum_{j=1}^J t_j w_j (z_B^j)^\top \Sigma_n \hat{x}_B \\ &= \frac{\alpha}{\sqrt{k}}\sum_{j=1}^J t_j w_j \frac{(z_B^j)^\top \Sigma_n \hat{x}_B}{\sqrt{k}} = \frac{\alpha}{\sqrt{k}}\sum_{j=1}^J t_j w_j \frac{(z_B^j)^\top \Sigma_n \hat{x}_B}{\sqrt{\hat{x}^\top P \hat{x}}} = \frac{\alpha}{\sqrt{k}}\sum_{j=1}^J t_j w_j \frac{x_B^\top \Sigma_n z_B^j}{\sqrt{x^\top P x}}.\end{aligned}$$

Therefore, the MILP problem $(\mathcal{P}_\alpha)$ finds an approximate solution of problem (37), when $\frac{\alpha}{\sqrt{k^*}} \approx 1$. Finally, the constraints on z^j along with $z^1, \cdots, z^J \in \mathbb{R}_+^{r\times r}$ implies that $z^1, \cdots, z^J \in \{0,1\}^{r\times r}$. $\square$

PROPOSITION 9. *Let* $Q_n \overset{\text{def}}{=} \frac{A^\top A}{h^\top h} + \begin{pmatrix}\Sigma_Z & 0\\ 0 & \Sigma_n + \Sigma_B\end{pmatrix}$. *The E-optimal design policy* $\min\limits_{x_n\in\mathcal{X}} \log\det(\Sigma_{n+1})$ *can be computed by the following mixed-integer second order cone optimization:*

$$\begin{aligned}\min_{x_n\in\mathcal{X},\ V\geq 0,\ z}\quad & z\\ s.t\quad & \left\|\begin{matrix}2Q_n^{1/2}x_n\\ tr(\Sigma_n V) - z\end{matrix}\right\|_2 \leq tr(\Sigma_n V) + z,\\ & (x_{B,n})_k + (x_{B,n})_l - 1 \leq V_{kl} \leq \min\{(x_{B,n})_k, (x_{B,n})_l\}.\end{aligned}$$

Sketch of Proof of Proposition 9 (see Section EC.5). From equation (10) and $\Sigma_n \succ 0$, we get $\log\det(\Sigma_{n+1}) = \log\left(1 - \frac{1}{x_n^\top Q_n x_n} x_{B,n}{}^\top \Sigma_n x_{B,n}\right) + \log\det(\Sigma_n)$. Since at stage n, Σ_n is a constant and log is an increasing function, to minimize $\log\det(\Sigma_{n+1})$ it is sufficient to maximize $\frac{x_{B,n}{}^\top \Sigma_n x_{B,n}}{x_n^\top Q_n x_n}$. Defining $V = x_{B,n} x_{B,n}{}^\top$ yields $x_{B,n}{}^\top \Sigma_n x_{B,n} = \mathrm{tr}(\Sigma_n V)$. Hence, we have minimizing z subject to $\frac{x_n^\top Q_n x_n}{\mathrm{tr}(\Sigma_n V)} \leq z$ or equivalently $\left\|\begin{matrix}2Q_n^{1/2}x_n\\ \mathrm{tr}(\Sigma_n V) - z\end{matrix}\right\|_2 \leq \mathrm{tr}(\Sigma_n V) + z$. $\square$

References

Abe N, Verma N, Apte C, Schroko R (2004) Cross channel optimized marketing by reinforcement learning. *Proceedings of the 10th ACM International Conference on Knowledge Discover and Data Mining* 1–6.

Adams RP, Murray I, MacKay DJ (2009) Tractable nonparametric Bayesian inference in Poisson processes with gaussian process intensities. *Proceedings of the 26th International Conference on Machine Learning* 1–8.

Agrawal S, Goyal N (2012) Analysis of thompson sampling for the multi-armed bandit problem. *The 25th Annual Conference on Learning Theory* 23:39.1–39.26.

Barajas J, Akella R, Holtan M, Flores A (2016) Experimental designs and estimation for online display advertising attribution in marketplaces. *Marketing Science* 35(3):465–483.

Basu AK, Basu A, Batra R (1995) Modeling the response pattern to direct marketing campaigns. *Journal of Marketing Research* 32:204–212.

Bauer LC (1987) Direct response advertising: Forecasting responses over time. *Journal of Direct Marketing* 1:38–50.

Bauer LC (1991) Logistic versus decaying exponential equations for describing mail survey response curves: A conceptual rationale and reanalysis. *Journal of Direct Marketing* 5:15–26.

Bechhofer R, Santner T, Goldsman D (1995) *Design and Analysis of Experiments for Statistical Selection, Screening and Multiple Comparisons* (John Wiley & Sons, New York).

Bertsekas DP (2012) *Dynamic Programming and Optimal Control, Vol II* (Athena Scientific), fourth edition.

Bertsimas D, Mersereau AJ (2007) A learning approach for interactive marketing to a customer segment. *Operations Research* 55(6):1120–1135.

Bose I, Chen X (2009) Quantitative models for direct marketing: A review from systems perspective. *European Journal of Operational Research* 195:1–16.

Boyd S (2004) *Convex Optimization* (Cambridge University Press), first edition.

Chapelle O, Li L (2012) An empirical evaluation of thompson sampling. *Advances in Neural Information Processing Systems* 24:2249–2257.

Charski M (2013) Analytics of customer data drives personalized marketing campaigns. *DataInformed Big Data and Analytics in the Enterprise* 1–2.

Chick SE (2006) Subjective probability and Bayesian methodology. Henderson SG, Nelson BL, eds., *Handbooks of Operations Research and Management Science, vol. 13: Simulation*, 225–258 (North-Holland Publishing, Amsterdam).

Chick SE, Branke J, Schmidt C (2010) Sequential sampling to myopically maximize the expected value of information. *INFORMS Journal on Computing* 22(1):71–80.

Chick SE, Frazier P (2013) Sequential sampling with economics of selection procedures. *Management Science* 58(3):550–569.

Ching AT, Erdem T, Keane MP (2013) Learning models: An assessment of progress, challenges, and new developments. *Marketing Science* 32(6):913–938.

Chun YH (2012) Monte Carlo analysis of estimation methods for the prediction of customer response patterns in direct marketing. *European Journal of Operational Research* 217:673–678.

Chun YH, Jung Y (2015) Predictive modeling of customer response behavior in direct marketing. Warkentin M, ed., *The Best Thinking in Business Analytics from the Decision Sciences Institute*, 1–16.

Cobanoglu C, Cobanoglu N (2003) The effect of incentives in web surveys: Application and ethical considerations. *International Journal of Market Research* 45:475–488.

Coussement K, Buckinx W (2011) A probability-mapping algorithm for calibrating the posterior probabilities: A direct marketing application. *European Journal of Operational Research* 214:732–738.

Defourny B, Ryzhov I, Powell W (2015) Optimal information blending with measurements in the l2 sphere. *Mathematics of Operations Research* 40(4):1060–1088.

Farris P, Bendle N, Pfeifer PE, Reibstein DJ (2015) *Marketing Metrics: The Manager's Guide to Measuring Marketing Performance* (Pearson FT Press), third edition.

Foedermayr EK, Diamantopoulos A (2008) Market segmentation in practice: Review of empirical studies, methodological assessment and agenda for future research. *Journal of Strategic Marketing* 16(3):223–265.

Fox RJ, Reddy SK, Rao B (1997) Modeling response to repetitive promotional stimuli. *Journal of the Academy of Marketing Science* 25(3):242–255.

Frazier PI, Powell WB, Dayanik S (2008) A knowledge gradient policy for sequential information collection. *SIAM Journal on Control and Optimization* 47:2410–2439.

Frazier PI, Powell WB, Dayanik S (2009) The Knowledge-Gradient policy for correlated normal beliefs. *INFORMS Journal on Computing* 21(4):599–613.

Gittins J, Glazebrook KD, Weber RR (2011) *Multi-Armed Bandit Allocation Indices* (John Wiley & Sons, Chichester, UK), second edition.

Golub GH, Loan CFV (1996) *Matrix Computation* (Hopkins Fulfillment Service), third edition.

Goyat S (2011) The basis of market segmentation: a critical review of literature. *European Journal of Business and Management* 3(9):1–11.

Graf S, Luschgy H (2000) *Foundations of Quantization for Probability Distributions* (Springer-Verlag, Berlin, Germany).

Gupta S, Zeithaml V (2006) Customer metrics and their impact on financial performance. *Marketing Science* 25(6):718–739.

Gupta SS, Miescke KJ (1996) Bayesian look ahead one-stage sampling allocations for selection of the best population. *Journal of Statistical Planning and Inference* 54(2):229–244.

Gurobi Optimization I (2015) Gurobi optimizer reference manual. URL `http://www.gurobi.com`.

Han B, Ryzhov IO, Defourny B (2015) Optimal learning in linear regression with combinatorial feature selection. *INFORMS Journal on Computing* 1–37.

Hanssens DM, Leeflang PSH, Wittink DR (2005) Market response models and marketing practice. *Applied Stochastic Models In Business And Industry* 21:423–434.

Hill RW (1981) Using s-shaped curves to predict response rates. *Journal of Marketing Research* 18:240–242.

Huxley SJ (1980) Predicting response speed in mail surveys. *Journal of Marketing Research* 17:63–68.

Jones DR, Schonlau M, Welch W (1998) Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* 13:455–492.

Klapdor S, Anderl EM, von Wangenheim F, Schumann JH (2014) Finding the right words: The influence of keyword characteristics on performance of paid search campaigns. *Journal of Interactive Marketing* 28:285–301.

Kottas A, Sanso B (2007) Bayesian mixture modeling for spatial Poisson process intensities, with applications to extreme value analysis. *Journal of Statistical Planning and Inference* 137(10):3151–3163.

Kunreuther H, Meyer R, Michel-Kerjan E (2014) To get people to buy insurance, change how you talk about risk. *Bloomberg Business* 1–2.

Lilien GL, Rangaswamy A, Bruyn AD (2012) *Principles of Marketing Engineering* (DecisionPro, Inc.), second edition.

Lin S, Zhang J, Hauser JR (2015) Learning from experience, simply. *Marketing Science* 34(1):1–9.

Lu T, Boutilier C (2014) Dynamic segmentation for large-scale marketing optimization. *International Conference on Machine Learning* 1–8.

McCrary M (2009) Enhanced customer targeting with multi-stage models: Predicting customer sales and profit in the retail industry. *Journal of Targeting, Measurement and Analysis for Marketing* 17:273–295.

McGranahan D, Catlin T, Ray S (2013) Beyond price: The rise of customer-centric marketing in insurance. *McKinsey & Company, Financial Services Practice Report* 1–32.

Moro S, Cortez P, Rita P (2014) A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems* 62:22–31.

Moro S, Laureano R, Cortez P (2011) Using data mining for bank direct marketing: An application of the crisp-dm methodology. *Proceedings of the European Simulation and Modelling Conference* 117–121.

MOSEK ApS (2015) *The MOSEK optimization toolbox for MATLAB manual. Version 7.1 (Revision 28).* URL `http://docs.mosek.com/7.1/toolbox/index.html`.

Negoescu DM, Frazier PI, Powell WB (2011) The knowledge gradient algorithm for sequencing experiments in drug discovery. *INFORMS Journal on Computing* 23:346–363.

Nesamoney D (2015) *Personalized Digital Advertising: How Data and Technology Are Transforming How We Market* (Pearson FT Press), first edition.

Pages G, Wilbertz B (2012) Optimal delaunay and voronoi quantization schemes for pricing american style options. Carmona RA, Moral PD, Hu P, Oudjane N, eds., *Numerical Methods in Finance*, 171–213.

Powell WB (2011) *Approximate Dynamic Programming: Solving the Curses of Dimensionality* (Wiley Series in Probability and Statistics), second edition.

Puterman ML (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, Inc. New York, NY, USA).

Rossi PE, Allenby G (2003) Bayesian statistics and marketing. *Marketing Science* 22:304–328.

Russo D, Roy BV (2014) An information-theoretic analysis of thompson sampling. *Journal of Machine Learning Research* 1–30.

Ryzhov IO, Powell W, Frazier PI (2012) The knowledge-gradient algorithm for a general class of online learning problems. *Operations Research* 60(1):180–195.

Ryzhov IO, Powell WB (2012a) Information collection for linear programs with uncertain objective coefficients. *SIAM Journal on Optimization* 22(4):1344–1368.

Ryzhov IO, Powell WB (2012b) *Optimal Learning* (Wiley), first edition.

Schultz DE, Barnes BE, Schultz HF, Azzaro M (2009) *Building Customer-brand Relationships* (Routledge), first edition.

Schwartz EM, Bradlow E, Fader P (2016) Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science* In press.

Tellis GJ, Franses PH (2006) Optimal data interval for estimating advertising response. *Marketing Science* 25(3):217–229.

Wainwright MJ, Jordan MI, et al. (2008) Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning* 1(1–2):1–305.

Webber H (1998) *Divide and Conquer: Target Your Customers Through Market Segmentation* (Wiley).

Wedel M, Kamakura W (2000) *Market segmentation: Conceptual and methodological foundations* (Boston, MA: Kluwer Academic Publishers), second edition.

Wedel M, Kamakura WA (2002) Introduction to the special issue on market segmentation. *Journal of classification* 19:179–182.

Weinstein A (2006) A strategic framework for defining and segmenting markets. *Journal of Strategic Marketing* 14(2):115–127.

Whelan D, O'Neill S (2014) Customer loyalty in P&C insurance. *Bain & Company Report* 1–16.

Electronic Companion

EC.1. Proof of Proposition 1

We start by establishing a few intermediate results.

LEMMA EC.1. *Suppose*

[A] $\Sigma_{n+1}^{-1}=\Sigma_n^{-1}+\alpha x_{B,n}{x_{B,n}}^\top$ *for some* $\alpha>0$.

Then

$$|\Sigma_{n+1}|^{-\frac{1}{2}}=|\Sigma_n|^{-\frac{1}{2}}\sqrt{1+\alpha {x_{B,n}}^\top\Sigma_n x_{B,n}}. \tag{EC.1}$$

Proof. It follows from assumption [A] and the fact that for any invertible square matrix X, we have $|X+AB|=|X||I+BX^{-1}A|$ and $|A^{-1}|=|A|^{-1}$, that

$$|\Sigma_{n+1}|^{-1}=|\Sigma_{n+1}^{-1}|=|\Sigma_n^{-1}+\alpha x_{B,n}{x_{B,n}}^\top|=|\Sigma_n^{-1}||1+\alpha {x_{B,n}}^\top\Sigma_n x_{B,n}|.$$

Positive semidefiniteness of Σ_Z, Σ_B, and Σ_n implies that $1+\alpha {x_{B,n}}^\top\Sigma_n x_{B,n}>0$. This completes the proof of equation (EC.1). □

LEMMA EC.2. *Let* Σ_{n+1} *be defined as in equation* (10)*. Then* [A] *holds for* $\alpha=(1+x_{Z,n}^\top\Sigma_Z x_{Z,n}+x_{B,n}^\top\Sigma_B x_{B,n})^{-1}$.

Proof. It follows from the matrix inversion lemma, $(A+UCV)^{-1}=A^{-1}-A^{-1}U(C^{-1}+VA^{-1}U)^{-1}VA^{-1}$, that

$$\begin{aligned}\Sigma_{n+1}^{-1}&=\left(\Sigma_n-\frac{\Sigma_n x_n {x_n}^\top\Sigma_n}{1+{x_n}^\top\Sigma_n x_n+x_{Z,n}^\top\Sigma_Z x_{Z,n}+{x_{B,n}}^\top\Sigma_B x_{B,n}}\right)^{-1}\\&=\Sigma_n^{-1}+\Sigma_n^{-1}\Sigma_n x_n\left(1+x_{Z,n}^\top\Sigma_Z x_{Z,n}+{x_n}^\top\Sigma_B x_n\right)^{-1}{x_n}^\top\Sigma_n\Sigma_n^{-1}=\Sigma_n^{-1}+\alpha x_n {x_n}^\top.\end{aligned}\tag{EC.2}$$

This completes the proof. □

LEMMA EC.3. *Let* [A] *hold. In addition, suppose*

[B] there exists L *such that* $\theta_{n+1}=\theta_n+L$.

[C] there exists $\gamma>0$ *such that* $L=\frac{\gamma\ell}{\alpha^{-1}+{x_{B,n}}^\top\Sigma_n x_{B,n}}\Sigma_n x_{B,n}$.

Then,

$$\begin{aligned}(\mu_B-\theta_{n+1})^\top\Sigma_{n+1}^{-1}(\mu_B-\theta_{n+1})&=(\mu_B-\theta_n)^\top\Sigma_n^{-1}(\mu_B-\theta_n)+\alpha\left(\ell\gamma-(\mu_B-\theta_n)^\top x_{B,n}\right)^2\\&\quad-\frac{\gamma^2\ell^2}{{x_{B,n}}^\top\Sigma_n x_{B,n}+\alpha^{-1}}.\end{aligned}\tag{EC.3}$$

Proof. From Assumption [B],

$$(\mu_B-\theta_{n+1})^\top\Sigma_{n+1}^{-1}(\mu_B-\theta_{n+1})=(\mu_B-\theta_n)^\top\Sigma_{n+1}^{-1}(\mu_B-\theta_n)-2(\mu_B-\theta_n)^\top\Sigma_{n+1}^{-1}L+L^\top\Sigma_{n+1}^{-1}L. \quad \text{(EC.4)}$$

Next, we apply Assumption [A] for each term in (EC.4). First,

$$(\mu_B-\theta_n)^\top\Sigma_{n+1}^{-1}(\mu_B-\theta_n)=(\mu_B-\theta_n)^\top\Sigma_n^{-1}(\mu_B-\theta_n)+\alpha(\mu_B-\theta_n)^\top x_{B,n}{x_{B,n}}^\top(\mu_B-\theta_n). \quad \text{(EC.5)}$$

From Assumptions [A] and [C], we have

$$\begin{aligned}L^\top\Sigma_{n+1}^{-1}L&=\left(\frac{\gamma\ell}{{x_{B,n}}^\top\Sigma_n x_{B,n}+\alpha^{-1}}\right)^2{x_{B,n}}^\top\Sigma_n\left(\Sigma_n^{-1}+\alpha x_{B,n}{x_{B,n}}^\top\right)\Sigma_n x_{B,n} \quad \text{(EC.6)}\\&=\left(\frac{\gamma\ell}{{x_{B,n}}^\top\Sigma_n x_{B,n}+\alpha^{-1}}\right)^2\alpha\left(\alpha^{-1}+{x_{B,n}}^\top\Sigma_n x_{B,n}\right){x_{B,n}}^\top\Sigma_n x_{B,n}\\&=\frac{\alpha\gamma^2\ell^2{x_{B,n}}^\top\Sigma_n x_{B,n}}{{x_{B,n}}^\top\Sigma_n x_{B,n}+\alpha^{-1}}=\alpha\gamma^2\ell^2-\frac{\gamma^2\ell^2}{{x_{B,n}}^\top\Sigma_n x_{B,n}+\alpha^{-1}}.\end{aligned}$$

Next, we have

$$\begin{aligned}(\mu_B-\theta_n)^\top\Sigma_{n+1}^{-1}L&=(\mu_B-\theta_n)^\top\left(\Sigma_n^{-1}+\alpha x_{B,n}{x_{B,n}}^\top\right)\frac{\gamma\ell}{(\alpha^{-1}+{x_{B,n}}^\top\Sigma_n x_{B,n})}\Sigma_n x_n\\&=\frac{\gamma\ell}{\alpha^{-1}+{x_{B,n}}^\top\Sigma_n x_{B,n}}(\mu_B-\theta_n)^\top\left(\Sigma_n^{-1}+\alpha x_{B,n}{x_{B,n}}^\top\right)\Sigma_n x_{B,n}\\&=\frac{\gamma\ell}{(\alpha^{-1}+{x_{B,n}}^\top\Sigma_n x_{B,n})}(\mu_B-\theta_n)^\top x_{B,n}\left(1+\alpha{x_{B,n}}^\top\Sigma_n x_{B,n}\right)\\&=\alpha\gamma\ell(\mu_B-\theta_n)^\top x_n. \quad \text{(EC.7)}\end{aligned}$$

By using equations (EC.5), (EC.6), (EC.7) in (EC.4), we have

$$\begin{aligned}&(\mu_B-\theta_{n+1})^\top\Sigma_{n+1}^{-1}(\mu_B-\theta_{n+1})=(\mu_B-\theta_n)^\top\Sigma_n^{-1}(\mu_B-\theta_n)\\&\quad+\alpha\left[(\mu_B-\theta_n)^\top x_{B,n}{x_{B,n}}^\top(\mu_B-\theta_n)-2\ell\gamma(\mu_B-\theta_n)^\top x_{B,n}+\gamma^2\ell^2\right]-\frac{\gamma^2\ell^2}{{x_{B,n}}^\top\Sigma_n x_{B,n}+\alpha^{-1}}\\&=(\mu_B-\theta_n)^\top\Sigma_n^{-1}(\mu_B-\theta_n)+\alpha\left(\ell\gamma-(\mu_B-\theta_n)^\top x_{B,n}\right)^2-\frac{\gamma^2\ell^2}{{x_{B,n}}^\top\Sigma_n x_{B,n}+\alpha^{-1}},\end{aligned}$$

which leads to equation (EC.3). □

Next we apply Lemmas EC.1, EC.3, and EC.2 to prove Proposition 1.

Proof of Proposition 1. Since (μ_B,ρ) follows a multivariate normal-gamma distribution with parameters $(\theta_n,\Sigma_n,a_n,b_n)$, the joint density is given by

$$\begin{aligned}\Pr(\mu_B,\rho|S_n)&=\Pr(\mu_B|\rho,\theta_n,\Sigma_n)\ \Pr(\rho|a_n,b_n)\\&=\left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}}|\Sigma_n|^{-1/2}e^{-\frac{\rho}{2}(\mu_B-\theta_n)^\top\Sigma_n^{-1}(\mu_B-\theta_n)}\times\frac{b_n^{a_n}}{\Gamma(a_n)}\rho^{a_n-1}e^{-b_n\rho}, \quad \text{(EC.8)}\end{aligned}$$

where $\Gamma(\cdot)$ is the gamma function.

Let η_{n+1} and κ_{n+1} be the observations corresponding to the measurement $x_n=(x_{Z,n},x_{B,n})$. It follows from $\zeta\sim\mathcal{N}(\mu_Z,\Sigma_Z)$ and $\beta\sim\mathcal{N}(\mu_B,\Sigma_B)$, and the assumption that ζ, β, and ϵ are independent, that for every given $(x_{Z,n},x_{B,n})$,

$$\left\{\zeta^\top x_Z+\beta^\top x_B+\epsilon\mid\mu_B,\rho\right\}\ \sim\mathcal{N}\left(\mu_Z^\top x_{Z,n}+\mu_B^\top x_{B,n}\ ,\ \frac{1}{\rho}\left(1+x_{Z,n}^\top\Sigma_Z x_{Z,n}+x_{B,n}^\top\Sigma_B x_{B,n}\right)\right).$$

Define, $\hat{\rho} \stackrel{\text{def}}{=} \frac{1}{1+x_{Z,n}^{\top}\Sigma_Z x_{Z,n}+x_{B,n}^{\top}\Sigma_B x_{B,n}}\rho = \frac{1}{\hat{\sigma}_n}\rho$, where $\hat{\sigma}_n = 1 + x_{Z,n}^{\top}\Sigma_Z x_{Z,n} + x_{B,n}^{\top}\Sigma_B x_{B,n}$. Therefore, under model (2), we have

$$\left\{\eta_{n+1} = \kappa_{n+1}\left(\zeta^{\top}x_{Z,n} + \beta^{\top}x_{B,n} + \epsilon\right)|\mu_B, \rho, x_n, \kappa_{n+1}\right\} \sim \mathcal{N}\left(\kappa_{n+1}\left(\mu_Z^{\top}x_{Z,n} + \mu_B^{\top}x_{B,n}\right), \frac{\kappa_{n+1}^2}{\hat{\rho}}\right). \quad \text{(EC.9)}$$

Consequently, given $\kappa_{n+1} > 0$,

$$\Pr(\eta_{n+1}|\mu_B, \rho, x_n, \kappa_{n+1}) = \frac{\sqrt{\hat{\rho}}}{\sqrt{2\pi}\kappa_{n+1}} \exp\left(-\frac{\hat{\rho}}{2\kappa_{n+1}^2}\left(\eta_{n+1} - \kappa_{n+1}(\mu_Z^{\top}x_{Z,n} + \mu_B^{\top}x_{B,n})\right)^2\right). \quad \text{(EC.10)}$$

Bayes' rule implies that

$$\begin{aligned}
\Pr\left(\mu_B, \rho|S_n, x_n, \kappa_{n+1}, \eta_{n+1}\right) &\propto \Pr\left(\mu_B, \rho|S_n, x_n, \kappa_{n+1}\right) \times \Pr\left(\eta_{n+1}|\mu_B, \rho, S_n, x_n, \kappa_{n+1}\right) \\
&= \Pr\left(\mu_B, \rho|S_n\right) \times \Pr\left(\eta_{n+1}|\mu_B, \rho, x_n, \kappa_{n+1}\right) \qquad \text{(EC.11)} \\
&= \left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} |\Sigma_n|^{-1/2} e^{-\frac{\rho}{2}(\mu_B-\theta_n)^{\top}\Sigma_n^{-1}(\mu_B-\theta_n)} \times \frac{b_n^{a_n}}{\Gamma(a_n)}\rho^{a_n-1}e^{-b_n\rho} \qquad \text{(EC.12)} \\
&\times \frac{\sqrt{\hat{\rho}}}{\sqrt{2\pi}\kappa_{n+1}} e^{-\frac{\hat{\rho}}{2}\left(\ell+(\theta_n-\mu_B)^{\top}x_{B,n}\right)^2},
\end{aligned}$$

where $\ell \stackrel{\text{def}}{=} \frac{\eta_{n+1}}{\kappa_{n+1}} - (\mu_Z^{\top}x_{Z,n} + \theta_n^{\top}x_{B,n})$. The equality (EC.11) comes from the fact that the realized frequency κ_{n+1} does not impact the distribution of μ_B and ρ. Thus, $\Pr\left(\mu_B, \rho|S_n, x_n, \kappa_{n+1}\right) = \Pr\left(\mu_B, \rho|S_n\right)$. In addition, given μ_B and ρ, the description of η_{n+1} implies that $p\left(\eta_{n+1}|\mu_B, \rho, S_n, x_n, \kappa_{n+1}\right) = p\left(\eta_{n+1}|\mu_B, \rho, x_n, \kappa_{n+1}\right)$. Equation (53) comes from (EC.8) and (EC.10).

Given equation (9), Assumptions in Lemma EC.3 holds for $\gamma = 1$ and $\alpha = \hat{\sigma}^{-1}$. Therefore, it follows from Lemma EC.3 that

$$(\mu_B - \theta_{n+1})^{\top}\Sigma_{n+1}^{-1}(\mu_B - \theta_{n+1}) = (\mu_B - \theta_n)^{\top}\Sigma_n^{-1}(\mu_B - \theta_n) + \hat{\sigma}^{-1}\left(\ell - (\mu_B - \theta_n)^{\top}x_{B,n}\right)^2 - 2(b_{n+1} - b_n).$$

Here, we used equation (12) to get $2(b_{n+1} - b_n) = \frac{\gamma^2\ell^2}{x_{B,n}{}^{\top}\Sigma_n x_{B,n} + \alpha^{-1}}$. Therefore,

$$\begin{aligned}
&-\frac{\rho}{2}(\mu_B - \theta_n)^{\top}\Sigma_n^{-1}(\mu_B - \theta_n) - b_n\rho - \frac{\hat{\rho}}{2}\left(\ell - (\mu_B - \theta_n)^{\top}x_{B,n}\right)^2 \qquad \text{(EC.13)} \\
&= -\frac{\rho}{2}(\mu_B - \theta_{n+1})^{\top}\Sigma_{n+1}^{-1}(\mu_B - \theta_{n+1}) - b_{n+1}\rho.
\end{aligned}$$

Hence,

$$\begin{aligned}
&\left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} |\Sigma_n|^{-1/2} e^{-\frac{\rho}{2}(\mu_B-\theta_n)^{\top}\Sigma_n^{-1}(\mu_B-\theta_n)} \times \frac{b_n^{a_n}}{\Gamma(a_n)}\rho^{a_n-1}e^{-b_n\rho}\frac{\sqrt{c\rho}}{\sqrt{2\pi}\kappa_{n+1}}e^{-\frac{c\rho}{2}\left(\ell+(\theta_n-\mu_B)^{\top}x_{B,n}\right)^2} \\
&= \left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} \frac{|\Sigma_{n+1}|^{-\frac{1}{2}}}{\sqrt{1+\alpha x_n{}^{\top}\Sigma_n x_n}} e^{-\frac{\rho}{2}(\mu_B-\theta_{n+1})^{\top}\Sigma_{n+1}^{-1}(\mu_B-\theta_{n+1})} \times \frac{\sqrt{c}b_n^{a_n}}{\sqrt{2}\kappa_{n+1}B(a_n, \frac{1}{2})\Gamma(a_{n+1})}\rho^{a_{n+1}-1}e^{-b_{n+1}\rho} \\
&= c_0\left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} |\Sigma_{n+1}|^{-\frac{1}{2}} e^{-\frac{\rho}{2}(\mu_B-\theta_{n+1})^{\top}\Sigma_{n+1}^{-1}(\mu_B-\theta_{n+1})} \times \frac{b_{n+1}^{a_{n+1}}}{\Gamma(a_{n+1})}\rho^{a_{n+1}-1}e^{-b_{n+1}\rho}.
\end{aligned}$$

where we used equation (11), $a_{n+1} = a_n + \frac{1}{2}$, and the fact that $B(a_n, \frac{1}{2}) = \frac{\Gamma(a_n)\Gamma(\frac{1}{2})}{\Gamma(a_n+\frac{1}{2})} = \frac{\sqrt{\pi}\Gamma(a_n)}{\Gamma(a_{n+1})}$, where $B(a_n, \frac{1}{2})$ is the beta function. In the last equation, the term c_0 is given by

$$c_0 \stackrel{\text{def}}{=} \frac{\sqrt{c} b_n^{a_n}}{\sqrt{1+\alpha x_{B,n}{}^\top \Sigma_n x_{B,n}} b_{n+1}^{a_{n+1}} \sqrt{2} \kappa_{n+1} B(a_n, \frac{1}{2})}. \tag{EC.14}$$

Note that c_0 is constant with respect to ρ and μ_B. Thus, we arrive at

$$p\left(\mu_B, \rho | S_n, x_n, \kappa_{n+1}, \eta_{n+1}\right) \propto \left(\frac{\rho}{2\pi}\right)^{\frac{r}{2}} |\Sigma_{n+1}|^{-1/2} e^{-\frac{\rho}{2}(\mu_B - \theta_{n+1})^\top \Sigma_{n+1}^{-1}(\mu_B - \theta_{n+1})} \frac{b_{n+1}^{a_{n+1}}}{\Gamma(a_{n+1})} \rho^{a_{n+1}-1} e^{-b_{n+1}\rho},$$

which is the normal-gamma density with parameters $S_{n+1} = (\theta_{n+1}, \Sigma_{n+1}, a_{n+1}, b_{n+1})$. □

EC.2. Proof of Proposition 3

Proof. From the description of η_{n+1} in (2), we have

$$\{\eta_{n+1} \mid S_n, \rho, x_n, \mu_B, \kappa_{n+1}\} \sim \mathcal{N}\left(\kappa_{n+1}(\mu_Z^\top x_{Z,n} + \mu_B^\top x_{B,n}), \frac{\kappa_{n+1}^2}{\rho} \hat{\sigma}_n\right). \tag{EC.15}$$

Recall that $\hat{\sigma}_n = 1 + x_{Z,n}^\top \Sigma_Z x_{Z,n} + x_{B,n}^\top \Sigma_B x_{B,n}$. Thus from (EC.15) and $\mu_B | S_n, \rho \sim \mathcal{N}(\theta_n, \frac{1}{\rho}\Sigma_n)$, the random variable $\eta_{n+1} \mid S_n, \rho, x_n, \kappa_{n+1}$ follows a multiplicative model whose mean is distributed according to another Gaussian distribution. This implies that the Gaussian distribution

$$Z \stackrel{\text{def}}{=} \{\eta_{n+1} \mid S_n, \rho, x_n, \mu_B, \kappa_{n+1}\} - \kappa_{n+1}(\mu_Z^\top x_{Z,n} + \mu_B^\top x_{B,n}) \sim \mathcal{N}\left(0, \frac{\kappa_{n+1}^2}{\rho} \hat{\sigma}_n\right),$$

is independent of $\kappa_{n+1}(\mu_Z^\top x_{Z,n} + \mu_B^\top x_{B,n})$. Therefore,

$$\begin{aligned}\{\eta_{n+1} \mid S_n, \rho, x_n, \kappa_{n+1}\} &= Z + \kappa_{n+1}(\mu_Z^\top x_{Z,n} + \mu_B^\top x_{B,n}) \\ &\sim \mathcal{N}\left(0 + \kappa_{n+1}(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n}), \frac{\kappa_{n+1}^2}{\rho}\hat{\sigma}_n + \frac{\kappa_{n+1}^2}{\rho} x_{B,n}{}^\top \Sigma_n x_{B,n}\right).\end{aligned}$$

Given $k_{n+1} = k > 0$, we have

$$\begin{aligned}\Pr\left(\eta_{n+1} = \eta \mid S_n, x_n, \kappa_{n+1} = k\right) &= \int_0^\infty \Pr\left(\eta_{n+1} \mid \rho = r, S_n, x_n, \kappa_{n+1} = k\right) \Pr(\rho = r) dr \\ &= \int_0^\infty \frac{\sqrt{r}}{k\sqrt{2\pi\left(\hat{\sigma}_n + x_{B,n}^\top \Sigma_n x_{B,n}\right)}} e^{-\frac{r\left(\eta - k(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n})\right)^2}{2k^2(\hat{\sigma}_n + x_{B,n}{}^\top \Sigma_n x_{B,n})}} \times \frac{b_n^{a_n}}{\Gamma(a_n)} r^{a_n - 1} e^{-b_n r} dr.\end{aligned}$$

Define

$$\hat{\eta}_{n+1} \stackrel{\text{def}}{=} \sqrt{\frac{a_n}{\kappa^2 b_n(\hat{\sigma}_n + x_{B,n}{}^\top \Sigma_n x_{B,n})}} \left(\eta_{n+1} - \kappa(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n})\right). \tag{EC.16}$$

Therefore, given (S_n, x_n, κ_{n+1}), $\hat{\eta}_{n+1} = \hat{\eta}$ if and only if $\eta_{n+1} = g(\hat{\eta})$, where

$$g(\hat{\eta}) \stackrel{\text{def}}{=} k(\mu_Z^\top x_{Z,n} + \theta_n^\top x_{B,n}) + k\sqrt{\frac{b_n\left(\hat{\sigma}_n + x_{B,n}{}^\top \Sigma_n x_{B,n}\right)}{a_n}} \hat{\eta}.$$

Therefore, $\Pr(\hat{\eta}_{n+1} = \hat{\eta} | S_n, x_n, \kappa_{n+1}) = \Pr(\eta_{n+1} = g(\hat{\eta}) | S_n, x_n, \kappa_{n+1}) \times g'(\hat{\eta})$. Denote

$$\hat{\sigma} \stackrel{\text{def}}{=} \hat{\sigma}_n + {x_{B,n}}^\top \Sigma_n x_{B,n}. \tag{EC.17}$$

Hence, we have

$$\begin{aligned}
\Pr(\hat{\eta}_{n+1} = \hat{\eta} \mid S_n, x_n, \kappa_{n+1} = k) &= \int_0^\infty \frac{\sqrt{r}}{k\sqrt{2\pi\hat{\sigma}}} e^{-\frac{r\left(k\sqrt{\frac{b_n\hat{\sigma}}{a_n}}\hat{\eta}\right)^2}{2k^2\hat{\sigma}}} \times \frac{b_n^{a_n}}{\Gamma(a_n)} r^{a_n-1} e^{-b_n r} dr \times k\sqrt{\frac{b_n\hat{\sigma}}{a_n}} \\
&= \frac{1}{\Gamma(a_n)\sqrt{2a_n\pi}} \int_0^\infty e^{-rb_n\left(1+\frac{\hat{\eta}^2}{2a_n}\right)} b_n^{a_n+\frac{1}{2}} r^{a_n-\frac{1}{2}} dr \\
&= \frac{1}{\Gamma(a_n)\sqrt{2a_n\pi}} \left(1+\frac{\hat{\eta}^2}{2a_n}\right)^{-a_n-\frac{1}{2}} \int_0^\infty e^{-x} x^{a_n-\frac{1}{2}} dx \\
&= \frac{1}{\Gamma(a_n)\sqrt{2a_n\pi}} \left(1+\frac{\hat{\eta}^2}{2a_n}\right)^{-a_n-\frac{1}{2}} \Gamma\left(a_n+\frac{1}{2}\right).
\end{aligned}$$

The last line is the pdf of the standard Student's t-distribution with $\nu = 2a_n$ degrees of freedom. Hence, $\hat{\eta}_{n+1} | S_n, x_n, \kappa_{n+1} = k > 0$ follows a standard Student's t-distribution with $\nu = 2a_n$ degrees of freedom. This completes the proof of (16). $\square$

EC.3. Proof of Proposition 4

Proof. We prove by backward induction on n. For the function J as in (8), problem (18) is reduced to

$$\sup_{\pi\in\Pi} \mathbb{E}^\pi \left[\max_{x\in\mathcal{X}} \lambda^{\mathscr{M}}(x) \left(\mu_Z^\top x_Z + \theta_N^\top x_B\right)\right]. \tag{EC.18}$$

Recall that $\theta_N := \mathbb{E}[\mu_B | S_N]$.For $n = N-1$, for any $x_{N-1} \in \mathcal{X}$,

$$\begin{aligned}
Q_{N-1}(S_{N-1}, x_{N-1}) &= \mathbb{E}\left[V_N\left(S^{\mathscr{M}}\left(S_{N-1}, x_{N-1}, (\kappa_N, \eta_N)\right)\right)\right] \\
&= \mathbb{E}\left[\max_{y\in\mathcal{X}} \lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_N^\top y_B\right)\right] \\
&= \mathbb{E}\left[\max_{y\in\mathcal{X}} \lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_{N-1}^\top y_B + \sqrt{\frac{b_{N-1}}{a_{N-1}\hat{\sigma}}} y_B^\top \Sigma_{N-1} x_{B,N-1} T_{2a_{N-1}}\right)\right] \\
&\geq \max_{y\in\mathcal{X}} \mathbb{E}\left[\lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_{N-1}^\top y_B + \sqrt{\frac{b_{N-1}}{a_{N-1}\hat{\sigma}}} y_B^\top \Sigma_{N-1} x_{B,N-1} T_{2a_{N-1}}\right)\right] \\
&= \max_{y\in\mathcal{X}} \left\{\lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_{N-1}^\top y_B + \sqrt{\frac{b_{N-1}}{a_{N-1}\hat{\sigma}}} y_B^\top \Sigma_{N-1} x_{B,N-1} \mathbb{E}[T_{2a_{N-1}}]\right)\right\} \\
&= \max_{y\in\mathcal{X}} \lambda^{\mathscr{M}}(y)\left(\mu_Z^\top y_Z + \theta_{N-1}^\top y_B\right) = V_N(S_{N-1}).
\end{aligned}$$

Here, we used the notation $\hat{\sigma}$ defined in (EC.17), i.e., $\hat{\sigma} = \hat{\sigma}_{N-1}(x_{Z,N-1}, x_{B,N-1}) + x_{B,N-1}^\top \Sigma_{N-1} x_{B,N-1}$. In the last equation, we used the fact that $\mathbb{E}[T_{2a_{N-1}}] = 0$.

Next, we prove the induction step. Suppose the statement is true for $n+1$, we prove it for n. For any $x_n \in \mathcal{X}$,

$$
\begin{aligned}
Q_n(S_n, x_n) &= \mathbb{E}\left[V_{n+1}\left(S^{\mathscr{M}}\left(S_n, x_n, (\kappa_{n+1}, \eta_{n+1})\right)\right)\right], \\
&= \mathbb{E}\left[\max_{y \in \mathcal{X}} \ \mathbb{E}\left[V_{n+2}\left(S^{\mathscr{M}}\left(S^{\mathscr{M}}\left(S_n, x_n, (\kappa_{n+1}, \eta_{n+1})\right), y, (\kappa_y, \eta_y)\right)\right)\right]\right] && \text{(EC.19)} \\
&\geq \max_{y \in \mathcal{X}} \ \mathbb{E}\left[\mathbb{E}\left[V_{n+2}\left(S^{\mathscr{M}}\left(S^{\mathscr{M}}\left(S_n, x_n, (\kappa_{n+1}, \eta_{n+1})\right), y, (\kappa_y, \eta_y)\right)\right)\right]\right] && \text{(EC.20)} \\
&= \max_{y \in \mathcal{X}} \ \mathbb{E}\left[\mathbb{E}\left[V_{n+2}\left(S^{\mathscr{M}}\left(S^{\mathscr{M}}(S_n, y, (\kappa_y, \eta_y)), x_n, (\kappa_{n+1}, \eta_{n+1})\right)\right)\right]\right] && \text{(EC.21)} \\
&= \max_{y \in \mathcal{X}} \ \mathbb{E}\left[Q_{n+1}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right), x_n\right)\right]. && \text{(EC.22)}
\end{aligned}
$$

where (κ_y, η_y) is the observed outcome of the marketing campaign characterized by the input feature vector y. Equality (EC.19) comes from the optimality equation (21). In equation (EC.21), we use that the belief model and the induced transition function $S^{\mathscr{M}}$ satisfy

$$
S^{\mathscr{M}}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right), x, (\kappa_x, \eta_x)\right) = S^{\mathscr{M}}\left(S^{\mathscr{M}}\left(S_n, x, (\kappa_x, \eta_x)\right), y, (\kappa_y, \eta_y)\right), \ \forall \ y, x \in \mathcal{X}.
$$

The induction hypothesis implies that

$$
Q_{n+1}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right), x_n\right) \geq V_{n+2}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right)\right), \quad a.s., \quad \forall \ y \in \mathcal{X}.
$$

By applying this inequality in (EC.22), we arrive at

$$
\begin{aligned}
Q_n(S_n, x_n) &\geq \max_{y \in \mathcal{X}} \mathbb{E}\left[Q_{n+1}\left(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right), x_n\right)\right] \\
&\geq \max_{y \in \mathcal{X}} \mathbb{E}\left[V_{n+2}(S^{\mathscr{M}}\left(S_n, y, (\kappa_y, \eta_y)\right))\right] = V_{n+1}(S_n), \quad \forall \ x_n \in \mathcal{X}.
\end{aligned}
$$

This shows the validity of the statement for n, and completes the proof of the first part. When the one more marketing campaign can be measured according to the optimal policy π^*, we have $Q_n(S_n, x_n^*) = V_n(S_n)$. This along with the inequality $Q_n(S_n, x_n) \geq V_{n+1}(S_n)$ for any $x_n \in \mathcal{X}$, implies that $V_n(S_n) \geq V_{n+1}(S_n)$. □

EC.4. Proof of Proposition 5

Proof. It follows from equation (20) that $V_N^{\pi}(S_n) := U(\mathbb{E}[\mu_B | S_N = S_n]) = U(\theta_n)$. Therefore,

$$
\begin{aligned}
v_x^{\mathrm{KG}, n} &= \mathbb{E}\left[V_N\left(S^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) - V_N(S_n) \mid S_n, x\right] \\
&= \mathbb{E}\left[V_N\left(S^{\mathscr{M}}\left(S_n, x, (\kappa_{n+1}, \eta_{n+1})\right)\right) - U(\theta_n) \mid S_n, x\right] \\
&= \mathbb{E}\left[U\left(\theta^{\mathscr{M}}\left(S_n, x, (\kappa_{n+1}, \eta_{n+1})\right)\right) \mid S_n, x\right] - U(\theta_n). && \text{(EC.23)}
\end{aligned}
$$

where $\theta^{\mathscr{M}}(S_n, x, (\eta_{n+1}, \kappa_{n+1})) = \mathbb{E}\left[\mu_B \mid S_N = S^{\mathscr{M}}(S_n, x, (\eta_{n+1}, \kappa_{n+1}))\right]$. Therefore,

$$\begin{aligned}
&\mathbb{E}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x\right] \\
&= \mathbb{E}_{\kappa_{n+1}}\left[\mathbb{E}_{\eta_{n+1}}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x, \kappa_{n+1}\right]\right] \\
&= \Pr(\kappa_{n+1} = 0|x) \times \mathbb{E}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x, \kappa_{n+1} = 0\right] \\
&\quad + \Pr(\kappa_{n+1} > 0|x) \times \mathbb{E}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x, \kappa_{n+1} > 0\right] \\
&= e^{-\lambda^{\mathscr{M}}(x)} U(\theta_n) + (1 - e^{-\lambda^{\mathscr{M}}(x)}) \mathbb{E}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x, \kappa_{n+1} > 0\right]. \quad \text{(EC.24)}
\end{aligned}$$

Here, we used the observation that $\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1})) = \theta_n$, given $\kappa_{n+1} = 0$. Consequently $\mathbb{E}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x, \kappa_{n+1} = 0\right] = U(\theta_n)$.

From Proposition 3 and given $\kappa_{n+1} > 0$, η_{n+1} follows the Student's t-distribution. In particular, using equation (17), we have

$$\theta^{\mathscr{M}}(S_n, x, (\eta_{n+1}, \kappa_{n+1})) = \theta_{n+1} = \theta_n + \sqrt{\frac{b_n}{a_n(\hat{\sigma}_n(x) + x_B^\top \Sigma_n x_B)}} \Sigma_n x_B T_{2a_n}.$$

Hence,

$$\begin{aligned}
&\mathbb{E}\left[U\left(\theta^{\mathscr{M}}(S_n, x, (\kappa_{n+1}, \eta_{n+1}))\right) \mid S_n, x, \kappa_{n+1} > 0\right] \\
&= \mathbb{E}\left[U\left(\theta_n + \sqrt{\frac{b_n}{a_n(\hat{\sigma}_n(x) + x_B^\top \Sigma_n x_B)}} \Sigma_n x_B T_{2a_n}\right) \mid S_n, x, \kappa_{n+1} > 0\right] \\
&= \mathbb{E}\left[\max_{y \in \mathcal{X}} \lambda(y) \left(\mu_Z^\top y_Z + \left(\theta_n + \sqrt{\frac{b_n}{a_n(\hat{\sigma}_n(x) + x_B^\top \Sigma_n x_B)}} \Sigma_n x_B T_{2a_n}\right)^\top y_B\right) \mid S_n, x, \kappa_{n+1} > 0\right] \\
&= \mathbb{E}\left[\max_{y \in \mathcal{X}} \left(p_y^n + q_y^n(x) T_{2a_n}\right) \mid S_n, x, \kappa_{n+1} > 0\right], \quad \text{(EC.25)}
\end{aligned}$$

where p_y^n and $q_y^n(x_n)$ are given in equations (27) and (28). Applying equation (56) in (55), and finally in equation (54) completes the derivation of (26). □

EC.5. Proof of Proposition 9

Proof. For Σ_{n+1} as in equation (10), and under the assumption that Σ_n is positive definite, we have (e.g., see Boyd (2004)):

$$\begin{aligned}
\log\det(\Sigma_{n+1}) &= \log\det\left(\Sigma_n - \frac{\Sigma_n x_{B,n} x_{B,n}{}^\top \Sigma_n}{x_n^\top Q_n x_n}\right) \\
&= \log\det\left(\Sigma_n^{1/2}\left(I - \frac{1}{x_n^\top Q_n x_n} \Sigma_n^{-1/2} \Sigma_n x_{B,n} x_{B,n}{}^\top \Sigma_n \Sigma_n^{-1/2}\right) \Sigma_n^{1/2}\right) \\
&= \log\left(1 - \frac{1}{x_n^\top Q_n x_n} \lambda_{\max}\left(\Sigma_n^{1/2} x_{B,n} x_{B,n}{}^\top \Sigma_n^{1/2}\right)\right) + \log\det(\Sigma_n) \\
&= \log\left(1 - \frac{1}{x_n^\top Q_n x_n} x_{B,n}{}^\top \Sigma_n x_{B,n}\right) + \log\det(\Sigma_n).
\end{aligned}$$

Here, we used the fact that $\Sigma_n^{1/2} x_{B,n} {x_{B,n}}^\top \Sigma_n^{1/2}$ is a matrix of rank one and thus has at most one non-zero eigenvalue ${x_{B,n}}^\top \Sigma_n x_{B,n}$. Since at stage n, Σ_n is a constant and log is an increasing function, to minimize $\log\det(\Sigma_{n+1})$ it is sufficient to maximize $\frac{{x_{B,n}}^\top \Sigma_n x_{B,n}}{x_n^\top Q_n x_n}$. By introducing $V = x_{B,n} {x_{B,n}}^\top$, we have ${x_{B,n}}^\top \Sigma_n x_{B,n} = \mathrm{tr}(\Sigma_n x_{B,n} {x_{B,n}}^\top) = \mathrm{tr}(\Sigma_n V)$. Thus, we arrive at minimizing z such that $\frac{x_n^\top Q_n x_n}{\mathrm{tr}(\Sigma_n V)} \leq z$. The recent constraint is equivalent to the second-order cone constraint

$$\left\| \begin{matrix} 2Q_n^{1/2} x_n \\ \mathrm{tr}(\Sigma_n V) - z \end{matrix} \right\|_2 \leq \mathrm{tr}(\Sigma_n V) + z.$$

□