

Chen, Le-yu; Lee, Sokbae

Working Paper

Best subset binary prediction

cemmap working paper, No. CWP50/17

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Chen, Le-yu; Lee, Sokbae (2017) : Best subset binary prediction, cemmap working paper, No. CWP50/17, Centre for Microdata Methods and Practice (cemmap), London, <https://doi.org/10.1920/wp.cem.2017.5017>

This Version is available at:

<https://hdl.handle.net/10419/189780>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Best subset binary prediction

Le-Yu Chen
Sokbae Lee

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP50/17

Best Subset Binary Prediction*

Le-Yu Chen[†]

Institute of Economics, Academia Sinica

Sokbae Lee[‡]

Department of Economics, Columbia University

Centre for Microdata Methods and Practice, Institute for Fiscal Studies

November 2017

Abstract

We consider a variable selection problem for the prediction of binary outcomes. We study the best subset selection procedure by which the explanatory variables are chosen by maximizing Manski (1975, 1985)'s maximum score type objective function subject to a constraint on the maximal number of selected variables. We show that this procedure can be equivalently reformulated as solving a mixed integer optimization (MIO) problem, which enables computation of the exact or an approximate solution with a definite approximation error bound. In terms of theoretical results, we obtain non-asymptotic upper and lower risk bounds when the dimension of potential covariates is possibly much larger than the sample size. Our upper and lower risk bounds are minimax rate-optimal when the maximal number of selected variables is fixed and does not increase with the sample size. We illustrate usefulness of the best subset binary prediction approach via Monte Carlo simulations and an empirical application of the work-trip transportation mode choice.

*We are grateful to the co-editor, Jianqing Fan, an associate editor and three anonymous referees for constructive comments and suggestions. We also thank Stefan Hoderlein, Joel Horowitz, Shakeeb Khan, Toru Kitagawa, Arthur Lewbel, and participants in 2017 Asian Meeting of the Econometric Society and 2017 annual conference of the International Association for Applied Econometrics for helpful comments. This work was supported by the Ministry of Science and Technology, Taiwan (MOST106-2410-H-001-015-), Academia Sinica (Career Development Award research grant), the European Research Council (ERC-2014-CoG-646917-ROMIA), the UK Economic and Social Research Council (ES/P008909/1 via CeMMAP), and the British Academy (International Partnership and Mobility Scheme Grant, reference number PM140162).

[†]E-mail: lychen@econ.sinica.edu.tw

[‡]E-mail: sl3841@columbia.edu

Keywords: *binary choice, maximum score estimation, best subset selection, ℓ_0 -constrained maximization, mixed integer optimization, minimax optimality, finite sample property*

JEL codes: C52, C53, C55

1 Introduction

Prediction of binary outcomes is an important topic in economics and various scientific fields. Let $Y \in \{0, 1\}$ be the binary outcome of interest and W a vector of covariates for predicting Y . Assume that the researcher has a training sample of n independent identically distributed (i.i.d.) observations $(Y_i, W_i)_{i=1}^n$ of (Y, W) . For $w \in \mathcal{W}$, let

$$b_\theta(w) \equiv 1 \{w'\theta \geq 0\}, \quad (1.1)$$

where \mathcal{W} is the support of W_i , θ is a vector of parameters, and $1 \{\cdot\}$ is an indicator function that takes value 1 if its argument is true and 0 otherwise.

One reasonable prediction rule is to choose θ such that it maximizes the probability of making the correct prediction $P[Y = b_\theta(W)]$. However, this is infeasible in practice since the joint distribution of (Y, W) is unknown. A natural sample analog is to maximize the sample average score which equals the proportion of correct predictions under the prediction rule (1.1) in the training sample. This maximization problem is equivalent to the maximum score estimation in binary response models and is pioneered by Manski (1975, 1985). Thus, we call the corresponding prediction rule the *maximum score* prediction rule. See Manski and Thompson (1989), Jiang and Tanner (2010), and Elliott and Lieli (2013) for prediction in the maximum score approach.

This prediction problem has the same structure as the binary classification problem, which is extensively studied in the statistics and machine learning literature. For example, see the classic work of Devroye, Györfi, and Lugosi (1996) among many others. In this literature, the empirical risk minimization (ERM) classifier over the class of binary classifiers specified by (1.1) is defined as a minimizer of $1 - S_n(\theta)$. In other words, the ERM classification rule is identical to the maximum score prediction rule.

In this paper, we address the covariate selection issue in the framework of predicting the binary outcome Y_i using the class of linear threshold-crossing prediction rules $b_\theta(W_i)$ defined by (1.1). We study the *best subset* selection procedure by which the covariates are chosen among a collection of candidate explanatory variables by maximizing the empirical score subject to a constraint on the maxi-

mal number of selected variables. In other words, we investigate theoretical and numerical properties of the ℓ_0 -norm constrained maximum score prediction rules.¹

To the best of our knowledge, Greenshtein (2006) and Jiang and Tanner (2010) are the only existing papers in the literature that explicitly considered the same prediction problem as ours. Greenshtein (2006) considered a general loss function that includes maximum score prediction as a special case in the i.i.d. setup. Greenshtein (2006) focused on a high dimensional case and established conditions under which the excess risk converges to zero as $n \rightarrow \infty$. Jiang and Tanner (2010) focused on the prediction of time series data and obtained an upper bound for the excess risk. Neither Greenshtein (2006) nor Jiang and Tanner (2010) provided any numerical results for the best subset maximum score prediction rule. In contrast, we focus on cross-sectional applications and emphasize computational aspects.

The main contributions of this paper are twofold: first, we show that the best subset maximum score prediction rule is minimax rate-optimal and second, we demonstrate that it can be implemented via mixed integer optimization. The first contribution is theoretical and builds on the literature of empirical risk minimization (Tsybakov, 2004; Massart and Nédélec, 2006, in particular). Specifically, we obtain *non-asymptotic* upper and lower risk bounds when the dimension of potential covariates is possibly much larger than the sample size n . Our upper and lower risk bounds are minimax rate-optimal when the maximal number of selected variables is fixed and does not increase with n . The existing results of finite-sample upper and lower risk bounds for the binary prediction problem focus on the case where there is no variable selection and the set of covariates is fixed and low-dimensional. Our risk bound results extend to the setup under the ℓ_0 -norm constraint when the set of potential covariates is high-dimensional.²

The second contribution is computational. We face two kinds of computational challenges. One challenge comes from the nature of the objective function and the other is from the best subset selection. The maximum score objective function is a piecewise constant function whose range set contains only finitely many points. Hence the maximum of the score maximization problem is always attained yet the maximizer is generally not unique. It is known that computing the maximum score estimates regardless of the presence of the ℓ_0 constraint is NP (non-deterministic polynomial-time)-hard (see, e.g., Johnson and Preparata, 1978). See Manski and

¹Here, the ℓ_0 -norm of a real vector refers to the number of non-zero components of the vector.

²Raskutti, Wainwright, and Yu (2011) developed minimax rate results for high-dimensional linear mean regressions. We have used in the derivation of our lower risk bound a technical lemma of their paper (Raskutti, Wainwright, and Yu, 2011, Lemma 4), which is based on the approximation theory literature. Nonetheless, Our results are not directly obtainable from Raskutti, Wainwright, and Yu (2011), who considered the least squares type objective function.

Thompson (1986) and Pinkse (1993) for first generation algorithms for maximum score estimation.

Our computation algorithm is based on the method of mixed integer optimization (MIO). Florios and Skouras (2008) provided compelling numerical evidence that the MIO approach is superior to the first-generation approaches. Kitagawa and Tetenov (2015) used a MIO formulation that is different from Florios and Skouras (2008) to solve maximum score type problems. The objective of interest in Kitagawa and Tetenov (2015) is to develop treatment choice rules by maximizing an empirical welfare criterion, which resembles the maximum score objective function. They derived minimax optimality and used the MIO formulation to implement their algorithm. Neither Florios and Skouras (2008) nor Kitagawa and Tetenov (2015) were concerned with the variable selection problem.

These second generation approaches are driven by developments in MIO solvers and also by availability of a much faster computer compared to the period when the first generation algorithms were proposed. In fact, fast developments in computing environments have made the MIO approach even more attractive now than 2000s. Florios and Skouras (2008) reported that they obtained the exact maximum score estimates using Horowitz (1993)'s data in 10.5 hours. In this application, the sample size was $n = 842$ and there were 4 parameters to estimate.

It is well known that use of a good and tighter parameter space can strengthen the performance of a global optimization procedure including the MIO approach. In this paper, we propose a data driven approach to refine the parameter space. Using a state-of-the-art MIO solver as well as a tailor-made heuristic to choose the parameter space, it took us less than 5 minutes to obtain the exact maximum score estimates using the same dataset with the same number of parameters to estimate.³ This is a dramatic improvement at the factor of more than 100 relative to the numerical performance reported in Florios and Skouras (2008). In other words, we demonstrate that hardware improvements combined with the advances in MIO solvers and also with a carefully chosen parameter space have made the maximum score approach empirically much more relevant now than ten years ago.

The second numerical challenge is concerned with constrained optimization with the ℓ_0 -norm constraint. It is well known that the ℓ_0 -norm constraint renders the variable selection problem NP-hard even in the regression setup where the objective function is convex and smooth (see, e.g., Natarajan, 1995; Bertsimas, King, and Mazumder, 2016). Recently Bertsimas, King, and Mazumder (2016) proposed a novel MIO approach to the best subset variable selection problem when least squares and least absolute deviation risks are concerned. They demonstrated

³This numerical result can be found in Online Appendix E of the paper.

that the MIO approach can efficiently deliver a provably optimal solution to the resulting ℓ_0 -norm constrained risk minimization problem for a variety of datasets with practical problem size. Our implementation of the best subset maximum score prediction rules combines insights from Bertsimas, King, and Mazumder (2016), Florios and Skouras (2008), and Kitagawa and Tetenov (2015). We present two alternative MIO solution methods that complement each other.

In practical applications, it is useful to consider an approximate solution by adopting an early termination rule. In our empirical application, by setting an explicitly pre-specified optimization error, we were able to obtain approximate maximum score estimates with Horowitz (1993)'s data in around 10 minutes when both an intercept term and one specific random covariate are always selected but 5 auxiliary variables were selected out of 9 potential covariates. This suggests that fast developments in computing environments will enable us to solve an empirical problem at a practically relevant scale in very near future. We provide additional numerical evidence in Monte Carlo experiments in a high-dimensional setup when the number of potential covariates is larger than the sample size.

The remainder of this paper is organized as follows. In Section 2, we describe our prediction rule. Section 3 establishes theoretical properties of the proposed prediction rule. In Section 4, we present computation algorithms using the MIO approach, and in Section 5, we conduct a simulation study on the performance of our prediction rule in both low and high dimensional variable selection problems. In Section 6, we illustrate usefulness of our prediction rule in the empirical application of work-trip mode choice using Horowitz (1993)'s data. We then conclude the paper in Section 7. Appendix A contains proofs of all theoretical results. Supplementary material of this paper are collated in online appendices.

2 A Best Subset Approach to Maximum Score Prediction of Binary Outcomes

In this section, we describe our proposal of the best subset maximum score prediction rule. Following Magnus and Durbin (1999) and Danilov and Magnus (2004), we distinguish between *focus covariates* that are always included in the prediction rule and *auxiliary covariates* of which we are less certain. We thus decompose the covariate vector W as $W = (X, Z)$, where X is a $(k + 1)$ -dimensional vector of focus covariates and Z is a p -dimensional vector of auxiliary covariates.

Noting that $b_{a\theta}(w) = b_{\theta}(w)$ for any positive real scalar a , we adopt the same scale normalization method as in Horowitz (1992) and Jiang and Tanner (2010)

by restricting the magnitude of the coefficient of one of the focus covariates to be unity. Specifically, write $X = (X_0, \tilde{X})$ where X_0 is a scalar variable and \tilde{X} is the remaining k -dimensional subvector of focus covariates. The parameter vector θ in (1.1) is decomposed accordingly as $\theta = (\alpha, \beta, \gamma)$, where $\alpha \in \{-1, 1\}$ and $(\beta, \gamma) \in \Theta$,⁴ which is a subset of \mathbb{R}^{k+p} . In this notation, the binary prediction rule has the following form:

$$b_{\alpha, \beta, \gamma}(w) = 1\{\alpha x_0 + \tilde{x}'\beta + z'\gamma \geq 0\} \text{ for } w \in \mathcal{W}. \quad (2.1)$$

We consider a parsimonious variable selection method by which the constituted prediction rule does not include more than a pre-specified number of auxiliary covariates. For any p dimensional real vector c , let $\|c\|_0 \equiv \sum_{j=1}^p 1\{c_j \neq 0\}$ be the ℓ_0 -norm of c . We carry out the ℓ_0 -norm constrained covariate selection procedure by solving the constrained maximization problem

$$\max_{(\alpha, \beta, \gamma) \in \{-1, 1\} \times \Theta_q} S_n(\alpha, \beta, \gamma), \quad (2.2)$$

where the objective function S_n is defined as

$$S_n(\theta) \equiv n^{-1} \sum_{i=1}^n 1\{Y_i = b_\theta(W_i)\} \quad (2.3)$$

and the ℓ_0 -norm constrained parameter space is given as

$$\Theta_q \equiv \{(\beta, \gamma) \in \Theta \subset \mathbb{R}^{k+p} : \|\gamma\|_0 \leq q\} \quad (2.4)$$

for a given positive integer q .

As we discussed in the introduction, solving for the exact maximizer for (2.2) is desirable yet can be computationally challenging. It is hence practically useful to consider an approximate solution, which is constructed below, to the maximization problem (2.2).

For any $\varepsilon \geq 0$, let $(\hat{\alpha}, \hat{\beta}, \hat{\gamma}) \in \{-1, 1\} \times \Theta_q$ be an approximate maximizer with ε tolerance level such that

$$S_n(\hat{\alpha}, \hat{\beta}, \hat{\gamma}) \geq \max_{(\alpha, \beta, \gamma) \in \{-1, 1\} \times \Theta_q} S_n(\alpha, \beta, \gamma) - \varepsilon \text{ almost surely.} \quad (2.5)$$

We refer to the prediction rule defined by $1\{\hat{\alpha}x_0 + \tilde{x}'\hat{\beta} + z'\hat{\gamma} \geq 0\}$ as the approximate best subset maximum score binary prediction rule.⁵ The value of ε can be specified for early termination of the solution algorithm to the problem (2.2). In Section 4,

⁴The parameter space of θ is defined as $\{-1, 1\} \times \Theta$, which is slightly incoherent notation.

⁵The dependence of $(\hat{\alpha}, \hat{\beta}, \hat{\gamma})$ on ε is suppressed for simplicity of notation.

we will present an algorithm that allows for computing an approximate solution to (2.2) within a definite approximation error bound specified by the tolerance level ε . In what follows, we use PRESCIENCE as shorthand for the approximate best subset maximum score binary prediction rule.⁶

Remark 1. In terms of model selection, there are several aspects one needs to consider. First, one needs to specify the covariate vector W . We recommend starting with a large set of covariates for W since we have a built-in model selection procedure. Second, it is necessary to decide which covariates belong to X (focus covariates) and Z (auxiliary covariates). What consists of auxiliary covariates depends on particular applications. The auxiliary covariates correspond to the part of the model specification the researcher is not sure about. For example, they could be some higher order terms or interaction terms. If a researcher does not have concrete ideas about the auxiliary covariates, we recommend letting the auxiliary covariates be all regressors except one the researcher is specifically interested in. Third, it is required to choose q (the ℓ_0 -norm constraint). The constant q is an important tuning parameter in our procedure. A particular choice of q can be motivated in some specific applications. Generally speaking, for the purpose of prediction, there is the standard tradeoff between flexibility, which require a larger q , and the risk of over-fitting, which pushes for a smaller q . We recommend using cross validation to choose q , as we will demonstrate in our empirical example and Monte Carlo experiments.

3 Theoretical Properties of PRESCIENCE

In this section, we study the theoretical properties of PRESCIENCE. Let F denote the joint distribution of (Y, W) . For $(\alpha, \beta, \gamma) \in \{-1, 1\} \times \Theta$, let

$$S(\alpha, \beta, \gamma) \equiv P(Y = b_{\alpha, \beta, \gamma}(W)). \quad (3.1)$$

Note that $S(\alpha, \beta, \gamma)$ depends on the joint distribution F . Given a cardinality bound q , let

$$S_q^* \equiv \sup_{(\alpha, \beta, \gamma) \in \{-1, 1\} \times \Theta_q} S(\alpha, \beta, \gamma). \quad (3.2)$$

That is, S_q^* is the supremum of $S(\alpha, \beta, \gamma)$ given the ℓ_0 -norm constraint.

Following the literature on empirical risk minimization (see, e.g., Devroye, Györfi, and Lugosi (1996), Lugosi (2002), Tsybakov (2004), Massart and Nédélec (2006), Greenshtein (2006) and Jiang and Tanner (2010) among many others), we

⁶It comes from the aPpRoximate bEst S(C)ubset maxImum scorE biNary prediCtion rule.

assess the predictive performance of PRESCIENCE by bounding the difference

$$U_n \equiv S_q^* - S(\widehat{\alpha}, \widehat{\beta}, \widehat{\gamma}), \quad (3.3)$$

where $(\widehat{\alpha}, \widehat{\beta}, \widehat{\gamma})$ is defined by (2.5) and depends on the given tolerance level ε . The difference U_n is non-negative by the definition of S_q^* . Hence, a good prediction rule will result in a small value of U_n with a high probability and also on average.

Throughout this section, we assume that $\varepsilon = 0$ for simplicity. Before presenting our theoretical results, we first introduce some notation. For any two real numbers a and b , let $a \vee b \equiv \max\{a, b\}$. Let $s \equiv k + q$ and

$$r_n \equiv s \ln(p \vee n) \vee 1. \quad (3.4)$$

Theorem 1. *Assume $s \geq 1$. Then for all $\sigma > 0$, there is a universal constant M_σ , which depends only on σ , such that*

$$P \left(U_n > 2\sqrt{\frac{M_\sigma r_n}{n}} \right) \leq \exp(-\sigma r_n), \quad (3.5)$$

provided that

$$(4s + 4) \ln(M_\sigma r_n) \leq r_n + (6s + 5) \ln 2. \quad (3.6)$$

Theorem 1 establishes that the tail probability of U_n decays exponentially in r_n . Moreover, this result is non-asymptotic: inequality (3.5) is valid for every sample size n for which condition (3.6) holds. By comparing the leading terms on both sides of inequality (3.6), we can see that condition (3.6) is satisfied, for instance, if

$$4[\ln(s) + \ln(\ln(p \vee n)) + \ln(M_\sigma)] \leq \frac{1}{2} \ln(p \vee n). \quad (3.7)$$

Hence, condition (3.6) is satisfied easily when $p \vee n$ takes a relatively large value compared to s . If $(p \vee n)$ diverges to infinity, then s can diverge at a sufficiently slow rate.

Theorem 1 implies that

$$E[U_n] = O \left(n^{-1/2} \sqrt{s \ln(p \vee n)} \right) = o(1), \quad (3.8)$$

provided that

$$s \ln(p \vee n) = o(n) \quad (3.9)$$

holds. This allows the case that

$$\ln p = O(n^\alpha) \text{ and } s = o(n^{1-\alpha}) \text{ for } 0 < \alpha < 1. \quad (3.10)$$

In other words, the predictive performance of PRESCIENCE remains good even when the number of potentially relevant covariates (p) grows exponentially, provided that the number of selected covariates (s) can only grow at a polynomial rate. Greenshtein and Ritov (2004) and Greenshtein (2006) consider the case where p grows at a polynomial rate. For this case, condition (3.9) implies that $s = o(n/\ln n)$, which coincides with the optimal sparsity rate established by Greenshtein and Ritov (2004) and Greenshtein (2006) under which a sequence of predictor selection procedures subject to the sparsity constraint can be shown to be persistent.

Remark 2. For the case with $\varepsilon > 0$, it is straightforward to modify theoretical results presented above such that the rate result (3.8) continues to hold provided that

$$\varepsilon = O\left(n^{-1/2}\sqrt{s \ln(p \vee n)}\right). \quad (3.11)$$

3.1 An Upper Bound under the Margin Condition

The result obtained in the previous subsection is derived under the i.i.d. setup but does not hinge on other regularity conditions on the underlying data generating distribution F . The result can be improved under additional assumptions on the distribution F . In this section, we consider a condition that is called the *margin condition* in the literature under which we may obtain a sharper result on the upper bound of $E[U_n]$. As before, the derived bound will be non-asymptotic.

It is necessary to introduce additional notation. Let

$$\mathcal{B}_q \equiv \{b_\theta : \theta \in \{-1, 1\} \times \Theta_q\}. \quad (3.12)$$

That is, \mathcal{B}_q is the class of all prediction rules in (2.1) with the ℓ_0 -norm constraint. For $w \in \mathcal{W}$, let

$$\eta(w) \equiv P(Y = 1|W = w), \quad (3.13)$$

$$b^*(w) \equiv 1 \{\eta(w) \geq 0.5\}. \quad (3.14)$$

For any measurable function $f : \mathcal{W} \mapsto \mathbb{R}$, let $\|f\|_1 = E[|f(W)|]$ denote the L_1 -norm of f . The functions η and b^* as well as the L_1 -norm $\|\cdot\|_1$ depend on the data

generating distribution F . For any indicator function $b : \mathcal{W} \mapsto \{0, 1\}$, let

$$\tilde{S}(b) \equiv P(Y = b(W)). \quad (3.15)$$

We now state the following regularity condition.

Condition 1 (Margin Condition). *There are some $\vartheta \geq 1$ and $h > 0$ such that, for every binary predictor $b : \mathcal{W} \mapsto \{0, 1\}$,*

$$\tilde{S}(b^*) - \tilde{S}(b) \geq h^\vartheta \|b^* - b\|_1^\vartheta. \quad (3.16)$$

Condition 1 is termed as the margin condition in the literature (see, e.g., Mammen and Tsybakov (1999), Tsybakov (2004) and Massart and Nédélec (2006)). For any binary predictor b ,

$$\tilde{S}(b^*) - \tilde{S}(b) = E[|2\eta(W) - 1| |b^*(W) - b(W)|], \quad (3.17)$$

so that $\tilde{S}(b)$ is maximized at $b = b^*$. Hence, Condition 1 implies that the functional $\tilde{S}(\cdot)$ has a well-separated maximum. See Proposition 1 of Tsybakov (2004) and the following discussions on the margin condition. Suppose that there exist universal positive constants C and α such that

$$P\{|\eta(W) - 1/2| \leq t\} \leq Ct^\alpha$$

for all $t > 0$. Then by modifying the proof of Proposition 1 of Tsybakov (2004) slightly, we can show that (3.16) holds with $\vartheta = (1 + \alpha)/\alpha$.

Recall that it is not necessary to assume (3.16) to establish the risk consistency, as shown in Theorem 1. We show below that we can obtain a tighter upper bound on $E[U_n]$ under (3.16). Let

$$\rho_n \equiv 1 \vee [\ln 2 + q \ln p + (s + 1) \ln(n + 1)]. \quad (3.18)$$

The next theorem, which is an application of Massart and Nédélec (2006, Theorem 2), establishes a finite-sample bound on $E[U_n]$ under the margin condition.

Theorem 2. *There are universal constants K and K' such that*

$$E[U_n] \leq \left[\tilde{S}(b^*) - \sup_{b \in \mathcal{B}_q} \tilde{S}(b) \right] + K' \left(\frac{K^2 \rho_n}{nh} \right)^{\vartheta/(2\vartheta-1)}, \quad (3.19)$$

provided that Condition 1 holds with

$$h \geq \left(\frac{K^2 \rho_n}{n} \right)^{\frac{1}{2\vartheta}}. \quad (3.20)$$

For $p \vee n$ sufficiently large, we have that $\rho_n \leq 5s \ln(p \vee n)$; thus, inequality (3.20) can hold under condition (3.9) in large samples, provided that h is fixed or does not go to zero too rapidly.

The first term on the right-hand side of inequality (3.19) represents the bias term. Equation (3.17) implies that there is no bias term, namely $\tilde{S}(b^*) = \sup_{b \in \mathcal{B}_q} \tilde{S}(b)$ if $b^* \in \mathcal{B}_q$. Therefore, Theorem 2 implies that

$$E[U_n] = O \left(\left[\frac{s \ln(p \vee n)}{nh} \right]^{\vartheta/(2\vartheta-1)} \right), \quad (3.21)$$

provided that $b^* \in \mathcal{B}_q$.⁷ The rate of convergence in (3.21) doubles that in (3.8) when h is fixed and $\vartheta = 1$. We notice that, if $b^* \notin \mathcal{B}_q$, the upper bound derived in Theorem 2 would asymptotically reduce to the non-zero bias term $\tilde{S}(b^*) - \sup_{b \in \mathcal{B}_q} \tilde{S}(b)$ and hence the margin condition alone does not suffice for deducing there is improved rate of convergence. Nevertheless, the rate result (3.8) still holds regardless of the validity of the presumption that $b^* \in \mathcal{B}_q$.

We now remark on the condition that $b^* \in \mathcal{B}_q$ in the context of the binary response model specified below. Suppose that the outcome Y is generated from a latent variable threshold crossing model (see, e.g., Manski, 1975, 1985):

$$Y = 1\{W'\theta^* \geq \xi\}, \quad (3.22)$$

where θ^* denotes the true data generating parameter vector and ξ is an unobserved latent variable whose distribution satisfies that

$$\text{Med}(\xi|W = w) = 0 \text{ for } w \in \mathcal{W}. \quad (3.23)$$

Let $\theta_0 \equiv \arg \sup_{\theta \in \{-1,1\} \times \Theta_q} S(\theta)$. For simplicity, assume that $\theta_0 \in \{-1,1\} \times \Theta_q$ so that the maximum is attained.

⁷For the case with $\varepsilon > 0$, it is also straightforward to modify Theorem 2 such that the rate result (3.21) continues to hold provided that

$$\varepsilon = O \left(\left[\frac{s \ln(p \vee n)}{nh} \right]^{\vartheta/(2\vartheta-1)} \right).$$

Proposition 1. *Assume that the model given by (3.22) and (3.23) is correctly specified. Suppose that Condition 1 holds. Then $b^* \in \mathcal{B}_q$ if and only if $W'\theta_0$ and $W'\theta^*$ have the same sign with probability 1.*

Manski (1988, Proposition 2) showed that, for the binary response model specified by (3.22) and (3.23), the true parameter value θ^* is identified relative to another value θ if and only if the event that $W'\theta$ and $W'\theta^*$ have different sign occurs with positive probability. Therefore, Proposition 1 implies that $b^* \in \mathcal{B}_q$ if and only if the “pseudo-true” value θ_0 is observationally equivalent to θ^* . In particular, this implies that $\theta^* = \theta_0$ if θ^* is point-identified.

It would be interesting to study the role of the bias when $b^* \neq \mathcal{B}_q$ using the framework of sieve estimation (Chen, 2007). As pointed by Elliott and Lieli (2013, Proposition 1), what matters is how well we can approximate the value of optimum $\sup_{\theta \in \{-1,1\} \times \Theta_q} S(\theta)$, not the optimizer $\arg \sup_{\theta \in \{-1,1\} \times \Theta_q} S(\theta)$. However, it would be much more demanding to develop non-asymptotic theory when the bias is present in our framework. We leave this as a topic for future research.

3.2 A Minimax Lower Bound under the Margin Condition

In this section, we derive a minimax lower bound under the margin condition. In particular, we focus on the case that $s = k + q$ is low-dimensional in that s does not grow with sample size n and also consider a sufficient condition for the margin condition.

Condition 1 is satisfied with $\vartheta = 1$ whenever

$$|2\eta(w) - 1| \geq h \text{ for } w \in \mathcal{W}. \quad (3.24)$$

Massart and Nédélec (2006) introduced (3.24) as an easily interpretable margin condition requiring that the conditional probability $\eta(w)$ should be bounded away from 1/2. Condition (3.24) holds under certain regularity assumptions on the binary response model as indicated in the following proposition.

Proposition 2. *Assume that the model given by (3.22) and (3.23) is correctly specified. Suppose that there are universal constants $\kappa_1 > 0$, $\kappa_2 > 0$ such that (i) $P(|W'\theta^*| \geq \kappa_1) = 1$ and (ii) there is some open interval T containing $(-\kappa_1, \kappa_1)$ such that $P(\xi \leq t | W = w)$ has a derivative (with respect to t) which is bounded below by κ_2 for every $t \in T$. Then condition (3.24) holds with $h = 2\kappa_1\kappa_2$.*

Conditions (i) and (ii) in Proposition 2 assume that $|W'\theta^*|$ is bounded away from zero and the density of ξ conditional on $W = w$ is bounded away from zero

in a neighborhood of zero. While the latter condition is mild, the former is non-trivial. Condition (i) can hold easily when all components of W are discrete, which is not uncommon in microeconomic applications of binary response models (see e.g., Komarova (2013) and Magnac and Maurin (2008)). In the presence of continuous covariates, this condition becomes more restrictive.

For any real vector u , let $\|u\|_E = \sqrt{u'u}$ denote the Euclidean norm of u . To state a minimax lower bound, we first define the following class of distributions.

Definition 1. For every $h \in (0, 1)$, let $\mathcal{P}(h, \mathcal{B}_q)$ denote the class of distributions F satisfying the following conditions: (i) $b^* \in \mathcal{B}_q$, (ii) condition (3.24) holds, and (iii) there are constants $c_u > 0$ and $c_l > 0$ such that, for any two vectors $\theta = (\alpha, \beta, \gamma)$, $\tilde{\theta} = (\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma}) \in \{-1, 1\} \times \Theta_q$ satisfying $\alpha = \tilde{\alpha}$ and $\beta = \tilde{\beta}$, it holds that

$$c_l \|\theta - \tilde{\theta}\|_E \leq \|b_\theta - b_{\tilde{\theta}}\|_1 \leq c_u \|\theta - \tilde{\theta}\|_E. \quad (3.25)$$

The first two conditions in the definition of $\mathcal{P}(h, \mathcal{B}_q)$ are already introduced in the previous subsection. The new condition (iii) imposes that the Euclidean norm $\|\theta - \tilde{\theta}\|_E$ is equivalent to the L_1 -norm $\|b_\theta - b_{\tilde{\theta}}\|_1$ for two values θ and $\tilde{\theta}$ that differ only in the components corresponding to the auxiliary covariate coefficients. This condition is concerned with restrictions on the distribution of the covariate vector W . The following proposition gives sufficient conditions for verifying this norm equivalence condition.

For any subset $J \subset \{1, \dots, p\}$, let Z_J denote the $|J|$ -dimensional subvector of $Z \equiv (Z^{(1)}, \dots, Z^{(p)})'$ formed by keeping only those elements $Z^{(j)}$ with $j \in J$. Let $\mathcal{I}_q \equiv \cup_{(\beta, \gamma) \in \Theta_q} \text{Supp}(\tilde{X}\beta + Z'\gamma)$, where $\text{Supp}(V)$ denotes the support of the random variable V .

Proposition 3. Suppose that s is fixed and does not grow with sample size n . Assume that there are positive real constants L_1 , L_2 and L_3 such that (a) the distribution of X_0 conditional on (\tilde{X}, Z) has a Lebesgue density that is bounded above by L_1 and bounded below by L_1^{-1} on \mathcal{I}_q , and (b) for any subset $J \subset \{1, \dots, p\}$ such that $|J| \leq 2q$, $P(\|Z_J\|_E \leq L_2) = 1$ and the smallest eigenvalue of $E(Z_J Z_J')$ is bounded below by L_3 . Then Condition (iii) stated in (3.25) holds with $c_u = L_1 L_2$ and $c_l = (L_1 L_2)^{-1} L_3$.

Condition (a) in Proposition 3 is mild. The first part of condition (b) holds with $L_2 = \bar{L}\sqrt{2q}$ if $\max_{j \in \{1, \dots, p\}} |Z^{(j)}| \leq \bar{L}$ with probability 1 for some universal positive constant \bar{L} . The second part of condition (b) is related to the sparse eigenvalue assumption used in the high dimensional regression literature (see, e.g. Raskutti, Wainwright, and Yu (2011)). For example, suppose that Z is a random vector

with mean zero and the variance-covariance matrix Σ whose (i, j) component is $(\Sigma)_{(i,j)} = \rho^{|i-j|}$ for some constant $\rho > 0$. Then the smallest eigenvalue of Σ is $1 - \rho$, independent of the dimension p . Thus, in this case, $E(Z_J Z_J')$ is bounded below by $1 - \rho$.

We now state the result on the minimax lower bound for the predictive performance of PRESCIENCE.

Theorem 3. *Assume the parameter space Θ in (2.4) satisfies that there is a universal constant $\kappa > 0$ such that*

$$\max_{j \in \{1, \dots, p\}} |\gamma_j| \leq \kappa, \quad (3.26)$$

where γ_j denotes the j th component of γ . Suppose p and q are even numbers and $q < 2p/3$. Let $\phi \equiv 0.71$. Then, for any binary predictor $\hat{b} : \mathcal{W} \mapsto \{0, 1\}$, which is in the set \mathcal{B}_q and is constructed based on the data $(Y_i, W_i)_{i=1}^n$, we have that

$$\sup_{F \in \mathcal{P}(h, \mathcal{B}_q)} E \left[S_q^* - \tilde{S}(\hat{b}) \right] \geq \frac{\phi q c_l (1 - \phi) (1 - h)}{32 n h c_u} \ln \left(\frac{p - q}{q/2} \right) \quad (3.27)$$

for $h \in (0, 1)$, which is defined in (3.24), such that

$$h \geq \left[\frac{\phi \sqrt{q} \ln \left(\frac{p - q}{q/2} \right)}{8 \sqrt{2} \kappa n c_u} \right]^{1/2}. \quad (3.28)$$

For any estimator \hat{b} taking value in \mathcal{B}_q , Theorem 3 implies that, as long as $\mathcal{P}(h, \mathcal{B}_q)$ is non-empty, there is some distribution F under which the average predictive risk $E \left[S_q^* - \tilde{S}(\hat{b}) \right]$ cannot be smaller than the lower bound term stated in (3.27).

Suppose that (s, c_l, c_u) are fixed and does not increase or decrease with n . Then this lower bound term is of order

$$O \left(\frac{(1 - h) \ln p}{n h} \right). \quad (3.29)$$

Comparing (3.29) to (3.21) evaluated at $\vartheta = 1$, we see that, if h is also a universal constant and p grows at a polynomial or exponential rate in n , then the upper and lower bound results induce the same convergence rate and hence the binary prediction method based on PRESCIENCE is rate-optimal in the minimax sense.

Remark 3. This minimax rate optimality is established under the assumption that s is fixed. Theorem 3 does not provide a rate-optimal lower bound when s diverges to infinity as $n \rightarrow \infty$, although it is a valid lower bound in any finite sample. It is

an interesting open question for future research to investigate minimax optimality when $s \rightarrow \infty$.

Remark 4. The assumption that p and q are even in Theorem 3 is innocuous for the minimax rate-optimality result. This assumption is made to invoke the known result (see Lemma 4 of Raskutti, Wainwright, and Yu (2011)) for the lower bound on the complexity of the ℓ_0 -ball. When p and/or q is odd, the lower bound result still holds since we can always consider $\mathcal{P}(h, \mathcal{B}'_q) \subset \mathcal{P}(h, \mathcal{B}_q)$, where \mathcal{B}'_q is a subspace of \mathcal{B}_q for which the parameter vector θ is confined to a lower dimensional space with dimension $p - 1$ and/or $q - 1$.

4 Implementation via Mixed Integer Optimization

We now present algorithms for solving the maximization problem (2.2). It is straightforward to see that solving (2.2) is the same as solving

$$\max \left\{ \max_{(\beta, \gamma) \in \Theta_q} S_n(1, \beta, \gamma), \max_{(\beta, \gamma) \in \Theta_q} S_n(-1, \beta, \gamma) \right\}.$$

In what follows, we focus on solving the sub-problem

$$\max_{(\beta, \gamma) \in \Theta_q} S_n(1, \beta, \gamma) \tag{4.1}$$

because the other case corresponding to $\alpha = -1$ can be solved by replacing the value of X_{0i} with that of $-X_{0i}$ and then applying the same solution method as developed for the case (4.1).

By (2.1) and noting that $Y_i \in \{0, 1\}$, solving the problem (4.1) amounts to solving

$$\max_{(\beta, \gamma) \in \Theta_q} n^{-1} \sum_{i=1}^n \left[(1 - Y_i) + (2Y_i - 1) 1\{X_{0i} + \tilde{X}'_i \beta + Z'_i \gamma \geq 0\} \right]. \tag{4.2}$$

We assume that the parameter space Θ is bounded and takes the polyhedral form:

$$\Theta = \{(\beta, \gamma) \in \mathbb{R}^{k+p} : A_1 \beta + A_2 \gamma \leq B\}$$

for some real constant matrices A_1 and A_2 and some real constant vector B . Let

$$\mathbf{C} \equiv \prod_{j=1}^p \left[\underline{\gamma}_j, \bar{\gamma}_j \right] \tag{4.3}$$

denote the smallest cube containing all values of γ in the pair (β, γ) confined by Θ .

Writing $\gamma = (\gamma_1, \dots, \gamma_p)$, we have that, if $(\beta, \gamma) \in \Theta$, then $-\infty < \underline{\gamma}_j \leq \gamma_j \leq \bar{\gamma}_j < \infty$ for $j \in \{1, \dots, p\}$. Let

$$M_i \equiv \max_{(\beta, \gamma) \in \Theta} \left| X_{0i} + \tilde{X}'_i \beta + Z'_i \gamma \right| \text{ for } i \in \{1, \dots, n\}. \quad (4.4)$$

Our implementation builds on the method of mixed integer optimization (in particular, Bertsimas, King, and Mazumder (2016), Florios and Skouras (2008), and Kitagawa and Tetenov (2015)) and present two alternative solution methods that complement each other. The values $(M_i)_{i=1}^n$ can be computed by formulating the maximization problem in (4.4) as linear programming problems, which can be easily and efficiently solved by modern numerical software such as MATLAB. Hence these values can be computed and stored beforehand as inputs to the algorithms that are used to solve the MIO problems described below.

4.1 Method 1

Our first solution method is based on an equivalent reformulation of the maximization problem (4.2) as the following constrained mixed integer optimization (MIO) problem:

$$\max_{(\beta, \gamma) \in \Theta, d_1, \dots, d_n, e_1, \dots, e_p} n^{-1} \sum_{i=1}^n [(1 - Y_i) + (2Y_i - 1) d_i] \quad (4.5)$$

subject to

$$(d_i - 1) M_i \leq X_{0i} + \tilde{X}'_i \beta + Z'_i \gamma < d_i (M_i + \delta), \quad i \in \{1, \dots, n\}, \quad (4.6)$$

$$e_j \underline{\gamma}_j \leq \gamma_j \leq e_j \bar{\gamma}_j, \quad j \in \{1, \dots, p\}, \quad (4.7)$$

$$\sum_{j=1}^p e_j \leq q, \quad (4.8)$$

$$d_i \in \{0, 1\}, \quad i \in \{1, \dots, n\}, \quad (4.9)$$

$$e_j \in \{0, 1\}, \quad j \in \{1, \dots, p\}, \quad (4.10)$$

where δ is a given small and positive real scalar (e.g. $\delta = 10^{-6}$ as in our numerical study).

We now explain the equivalence between (4.2) and (4.5). Given (β, γ) , the inequality constraints (4.6) and the dichotomization constraints (4.9) enforce that $d_i = 1\{X_{0i} + \tilde{X}'_i \beta + Z'_i \gamma \geq 0\}$ for $i \in \{1, \dots, n\}$. Therefore, maximizing the objective function in (4.2) for $(\beta, \gamma) \in \Theta$ subject to the constraints (4.6) and (4.9) is equivalent to solving the problem (4.2) using all covariates. This part of formulation is similar to the MIO formulation used by Kitagawa and Tetenov (2015, Section 5.1) for solving the maximum score type estimation problems without the variable

selection constraint.

Following Bertsimas, King, and Mazumder (2016), we implement the best subset variable selection feature through the additional constraints (4.7), (4.8) and (4.10). The on-off constraints (4.7) and (4.10) ensure that, whenever $e_j = 0$, the auxiliary covariate Z_j is excluded in the resulting PRESCIENCE. Finally, the cardinality constraint $\|\gamma\|_0 \leq q$ is enacted through the constraint (4.8), which restricts the maximal number of the binary controls e_j that can take value unity.

Modern numerical optimization solvers such as CPLEX, Gurobi, MOPS, Mosek and Xpress-MP can be used to effectively solve the MIO formulations (4.5) and (4.14) of the PRESCIENCE problem. Most of the solution algorithms employed by the MIO solvers can be viewed as complex and advanced refinements of the well-known branch-and-bound method for solving MIO problems.⁸ Along the branch-and-bound solution process, we can keep track of two important values: the best upper and lower bounds on the objective value of the MIO problem (4.5). The best lower bound corresponds to the objective function evaluated at the incumbent solution, which is the best feasible solution discovered so far. The best upper bound can be deduced by taking the maximum of the optimal objective values of all the linear programming relaxation formulations of the branching MIO sub-problems that have been solved so far. Let *MIO_gap* denote the difference between these two bounds. Note that the incumbent solution becomes optimal when the *MIO_gap* value reduces to zero.

We can use the *MIO_gap* value to solve for the ε -level PRESCIENCE introduced in Section 2. To see this, consider an early termination rule by which the solution algorithm is terminated whenever $MIO_gap \leq \varepsilon$ where ε is a given tolerance level. Let $(\hat{\beta}, \hat{\gamma})$ be the incumbent solution upon termination of the MIO solver. Because $(\hat{\beta}, \hat{\gamma})$ is in the feasible solution set of the problem (4.5), by constraints (4.7) and (4.8), we have that $\|\hat{\gamma}\|_0 \leq q$. Moreover, by constraints (4.6) and (4.9), we have that $\hat{d}_i = 1\{X_{0i} + \tilde{X}'_i \hat{\beta} + Z'_i \hat{\gamma} \geq 0\}$ for $i \in \{1, \dots, n\}$ so that $S_n(1, \hat{\beta}, \hat{\gamma})$ is equal to the objective function in (4.5) evaluated at $(\hat{d}_1, \dots, \hat{d}_n)$. Since (4.2) and (4.5) are equivalent maximization problems, it thus follows from the construction of *MIO_gap* value that

$$S_n(1, \hat{\beta}, \hat{\gamma}) \geq \max_{(\beta, \gamma) \in \Theta_q} S_n(1, \beta, \gamma) - MIO_gap. \quad (4.11)$$

Given the termination condition, we can therefore see that

$$S_n(1, \hat{\beta}, \hat{\gamma}) \geq \max_{(\beta, \gamma) \in \Theta_q} S_n(1, \beta, \gamma) - \varepsilon, \quad (4.12)$$

⁸See Online Appendix B for further details of the branch-and-bound method.

which yields an approximately optimal solution with the optimization tolerance level ε for the problem (4.1).

The PRESCIENCE can also be solved by incorporating the constraints (4.7), (4.8) and (4.10) in the MIO formulation of Florios and Skouras (2008) for solving the maximum score estimation problem. We now present this alternative MIO formulation below.

4.2 Method 2

Consider the constrained maximization problem:

$$\max_{(\beta, \gamma) \in \Theta_q} n^{-1} \sum_{i=1}^n 1 \left\{ (2Y_i - 1) (X_{0i} + \tilde{X}'_i \beta + Z'_i \gamma) \geq 0 \right\}. \quad (4.13)$$

The problem (4.13) without the constraint $\|\gamma\|_0 \leq q$ reduces to the type of maximum score estimation problem studied by Florios and Skouras (2008). The objective function in (4.13) coincides with that in (4.1) with probability 1 as long as the sum $X_0 + \tilde{X}'\beta + Z'\gamma$ is continuously distributed. This condition holds provided that the distribution of X_0 conditional on (\tilde{X}, Z) is continuous. With such a continuous covariate, we can also solve (4.1) by solving the following MIO formulation of (4.13):

$$\max_{(\beta, \gamma) \in \Theta, d_1, \dots, d_n, e_1, \dots, e_p} n^{-1} \sum_{i=1}^n d_i \quad (4.14)$$

subject to the constraints (4.7), (4.8), (4.9), (4.10), and

$$(1 - 2Y_i) (X_{0i} + \tilde{X}'_i \beta + Z'_i \gamma) \leq M_i (1 - d_i), \quad i \in \{1, \dots, n\}. \quad (4.15)$$

Florios and Skouras (2008) showed that maximizing the objective function in (4.14) for $(\beta, \gamma) \in \Theta$ subject to the constraints (4.9) and (4.15) is equivalent to solving the problem (4.13) using all covariates. This can be seen from the fact that the objective function of the MIO problem (4.14) is strictly increasing in d_i so that, given (β, γ) , it is optimal to set $d_i = 1 \{ (2Y_i - 1) (X_{0i} + \tilde{X}'_i \beta + Z'_i \gamma) \geq 0 \}$ under the constraints (4.9) and (4.15). Along similar arguments to those discussed for the problem (4.5), it is also straightforward to verify that the variable selection constraint $\|\gamma\|_0 \leq q$ is imposed through the constraints (4.7), (4.8) and (4.10). Therefore, the maximization problems (4.13) and (4.14) are equivalent.

4.3 Tightening the Parameter Space as a Warm Start to the MIO Formulation of the PRESCIENCE Problem

The MIO formulations (4.5) and (4.14) depend on the specification of the parameter space Θ . It is well known that use of a good and tighter parameter space can strengthen the performance of a global optimization procedure. Given an initial specification of Θ , we propose below a data driven approach to refine the parameter space.

Recall that $W_i = (X_{0i}, \tilde{X}_i, Z_i)$ is the entire covariate vector. Let \tilde{W}_i denote the vector (\tilde{X}_i, Z_i) . For $i \in \{1, \dots, n\}$, let \hat{P}_i be an estimate of $P_i \equiv P(Y_i = 1|W_i)$. Define the following sets recursively:

$$\Theta_1 \equiv \Theta, \bar{\Theta}_1 \equiv \left\{ (t_1, \dots, t_{k+p}) \in \Theta : t_1 \geq \hat{l}_1 \right\}$$

and, for $m \in \{2, \dots, k+p\}$,

$$\underline{\Theta}_m \equiv \left\{ (t_1, \dots, t_{k+p}) \in \Theta : \hat{l}_s \leq t_s \leq \hat{u}_s \text{ for } s \in \{1, \dots, m-1\} \right\}, \quad (4.16)$$

$$\bar{\Theta}_m \equiv \left\{ (t_1, \dots, t_{k+p}) \in \underline{\Theta}_m : t_m \geq \hat{l}_m \right\}. \quad (4.17)$$

where, for $j \in \{1, \dots, k+p\}$, the quantities \hat{l}_j and \hat{u}_j are defined respectively by

$$\hat{l}_j \equiv \min_{t \in \Theta_j} t_j \text{ subject to} \quad (4.18)$$

$$(X_{0i} + \tilde{W}_i' t)(\hat{P}_i - 0.5) \geq 0 \text{ for } i \in \{1, \dots, n\}. \quad (4.19)$$

$$\hat{u}_j \equiv \max_{t \in \bar{\Theta}_j} t_j \text{ subject to the constraints (4.19)}. \quad (4.20)$$

If the binary outcome Y_i is generated from the model specified by (3.22) and (3.23) and the conditional probability $P(Y = 1|W)$ is nonparametrically estimated, the interval $[\hat{l}_j, \hat{u}_j]$ is a nonparametric estimate of the identified set for the j th component of the parameter vector $t = (\beta, \gamma)$. In this case, the sign-matching constraints (4.19) can be regarded as the empirical counterparts of the inequalities stated in the set

$$\left\{ t \in \Theta : (X_{0i} + \tilde{W}_i' t)(P(Y_i = 1|W_i) - 0.5) \geq 0 \text{ almost surely} \right\},$$

which contains those t values that are observationally equivalent to the true data generating parameter value (see, e.g. Komarova, 2013; Chen and Lee, 2015). Our procedures for computing \hat{l}_j and \hat{u}_j are modified versions of Horowitz (1998, p. 62)'s linear programming formulations of the identified bounds on the parameter

components. The formulations (4.18) and (4.20) differ from those of Horowitz in that we further tighten the domain of t in these optimization problems by exploiting the information of the upper and lower bound values that have been solved so far.

When the covariate vector is of high dimension, nonparametric estimation of $P(Y = 1|W)$ would suffer from the curse of dimensionality problem. In this paper, we consider estimating this conditional probability by parametric methods such as the logit or probit approach. In Monte Carlo experiments and an empirical application, we estimate $(P_i)_{i=1}^n$ by the fitted choice probabilities from the logit regression of Y on all the covariates.

Noting that the parametric model for estimating $P(Y = 1|W)$ may be misspecified, for $\tau \geq 1$, we construct a conservative space $\widehat{\Theta}(\tau)$, which is a τ -enlargement of the space $\prod_{j=1}^{k+p} [\widehat{l}_j, \widehat{u}_j]$ as given below:

$$\widehat{\Theta}(\tau) \equiv \left\{ t \in \Theta : -\tau \left(|\widehat{l}_j| \vee |\widehat{u}_j| \right) \leq t_j \leq \tau \left(|\widehat{l}_j| \vee |\widehat{u}_j| \right) \text{ for } j \in \{1, \dots, k+p\} \right\}.$$

We can solve the MIO problems (4.5) and (4.14) with the refined parameter space $\widehat{\Theta}(\tau)$ in place of the original space Θ . Using the terminology used in Bertsimas, King, and Mazumder (2016), we shall refer to these refined MIO representations as the *warm-start* MIO formulations of the PRESCIENCE problem. The value of τ is treated as a tuning parameter for solving the warm-start MIO problems. The original formulations (4.5) and (4.14) based on the space Θ are referred to as the *cold-start* MIO formulations.

Computation of the refined space $\widehat{\Theta}(\tau)$ requires solving $2(k+p)$ simple linear programming problems.⁹ This task can be done very efficiently even when p is relatively large. On the other hand, the space $\widehat{\Theta}(\tau)$ is not always constructible since the problems (4.18) and (4.20) may not admit any feasible solution. This may occur due to the misspecification issue of using parametric choice probability estimates. Alternatively, it can also occur when the postulated binary response model specified by (3.22) and (3.23) itself is misspecified. As illustrated by Monte Carlo simulations and a real data application in Online Appendices D and E, when the refined space $\widehat{\Theta}(\tau)$ is available, solving the warm-start MIO formulations can be computationally far more efficient than solving their corresponding cold-start versions.

We conclude this subsection by commenting that our warm-start approach does

⁹Both the covering cube \mathbf{C} and the quantities $(M_i)_{i=1}^n$ depend on the input parameter space. Hence, for the warm-start formulations of (4.5) and (4.14), these objects are also computed under the refined space $\widehat{\Theta}(\tau)$.

not work well when p is greater than n . In this case, irrespective of the knowledge of the true choice probabilities, the dimension of the vector of unknown coefficients is larger than the number of inequalities given by the constraints (4.19) such that these constraints may become ineffective for tightening the original parameter bounds. It is a topic for future research how to devise a good warm-start option for the high-dimensional setup.

5 Simulation Study

In this section, we study the performance of the PRESCIENCE method in Monte Carlo experiments. Throughout this paper, we used the MATLAB implementation of the Gurobi Optimizer to solve the MIO problems for all computations. Moreover, all numerical computations were done on a desktop PC (Windows 7) equipped with 32 GB RAM and a CPU processor (Intel i7-5930K) of 3.5 GHz.¹⁰

Let $V = (V_1, \dots, V_{p+1})$ be a multivariate normal random vector with mean zero and covariance matrix Σ with its element $\Sigma_{i,j} = (0.25)^{|i-j|}$. The binary outcome is generated according to the following setup:

$$Y = 1\{W'\theta^* \geq \varepsilon\},$$

where θ^* denotes the value of the true data generating parameter vector, $W = (X, Z)$ is a $(p+2)$ dimensional covariate vector with the focus covariates $X = (X_0, \tilde{X}) = (V_1, 1)$ and the auxiliary covariates $Z = (V_2, \dots, V_{p+1})$,

$$\varepsilon = 0.25\sigma(W)\xi,$$

and ξ is a $N(0, 1)$ random variate independent of V . We set $\theta_1^* = 1$, $\theta_2^* = 0$, and $\theta_j^* = 0$ for $j \in \{4, \dots, p+2\}$. The coefficient θ_3^* is chosen to be non-zero such that, among the p auxiliary covariates, only the variable Z_1 is relevant in the data generating processes (DGP).

We consider the following two specifications for θ_3^* and $\sigma(W)$:

$$\text{DGP(i)} : \theta_3^* = -0.35 \text{ and } \sigma(W) = 1.$$

$$\text{DGP(ii)} : \theta_3^* = -1.5 \text{ and } \sigma(W) = 1 + 2(V_1 + V_2)^2 + (V_1 + V_2)^4.$$

¹⁰The MATLAB codes for implementing the PRESCIENCE approach are available from the authors via the website <https://github.com/LeyuChen/Best-Subset-Binary-Prediction>. This implementation requires the Gurobi Optimizer, which is freely available for academic purposes.

As before, the parameter vector θ in (1.1) is decomposed as $\theta = (\alpha, \beta, \gamma)$ where α , β and γ are coefficients associated with X_0 , \tilde{X} and Z , respectively. The parameter space for the PRESCIENCE approach is specified to be

$$\{(\alpha, \beta, \gamma) \in \mathbb{R}^{p+2} : \alpha = 1, (\beta, \gamma) = (\beta, \gamma_1, \dots, \gamma_p) \in [-10, 10]^{p+1}\} \quad (5.1)$$

over which we compute PRESCIENCE via solving its corresponding MIO problem.

There were 100 simulation repetitions in each Monte Carlo experiment. For each simulation repetition, we generated a training sample of n observations for estimating the coefficients θ and a validation sample of 5000 observations for evaluating the out-of-sample predictive performance. The training sample size n was set to be 100 for DGP(i) and 50 for DGP(ii). For each DGP setup, we performed simulations with both the low and high dimensional covariate configurations. For the low dimensional case, we set $p = 10$ for both DGP(i) and (ii). For the high dimensional case, we set $p = 200$ for DGP(i) and $p = 60$ for DGP(ii).

Let M be a given class of prediction methods. For each $m \in M$, let $\hat{\theta}(m)$ denote the coefficients computed under the prediction method m . Let *in_Score* denote the average of the in-sample objective values $S_n(\hat{\theta}(m))$ over all the simulation repetitions. In each simulation repetition, we approximated the out-of-sample objective value $S(\hat{\theta}(m))$ using the generated validation sample. Let *out_Score* denote the average of $S(\hat{\theta}(m))$ over all the simulation repetitions. It is straightforward to see that the theoretically best prediction rule $b^*(w)$ in this simulation design takes the form $b^*(w) = 1 \{w'\theta^* \geq 0\}$. Hence, we also assess the predictive performance of a given prediction method $m \in M$ by its relative score, which is ratio of the score evaluated at $\hat{\theta}(m)$ over that evaluated at θ^* . Let *in_RS* and *out_RS* respectively denote the average of in-sample relative scores $S_n(\hat{\theta}(m))/S_n(\theta^*)$ and that of out-of-sample relative scores $S(\hat{\theta}(m))/S(\theta^*)$ over all the simulation repetitions.

We also examine the variable selection performance of the prediction method. We say that a variable Z_j is effectively selected under the prediction method m if and only if the magnitude of $\hat{\theta}_{j+2}(m)$ is larger than a small tolerance level (e.g. 10^{-6} as used in our numerical study) which is distinct from zero in numerical computation. Let *Corr_sel* be the proportion of the auxiliary covariate Z_1 being effectively selected. Let *Orac_sel* be the proportion of obtaining an oracle variable selection outcome where, among all the auxiliary covariates, Z_1 was the only one that was effectively selected. Let *Num_irrel* denote the average number of effectively selected auxiliary covariates whose true DGP coefficients are zero.

We considered the following prediction methods:

$$\mathcal{M} = \{\{\text{PRESCIENCE}(q), q \in \{1, 2, 3\}\}, \text{PRESCIENCE_CV}, \text{logit_lasso}, \text{probit_lasso}\}, \quad (5.2)$$

where $\text{PRESCIENCE}(q)$ denotes the PRESCIENCE approach with a cardinality bound q imposed on the auxiliary covariates, PRESCIENCE_CV denotes the PRESCIENCE approach using a data driven value of $q \in \{1, 2, 3\}$ via the 5-fold cross validation procedure, and logit_lasso and probit_lasso respectively denote the ℓ_1 -penalized logit and probit maximum likelihood estimation (MLE) approaches (see e.g. Friedman, Hastie, and Tibshirani, 2010). Throughout this simulation study, we employed the cold-start MIO formulation (4.5) to solve the PRESCIENCE problems. For the simulation experiment with $p < n$, we computed the exact solution to each PRESCIENCE problem. For the high dimensional case with $p > n$, we solved for the PRESCIENCE solution with the tolerance level ε specified according to the rule

$$\varepsilon = \min\{0.05, 0.5\sqrt{\ln(p \vee n)/n}\}. \quad (5.3)$$

Note that this early termination rule is compatible with the order of magnitude stated in the condition (3.11) for the convergence rate result (3.8). For the logit_lasso and probit_lasso approaches, we used the MATLAB function **lassoglm** to implement these two penalized MLE approaches for which we calibrated the optimal lasso penalty parameter value among a sequence of 100 values via the 10-fold cross validation procedure. We used the default setup of **lassoglm** for constructing this tuning sequence.

5.1 Simulation Results for the Homoskedastic Error Design

We now present the simulation results under the setup of DGP(i), where the distribution of the error term ε is independent of the covariates W . First, we report the computational performance of our MIO solution algorithm to the $\text{PRESCIENCE}(q)$ problem. Table 1 gives the summary statistics of the MIO computation time in CPU seconds across simulation repetitions. From this table, we can see that the MIO problems for the PRESCIENCE computation were solved very efficiently in the DGP(i) simulations where the number of the auxiliary covariates could be the double of the sample size yet the maximum computation time was only around 5 minutes. It is also interesting to note that the PRESCIENCE computation time was not monotone in q . This feature might be due to the branching strategy

heuristics of the MIO branch-and-bound solution algorithms.

Table 1: PRESCIENCE(q) computation time (CPU seconds) under DGP(i)

q	$p = 10$			$p = 200$		
	1	2	3	1	2	3
mean	0.69	1.10	0.51	7.11	20.37	1.96
min	0.01	0.01	0.01	0.31	0.16	0.07
median	0.35	0.38	0.33	3.76	2.09	0.99
max	3.51	24.56	5.42	68.04	362.7	15.95

We next turn to the statistical performance of the binary prediction method. In Tables 2 and 3, we compare the aforementioned predictive and variable selection performance measures for the various prediction approaches given in (5.2). As shown in these two tables, regardless of p , the in-sample fit in terms of in_Score and in_RS for the PRESCIENCE(q) method increased with q . This finding is expected because the in-sample objective function (2.3) is monotone in q by design for the PRESCIENCE approach. Nonetheless, both tables indicate that out_Score and out_RS also declined as q increased, thus resounding with the known issue that in-sample overfitting may result in poor out-of-sample performance. When the true number of effective auxiliary covariates is unknown, one can use a data-driven choice of q via the cross validation procedure. From Tables 2 and 3, we find that the PRESCIENCE_CV approach indeed balanced well the in-sample and out-of-sample predictive performances. Moreover, its predictive performance measures were also comparable to those given by the `logit_lasso` and `probit_lasso` approaches.

We now discuss the variable selection results. Table 2 indicates that all the six prediction approaches had high $Corr_sel$ rates and hence were capable of effectively selecting the relevant covariate Z_1 . However, the good performance in the $Corr_sel$ criterion may come at the cost of overfitting. Therefore, we also have to take into account the performance in excluding irrelevant auxiliary covariates. The simulation design implies that the case with $q = 1$ is the most parsimonious PRESCIENCE setup that correctly specifies the number of effective auxiliary covariates in the DGP. Therefore, it is not surprising that PRESCIENCE(1) performed the best in terms of Num_irrel . We note that the PRESCIENCE_CV approach also performed very well in excluding the irrelevant variables. In fact, only the PRESCIENCE(1) and PRESCIENCE_CV approaches could yield a non-zero probability of inducing an oracle variable selection outcome whereas PRESCIENCE(q) with $q > 1$ and the `logit_lasso` and `probit_lasso` approaches always included irrele-

vant auxiliary covariates.

We also observe a similar pattern in the results of the DGP(i) setup with $p = 200$. From Table 3, we find that the PRESCIENCE_CV approach also balanced very well the requirement for including the relevant but excluding the irrelevant variables. It is also noted that the logit_lasso and probit_lasso approaches in this high dimensional simulation setup tended to select far more irrelevant variables than the PRESCIENCE approach, hence suffering from a larger extent of overfitting.

Table 2: Comparison of prediction methods under DGP(i) with $p = 10$

method	PRESCIENCE(q)			PRESCIENCE_CV	logit_lasso	probit_lasso
	$q = 1$	$q = 2$	$q = 3$			
<i>Corr_sel</i>	0.93	0.99	1	0.97	1	1
<i>Orac_sel</i>	0.93	0	0	0.51	0	0
<i>Num_irrel</i>	0.07	1.01	1.99	0.71	5.03	4.75
<i>in_Score</i>	0.948	0.964	0.974	0.960	0.947	0.944
<i>in_RS</i>	1.028	1.046	1.058	1.042	1.028	1.024
<i>out_Score</i>	0.904	0.901	0.898	0.903	0.904	0.904
<i>out_RS</i>	0.982	0.979	0.976	0.981	0.983	0.983

Table 3: Comparison of prediction methods under DGP(i) with $p = 200$

method	PRESCIENCE(q)			PRESCIENCE_CV	logit_lasso	probit_lasso
	$q = 1$	$q = 2$	$q = 3$			
<i>Corr_sel</i>	0.78	0.88	0.89	0.86	0.99	0.99
<i>Orac_sel</i>	0.78	0	0	0.51	0	0
<i>Num_irrel</i>	0.22	1.12	2.11	0.73	21.18	20.38
<i>in_Score</i>	0.943	0.965	0.972	0.957	0.981	0.977
<i>in_RS</i>	1.032	1.056	1.063	1.047	1.073	1.070
<i>out_Score</i>	0.893	0.891	0.883	0.895	0.876	0.876
<i>out_RS</i>	0.971	0.969	0.960	0.973	0.953	0.953

To save space, we present details of the simulation results under the setup of DGP(ii) in Online Appendix C. The results are similar to those under DGP(i).

6 An Illustrative Application

We illustrate usefulness of PRESCIENCE in the empirical application of work-trip mode choice. We used the transportation mode dataset analyzed by Horowitz (1993). This dataset has also been well studied for illustration of econometric methods developed for binary response models (e.g., see Florios and Skouras (2008), Benoit and Van den Poel (2012, Section 4.3), and the references therein). The previous literature focused on estimating slope coefficients in the binary response model; however, in this section, we are mainly interested in the numerical performance of alternative MIO algorithms and the result of covariate selection.

The data consist of 842 observations sampled randomly from the Washington, D.C., area transportation study. Each record in the dataset contains the following information for a single work trip of the traveler: the chosen transportation mode, the number of cars owned by the traveler's household (*CARS*), the transit out-of-vehicle travel time minus automobile out-of-vehicle travel time in minutes (*DOVTT*), the transit in-vehicle travel time minus automobile in-vehicle travel time in minutes (*DIVTT*) and the transit fare minus automobile travel cost in dollars (*DCOST*).

The dependent variable Y is the traveler's chosen mode of transportation such that $Y = 1$ if the choice is automobile and 0 otherwise. Following Florios and Skouras (2008), we standardized each of explanatory variables to have mean zero and unit variance. Following Horowitz (1993) and Florios and Skouras (2008), we specified the coefficient of *DCOST* to be unity and did not estimate that parameter. We set the focus covariates $X = (X_0, \tilde{X})$ to be $(DCOST, 1)$, where the constant term was included to capture the regression intercept and the parameter α was set to be unity. The resulting PRESCIENCE problem hence reduced to the maximization problem (4.1). We implemented the two MIO formulations developed in Section 4 for solving this problem. To compare their computational performance, we report the CPU time (in seconds) and the number of branch-and-bound nodes that the MIO solver had explored to reach the optimal solution. The former depends on both the computing hardware and software configurations whereas the latter only depends on the solution algorithms employed by the MIO solver.

For the auxiliary covariates, we set

$$\begin{aligned} Z = & (CARS, DOVTT, DIVTT, CARS \times DOVTT, DOVTT \times DIVTT, \\ & CARS \times DIVTT, CARS \times CARS, DOVTT \times DOVTT, \\ & DIVTT \times DIVTT). \end{aligned} \tag{6.1}$$

The covariate specification (6.1) allows us to approximate a smooth function of the variables ($CARS, DOVTT, DIVTT$) by its quadratic expansion.¹¹ We are interested in the data driven selection of these expansion terms through the PRESCIENCE procedure. In this setup, we have that $k = 1$ and $p = 9$. We specified the parameter space Θ to be $[-10, 10]^{10}$. We computed the PRESCIENCE for each $q \in \{1, 3, 5\}$. Since there are 10 unknown parameters in this setup, we solved for the PRESCIENCE solutions with a non-zero tolerance level which was also specified according to the rule (5.3). For $n = 842$, this amounts to setting the MIO tolerance level to be about 4.4%. To further reduce the computational cost, we adopted the warm-start strategy in the resulting MIO formulations. We set $\tau = 1.5$ and constructed $(\hat{P}_i)_{i=1}^n$ using the fitted choice probabilities from the logit regression of Y on all the covariates to derive the refined space $\hat{\Theta}(\tau)$ from the initial parameter space Θ .

Table 4: Refined parameter bounds ($\tau = 1.5$)

Covariate specification: $k = 1, p = 9$

Variable	lower bound	upper bound
Intercept	-10	10
$CARS$	-10	10
$DOVTT$	-9.8299	9.8299
$DIVTT$	-8.0158	8.0158
$CARS \times DOVTT$	-6.0306	6.0306
$DOVTT \times DIVTT$	-7.5870	7.5870
$CARS \times DIVTT$	-5.7873	5.7873
$CARS \times CARS$	-4.2513	4.2513
$DOVTT \times DOVTT$	-1.9552	1.9552
$DIVTT \times DIVTT$	-5.7297	5.7297

Table 4 presents the refined parameter bounds derived from $\hat{\Theta}(\tau)$ in this setup. The results of Table 4 indicate that the size (measured in the volume of a $(k + p)$ dimensional cube) of $\hat{\Theta}(\tau)$ is only about 0.99% of that of Θ ; thus, there is a considerable reduction in the parameter search space by using the refined space $\hat{\Theta}(\tau)$ in place of the original space Θ .

We now present in Table 5 the estimation results for the setup with the covariate specification (6.1). From Table 5, we can see that the two MIO formulations (4.5) and (4.14) yield the same set of selected variables across all the three cases of q . The parameter estimates computed from both formulations are also qualita-

¹¹In Online Appendix E, we report empirical results using $Z = (CARS, DOVTT, DIVTT)$.

Table 5: Results Using Quadratic Expansion

Covariate specification: $k = 1, p = 9$

q	1		3		5	
MIO formulation	(4.5)	(4.14)	(4.5)	(4.14)	(4.5)	(4.14)
focus covariates						
$DCOST$	1	1	1	1	1	1
Intercept	3.6374	3.2803	4.5191	3.1493	5.4587	2.9101
auxiliary covariates						
$CARS$	3.0404	2.4667	3.5782	2.4416	6.3592	2.1480
$DOVTT$	0	0	0.8390	0.3416	0	0
$DIVTT$	0	0	0	0	0	0
$CARS \times DOVTT$	0	0	0	0	-1.1798	-0.5332
$DOVTT \times DIVTT$	0	0	0	0	-3.7056	-0.4177
$CARS \times DIVTT$	0	0	-0.2744	0.1644	0	0
$CARS \times CARS$	0	0	0	0	0	0
$DOVTT \times DOVTT$	0	0	0	0	1.3282	0.1835
$DIVTT \times DIVTT$	0	0	0	0	2.7936	0.0744
in-sample performance						
maximized average score	0.8979	0.8979	0.9086	0.9086	0.9145	0.9097
MIO solver output						
MIO_gap	0.0428	0.0428	0.0428	0.0428	0.0428	0.0428
CPU time (in seconds)	55	259	65985	3931	566	778
branch-and-bound nodes	24547	71943	1521685	439394	208537	425053

tively similar in general. The variable $CARS$ remains to be selected in all these cases and its parameter estimate is also of the largest magnitude among all parameter estimates of the quadratic expansion variables. Moreover, there is very little loss in the goodness of fit from adopting only $CARS$ as the auxiliary covariate.

We now remark on the computational performance. For the cases of $q \in \{1, 5\}$, formulation (4.5) clearly outperformed formulation (4.14) in both the CPU time and the number of branch-and-bounds used in the computation; however, both MIO formulations performed quite well for these two variable selection cases. By contrast, for the case of $q = 3$, both approaches incurred more computational cost. In particular, it took around 18.3 hours and noticeably more branch-and-bound nodes to solve the formulation (4.5) in the $q = 3$ scenario. On the whole, these results suggest that the two MIO formulations (4.5) and (4.14) are valuable complements for implementing the PRESCIENCE procedures.

By construction, the maximized average score of the best subset variable selection approach increases with the specified value of q , which restricts the maximal

number of selected covariates in the resulting prediction rule. In this empirical application, our results presented so far indicate that the in-sample predictive performance of the parsimonious predictive model using an intercept term and only the two variables ($DCOST, CARS$) seems comparable to that of a more complex model using a richer set of covariates. We now investigate this issue further via the method of cross validation (CV).

We conducted the 5-fold CV analysis to assess the out-of-sample predictive performance of the PRESCIENCE methods for the selection of the variables specified by (6.1). As in Table 5, we considered the cases of $q \in \{1, 3, 5\}$ for which we implemented the corresponding PRESCIENCE procedures with a non-zero tolerance level. Because the training sample in each CV fold contains around 80% of the original observations, we set the MIO tolerance level for early termination to be about 4.9% by the rule (5.3).

Table 6: Summary of the 5-fold Cross Validation Results

Covariate specification: $k = 1, p = 9$


q	1		3		5	
MIO formulation	(4.5)	(4.14)	(4.5)	(4.14)	(4.5)	(4.14)
average in-sample performance						
MIO_gap	0.0454	0.0475	0.0475	0.0475	0.0475	0.0475
maximized objective value	0.8993	0.8993	0.9100	0.9112	0.9127	0.9148
average out-of-sample performance						
proportion of correct predictions	0.9026	0.8979	0.8884	0.8812	0.8932	0.8884

Table 6 summarizes the 5-fold CV results, which are based on the averages over the performance results computed in each CV fold. From Table 6, we can see that, for both MIO formulations, the in-sample maximized objective values were already very similar across the three cases of q though they did strictly increase with q . Moreover, irrespective of the MIO formulations, the parsimonious case of $q = 1$ had the best out-of-sample performance.

We now inspect the variables selected in each of the 5 CV folds. Figure 1 summarizes the parameter estimates computed in each CV fold for the 9 auxiliary covariates specified by (6.1). From this figure, we also note that $CARS$ was selected across all q cases in all CV folds. Its parameter estimate was also of a relatively large magnitude when compared to those of other selected variables. These cross-validation results further strengthen the finding that $CARS$ may be the most important predictive variable for the work-trip mode choice.

Figure 1: Auxiliary Covariate Parameter Estimates in each CV-fold

Covariate specification: $k = 1, p = 9$



Notes: For each panel of Figure 1, the values in the horizontal axis correspond to the indices of the 9 components of the auxiliary covariate vector defined by (6.1), whereas the vertical axis displays the values of the parameter estimates.

7 Concluding Remarks

In this paper, we consider the variable selection problem for predicting binary outcomes. We study the best subset selection procedure by which the covariates are chosen by maximizing the maximum score objective function subject to a constraint on the maximal number of selected covariates. We establish non-

asymptotic upper and lower risk bounds for the resulting best subset maximum score binary prediction rule when the dimension of potential covariates is possibly much larger than the sample size. The derived upper and lower risk bounds are minimax rate-optimal when the maximal number of selected variables is fixed and does not increase with the sample size. For implementation, we show that the variable selection problem of this paper can be equivalently formulated as a mixed integer optimization problem, which enables computation of the exact or an approximate solution with a definite approximation error bound.

The present paper takes the maximum score approach for the binary prediction problem. There is a large body of the literature that studies maximum score estimation in various other aspects since the seminal work by Manski (1975, 1985). In the context of semiparametric binary response models, advances of the maximum score approach have been made in terms of point identification (Manski, 1988), partial identification (Manski and Tamer, 2002; Komarova, 2013; Blevins, 2015; Chen and Lee, 2015), asymptotic distribution (Kim and Pollard, 1990; Seo and Otsu, 2017), panel data (Manski, 1987; Charlier, Melenberg, and van Soest, 1995; Abrevaya, 2000), time series (Moon, 2004; Guerre and Moon, 2006; de Jong and Woutersen, 2011), dynamic network formation (Graham, 2016), nonparametrically generated regressors (Chen, Lee, and Sung, 2014), and so on. The numerical approach employed in this paper can be adapted to these contexts.

Efficient computation is particularly demanding when it is necessary to obtain maximum score estimates repeatedly many times. This difficulty naturally arises, for example, in the context of resampling (Delgado, Rodriguez-Poo, and Wolf, 2001; Abrevaya and Huang, 2005; Lee and Pun, 2006; Patra, Seijo, and Sen, 2015) and change-point problems (Lee and Seo, 2008; Lee, Seo, and Shin, 2011). It would be an interesting future research topic to investigate numerical performance of our method (e.g. the warm-start procedure in Section 4.3) in these computation-intensive problems.

The maximum score approach has produced many offsprings: smoothed maximum score estimation (Horowitz, 1992, 2002), multinomial choice estimation (Matzkin, 1993; Fox, 2007), integrated maximum score estimation (Chen, 2010), Bayesian method (Benoit and Van den Poel, 2012), alternative estimation based on local nonlinear least squares (Blevins and Khan, 2013a,b; Khan, 2013), estimation using local polynomial smoothing (Chen and Zhang, 2015), and non-Bayesian Laplace type estimation (Jun, Pinkse, and Wan, 2015, 2017) among many others. Some of these alternative estimation methods are equipped with algorithms that are easier to compute than the maximum score estimation. It is an interesting open question how to accommodate a variable selection problem in these

alternative methods. One starting point can be the work by Jiang and Tanner (2010) who considered empirical risk minimization under the ℓ_0 constraint with Gibbs posterior as an alternative to maximum score prediction. It might be also interesting to consider penalized estimation with the ℓ_1 and/or ℓ_2 penalty. The maximum score approach is also closely related to the maximum utility estimation framework (Lieli and Nieto-Barthaburu, 2010; Lieli and White, 2010; Elliott and Lieli, 2013; Lieli and Springborn, 2013) for binary decision making under model uncertainty. It would be interesting to generalize our results to this framework. These are also topics for future research.

A Proofs of theoretical results

A.1 Proofs of Propositions 1, 2, and 3

Proof of Proposition 1. By (3.22) and (3.23), we have that

$$w'\theta^* \geq 0 \iff \eta(w) \geq 1/2 \iff b^*(w) = 1. \quad (\text{A.1})$$

Since the functional $\tilde{S}(b^*)$ is maximized at $b = b^*$ and, by Condition 1, this maximum is unique, we have that

$$b^* \in \mathcal{B}_q \iff b^* = b_{\theta_0}. \quad (\text{A.2})$$

By (1.1), we have that

$$b_{\theta_0}(w) = 1 \iff w'\theta_0 \geq 0. \quad (\text{A.3})$$

Thus it follows from (A.1), (A.2) and (A.3) that $b^* \in \mathcal{B}_q$ if and only if $W'\theta_0$ and $W'\theta^*$ have the same sign with probability 1. ■

Proof of Proposition 2. Let $G_w(t) \equiv P(\xi \leq t | W = w)$. By (3.22) and (3.23), we have that $\eta(w) = G_w(w'\theta^*)$ and $G_w(0) = 1/2$. By condition (ii) of Proposition 2, it hence follows that, if $w'\theta^* \geq \kappa_1$, then

$$\eta(w) \geq G_w(\kappa_1) \geq 1/2 + \kappa_1\kappa_2.$$

On the other hand, by similar arguments, we have that, if $w'\theta^* \leq -\kappa_1$, then

$$\eta(w) \leq G_w(-\kappa_1) \leq 1/2 - \kappa_1\kappa_2.$$

Combing these results and using condition (i) of Proposition 2, we thus deduce that condition (3.24) holds with $h = 2\kappa_1\kappa_2$. ■

Proof of Proposition 3. Consider any two vectors $\theta = (\alpha, \beta, \gamma)$, $\tilde{\theta} = (\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma}) \in \{-1, 1\} \times \Theta_q$ such that $\alpha = \tilde{\alpha}$ and $\beta = \tilde{\beta}$. We now prove this proposition for the case $\alpha = \tilde{\alpha} = 1$. The other case $\alpha = \tilde{\alpha} = -1$ can be proved using identical arguments and hence is omitted.

Assume $\alpha = \tilde{\alpha} = 1$. Note that

$$\begin{aligned} & \|b_\theta - b_{\tilde{\theta}}\|_1 \\ &= P\left(-\tilde{X}\beta - Z'\gamma \leq X_0 < -\tilde{X}\tilde{\beta} - Z'\tilde{\gamma}\right) + P\left(-\tilde{X}\beta - Z'\gamma > X_0 \geq -\tilde{X}\tilde{\beta} - Z'\tilde{\gamma}\right) \\ &= E\left[P\left(-\tilde{X}\beta - Z'\gamma \leq X_0 < -\tilde{X}\tilde{\beta} - Z'\tilde{\gamma} \mid \tilde{X}, Z\right) 1\{Z'\gamma \geq Z'\tilde{\gamma}\}\right] \\ & \quad + E\left[P\left(-\tilde{X}\beta - Z'\gamma > X_0 \geq -\tilde{X}\tilde{\beta} - Z'\tilde{\gamma} \mid \tilde{X}, Z\right) 1\{Z'\gamma \leq Z'\tilde{\gamma}\}\right]. \end{aligned}$$

By Condition (a) of Proposition 3, we hence have that

$$E\left[\|Z'(\gamma - \tilde{\gamma})\| L_1^{-1}\right] \leq \|b_{\theta_a} - b_{\theta_b}\|_1 \leq E\left[\|Z'(\gamma - \tilde{\gamma})\| L_1\right]. \quad (\text{A.4})$$

Let $J \equiv \{j \in \{1, \dots, p\} : \gamma_j \neq \tilde{\gamma}_j\}$. Since $\|\gamma\|_0 \leq q$ and $\|\tilde{\gamma}\|_0 \leq q$, we have that $|J| \leq 2q$. Therefore, $Z'(\gamma - \tilde{\gamma}) = Z'_J \delta_J$ where Z_J denotes the subvector of $Z \equiv (Z^{(1)}, \dots, Z^{(p)})'$ formed by keeping only those elements $Z^{(j)}$ with $j \in J$ and δ_J denotes the subvector of $\gamma - \tilde{\gamma}$ formed by keeping only those elements $(\gamma_j - \tilde{\gamma}_j)$ with $j \in J$.

By Condition (b) and Cauchy-Schwarz inequality, we have that with probability 1,

$$|Z'_J \delta_J| \leq L_2 \|\delta_J\|_E. \quad (\text{A.5})$$

and hence

$$\frac{\delta'_J Z_J Z'_J \delta_J}{L_2^2 \|\delta_J\|_E^2} \leq \frac{|Z'_J \delta_J|}{L_2 \|\delta_J\|_E} \leq 1. \quad (\text{A.6})$$

Using (A.6) and the assumption that the smallest eigenvalue of $E(Z_J Z'_J)$ is bounded below by L_3 , we thus have that

$$E\left[\|Z'_J \delta_J\|\right] \geq L_2^{-1} L_3 \|\delta_J\|_E. \quad (\text{A.7})$$

Noting that $\|\delta_J\|_E = \left\|\theta - \tilde{\theta}\right\|_E$ and combining (A.4), (A.5) and (A.7), we conclude that condition (3.25) holds with $c_u = L_1 L_2$ and $c_l = (L_1 L_2)^{-1} L_3$. ■

A.2 Proofs of Theorem 1

Recall the notation $\theta = (\alpha, \beta, \gamma)$. Let

$$G_n(\theta) \equiv \sqrt{n} [S_n(\theta) - S(\theta)].$$

We first present the following lemma, which will be used to prove Theorem 1.

Lemma 1. *For $t > 0$, there is a universal constant D such that*

$$P \left(\sup_{\theta \in \{-1, 1\} \times \Theta_q} |G_n(\theta)| > t \right) \leq 2 \binom{p}{q} \left(\frac{Dt}{\sqrt{4s+4}} \right)^{4s+4} e^{-2t^2}.$$

Proof. Let m be a subset of the index set $\{1, \dots, p\}$ such that m contains only q elements. Let \mathcal{M} be the collection of all such subsets. Note that $|\mathcal{M}| = \binom{p}{q}$. For $m \in \mathcal{M}$, let

$$\Gamma_m \equiv \{(\beta, \gamma) \in \Theta : \gamma_j = 0 \text{ for } j \notin m\}.$$

For any $t > 0$, we have that

$$\begin{aligned} P \left(\sup_{\theta \in \{-1, 1\} \times \Theta_q} |G_n(\theta)| > t \right) &\leq \sum_{m \in \mathcal{M}} P \left(\sup_{(\beta, \gamma) \in \Gamma_m} |G_n(1, \beta, \gamma)| > t \right) \\ &\quad + \sum_{m \in \mathcal{M}} P \left(\sup_{(\beta, \gamma) \in \Gamma_m} |G_n(-1, \beta, \gamma)| > t \right). \end{aligned} \quad (\text{A.8})$$

To complete the proof, it remains to derive the bounds on the tail probability terms on the right hand side of the inequality above.

Consider the function $f_\theta : \{0, 1\} \times \mathcal{W} \mapsto \{0, 1\}$ defined by

$$\begin{aligned} f_\theta(y, w) &\equiv 1 \{y = b_\theta(w)\} \\ &= 1 - y + (2y - 1) b_\theta. \end{aligned}$$

For $m \in \mathcal{M}$, let

$$\begin{aligned} F_m^+ &\equiv \{f_\theta : \theta \in \{1\} \times \Gamma_m\}, \\ F_m^- &\equiv \{f_\theta : \theta \in \{-1\} \times \Gamma_m\}. \end{aligned}$$

For each $m \in \mathcal{M}$, by Lemmas 9.6 and 9.9 of Kosorok (2008), the class of functions F_m^+ and F_m^- are both VC classes of functions with VC indices $V(F_m^+)$ and $V(F_m^-)$ satisfying that

$$V(F_m^+) = V(F_m^-) \leq 2s + 3.$$

For two measurable functions f and g and a given probability measure Q , define the (semi-) metric

$$d_Q(f, g) \equiv \sqrt{\int (f - g)^2 dQ}.$$

By Theorem 9.3 of Kosorok (2008), we have that, for some universal constant $K > 1$ and $0 < \epsilon < 1$,

$$\begin{aligned} & \sup_Q N(\epsilon, F_m^+, d_Q) \vee \sup_Q N(\epsilon, F_m^-, d_Q) \\ & \leq K(2s+3)(16e)^{2s+3}(\epsilon)^{-4(s+1)} \\ & \leq \left(\frac{\Lambda(s)}{\epsilon}\right)^{4s+4}, \end{aligned}$$

where

$$\Lambda(s) \equiv [K(2s+3)(16e)^{2s+3}]^{\frac{1}{4s+4}}$$

and, for a given class of functions F , $N(\epsilon, F, d_Q)$ denotes the minimal number of open balls (defined under the metric d_Q) of radius ϵ required to cover F .

Observe that

$$\Lambda(s) \leq 16Ke(2s+3)^{\frac{1}{4s+4}} \leq 32Ke, \quad (\text{A.9})$$

where the second inequality follows from the fact that $i \leq 2^i$ for all integer i .

Using (A.9), we can apply Theorem 1.3 of Talagrand (1994) to deduce that, for $t > 0$, there is a universal constant D such that

$$\begin{aligned} & P\left(\sup_{(\beta, \gamma) \in \Gamma_m} |G_n(1, \beta, \gamma)| > t\right) \vee P\left(\sup_{(\beta, \gamma) \in \Gamma_m} |G_n(-1, \beta, \gamma)| > t\right) \\ & \leq \left(\frac{Dt}{\sqrt{4s+4}}\right)^{4s+4} e^{-2t^2}. \end{aligned} \quad (\text{A.10})$$

Lemma 1 hence follows by combining inequalities (A.8) and (A.10). ■

We now prove Theorem 1.

Proof of Theorem 1. By Lemma 1 and using the fact that $\binom{p}{q} \leq p^q$, we have that, for some universal constant D and for $t > 0$,

$$P\left(\sup_{\theta \in \{-1, 1\} \times \Theta_q} |G_n(\theta)| > t\right) \leq 2p^q \left(\frac{Dt}{\sqrt{4s+4}}\right)^{4s+4} e^{-2t^2}. \quad (\text{A.11})$$

For $t \geq D$ and $s \geq 1$, the right hand side term of (A.11) can be further bounded above by

$$p^s t^{8s+8} 2^{-6s-5} e^{-2t^2} = e^{\lambda(s, p, t)}, \quad (\text{A.12})$$

where

$$\lambda(s, p, t) \equiv -2t^2 + (8s + 8) \ln t + s \ln p - (6s + 5) \ln 2.$$

For $\sigma > 0$, let

$$t = \sqrt{M_\sigma r_n} \quad (\text{A.13})$$

where

$$M_\sigma \equiv (1 + \sigma/2) \vee D^2. \quad (\text{A.14})$$

Note that $t \geq D$ by (A.13) and (A.14). Thus with the t value specified by (A.13), we have that

$$\lambda(s, p, t) \leq (-2M_\sigma + 1) r_n + (4s + 4) \ln(M_\sigma r_n) - (6s + 5) \ln 2 \leq -\sigma r_n \quad (\text{A.15})$$

where the second inequality follows from (A.14) and (3.6).

It is straightforward to see that

$$\begin{aligned} U_n &= \sup_{\theta \in \{-1, 1\} \times \Theta_q} S(\theta) - S(\widehat{\theta}) \\ &\leq \sup_{\theta \in \{-1, 1\} \times \Theta_q} |S_n(\theta) - S(\theta)| + \sup_{\theta \in \{-1, 1\} \times \Theta_q} S_n(\theta) - S(\widehat{\theta}) \\ &\leq \sup_{\theta \in \{-1, 1\} \times \Theta_q} |S_n(\theta) - S(\theta)| + S_n(\widehat{\theta}) - S(\widehat{\theta}) \\ &\leq 2 \sup_{(\alpha, \beta, \gamma) \in \{-1, 1\} \times \Theta_q} |S_n(\theta) - S(\theta)|. \end{aligned}$$

Hence, we have that

$$P \left(U_n > 2\sqrt{\frac{M_\eta r_n}{n}} \right) \leq P \left(\sup_{\theta \in \{-1, 1\} \times \Theta_q} |G_n(\theta)| > \sqrt{M_\eta r_n} \right). \quad (\text{A.16})$$

Therefore, Theorem 1 follows by putting together the results of A.11), (A.12) and (A.15) and then concluding that the right hand side term of (A.16) is bounded above by $e^{-\sigma r_n}$ under condition (3.6). ■

A.3 Proof of Theorem 2

We first introduce some notation which will be used in the proof Theorem 2. Let \mathcal{A} be a collection of subsets of \mathcal{W} . For any subset $S \subset \mathcal{W}$, let $T_{\mathcal{A}}(S)$ denote the trace of \mathcal{A} on S defined by

$$T_{\mathcal{A}}(S) = \{A \cap S : A \in \mathcal{A}\}.$$

If $\mathcal{A} = \cup_{j \in J} \mathcal{A}_j$, then we have

$$T_{\mathcal{A}}(S) \subset \cup_{j \in J} T_{\mathcal{A}_j}(S) \text{ for } S \subset \mathcal{W}. \quad (\text{A.17})$$

We now prove Theorem 2.

Proof of Theorem 2. By (1.1), (3.1), (3.2) and (3.15), we have that

$$S(\hat{\theta}) = \tilde{S}(b_{\hat{\theta}}) \text{ and } S_q^* = \sup_{b \in \mathcal{B}_q} \tilde{S}(b). \quad (\text{A.18})$$

Theorem 2 is an application of Theorem 2 of Massart and Nédélec (2006) to the binary prediction problem. Massart and Nédélec (2006, Section 2.4) showed how to apply their Theorem 2 to derive the risk upper bound for the empirical risk minimizer. Using their derived results (Massart and Nédélec (2006, p. 2340)) in their Theorem 2 for our setup, we conclude that, under Condition 1, there are universal constants K and K' such that

$$\begin{aligned} & E \left[\tilde{S}(b^*) - \tilde{S}(b_{\hat{\theta}}) \right] \\ & \leq 2 \left[\tilde{S}(b^*) - \sup_{b \in \mathcal{B}_q} \tilde{S}(b) \right] \\ & + K' \left[\left(\frac{K^2 (1 \vee E(H_{\mathcal{A}}))}{nh} \right)^{\theta/(2\theta-1)} \wedge \sqrt{\frac{K^2 (1 \vee E(H_{\mathcal{A}}))}{n}} \right], \end{aligned} \quad (\text{A.19})$$

where

$$\begin{aligned} \mathcal{A} & \equiv \{A_{\theta} : \theta \in \{-1, 1\} \times \Theta_q\}, \\ A_{\theta} & \equiv \{w \in \mathcal{W} : w'\theta \geq 0\}, \end{aligned}$$

and $H_{\mathcal{A}}$ is the random combinatorial entropy of \mathcal{A} defined by

$$H_{\mathcal{A}} = \ln (|T_{\mathcal{A}}(\{W_1, W_2, \dots, W_n\})|).$$

Using (3.3) and (A.18), it follows that

$$\begin{aligned} U_n & = \sup_{b \in \mathcal{B}_q} \tilde{S}(b) - \tilde{S}(b_{\hat{\theta}}) \\ & = \left[\sup_{b \in \mathcal{B}_q} \tilde{S}(b) - \tilde{S}(b^*) \right] + \left[\tilde{S}(b^*) - \tilde{S}(b_{\hat{\theta}}) \right]. \end{aligned} \quad (\text{A.20})$$

To complete the proof, it thus remains to derive an upper bound on the term $E(H_{\mathcal{A}})$.

Let m be a subset of the index set $\{1, \dots, p\}$ such that m contains only q elements. Let \mathcal{M} be the collection of all such subsets. Note that $|\mathcal{M}| = \binom{p}{q}$. For

$m \in \mathcal{M}$, let

$$\begin{aligned}\mathcal{A}_m^+ &\equiv \{A_\theta : \theta \in \{1\} \times \Gamma_m\}, \\ \mathcal{A}_m^- &\equiv \{A_\theta : \theta \in \{-1\} \times \Gamma_m\}, \\ \Gamma_m &\equiv \{(\beta, \gamma) \in \Theta : \gamma_j = 0 \text{ for } j \notin m\}.\end{aligned}$$

It is immediate to see that

$$\mathcal{A} = \cup_{m \in \mathcal{M}} (\mathcal{A}_m^+ \cup \mathcal{A}_m^-). \quad (\text{A.21})$$

For each $m \in \mathcal{M}$, by Lemmas 9.6 and 9.9 of Kosorok (2008), the family of sets \mathcal{A}_m^+ and \mathcal{A}_m^- are both VC classes of sets with VC indices $V(\mathcal{A}_m^+)$ and $V(\mathcal{A}_m^-)$ satisfying that

$$V(\mathcal{A}_m^+) = V(\mathcal{A}_m^-) \leq s + 2.$$

Hence by Corollary 1.3 of Lugosi (2002), we have that

$$|T_{\mathcal{A}_m^+}(\{W_1, W_2, \dots, W_n\})| \vee |T_{\mathcal{A}_m^-}(\{W_1, W_2, \dots, W_n\})| \leq (n + 1)^{s+1}. \quad (\text{A.22})$$

By (A.17), (A.21) and (A.22), we thus have that

$$\begin{aligned}H_{\mathcal{A}} &\leq \ln 2 + \ln \binom{p}{q} + (s + 1) \ln(n + 1) \\ &\leq \ln 2 + q \ln p + (s + 1) \ln(n + 1).\end{aligned} \quad (\text{A.23})$$

Theorem 2 therefore follows by combining the results (A.19), (A.20) and (A.23). ■

A.4 Proof of Theorem 3

Proof of Theorem 3. Define

$$R_n(h, \mathcal{B}_q) = \inf_{\widehat{b} \in \mathcal{B}_q} \sup_{F \in \mathcal{P}(h, \mathcal{B}_q)} E_F \left[\widetilde{S}(b^*) - \widetilde{S}(\widehat{b}) \right]$$

where the infimum is taken over the set of all binary predictors in \mathcal{B}_q that are constructed based on the data $(Y_i, W_i)_{i=1}^n$. By (3.2), it follows that $S_q^* = \widetilde{S}(b^*)$ under $F \in \mathcal{P}(h, \mathcal{B}_q)$. Thus Theorem 3 is proved once we show that

$$R_n(h, \mathcal{B}_q) \geq \frac{\phi q c_l (1 - \phi) (1 - h)}{32 n h c_u} \ln \left(\frac{p - q}{q/2} \right). \quad (\text{A.24})$$

For any indicator function $b \in \mathcal{B}_q$, let

$$\eta_b(w) \equiv [1 + (2b(w) - 1)h] / 2 \text{ for } w \in \mathcal{W}.$$

Let F_b denote the joint distribution of (Y, W) such that under F_b , the distribution of W satisfies condition (3.25), and Y conditional on $W = w$ follows a Bernoulli distribution with parameter $\eta_b(w)$ for every $w \in \mathcal{W}$. By the same arguments as those in the proof of Theorem 6 of Massart and Nédélec (2006, p. 2355), we can deduce that, for any finite subset \mathcal{C} of \mathcal{B}_q ,

$$\{F_b : b \in \mathcal{C}\} \subset \mathcal{P}(h, \mathcal{B}_q) \tag{A.25}$$

and

$$R_n(h, \mathcal{B}_q) \geq \frac{h}{2} \inf_{\hat{b} \in \mathcal{C}} \sup_{b \in \mathcal{C}} E_{F_b} \left[\left\| b - \hat{b} \right\|_1 \right]. \tag{A.26}$$

Consider the set

$$\mathcal{H} \equiv \{\gamma \in \{-1, 0, 1\}^p : \|\gamma\|_0 = q\}.$$

By Lemma 4 of Raskutti, Wainwright, and Yu (2011), we have that, for p, q even and $q < 2p/3$, there is a subset $\mathcal{A} \subset \mathcal{H}$ with cardinality

$$|\mathcal{A}| \geq \left(\frac{p-q}{q/2} \right)^{q/2} \tag{A.27}$$

such that

$$\|\gamma - \gamma'\|_0 \geq q/2 \text{ for all } \gamma, \gamma' \in \mathcal{A}. \tag{A.28}$$

Let

$$\mathcal{D} \equiv \{1\} \times \{\underline{0}\} \times \epsilon_n \mathcal{A} \tag{A.29}$$

where $\underline{0} = (0, \dots, 0)$ denotes the k -dimensional vector of which all elements take value 0, and $\epsilon_n > 0$ is a given sequence that will be chosen later. Note that, for any $\theta, \tilde{\theta} \in \mathcal{D}$,

$$\epsilon_n^2 \left\| \theta - \tilde{\theta} \right\|_0 \leq \left\| \theta - \tilde{\theta} \right\|_E^2 \leq 4\epsilon_n^2 \left\| \theta - \tilde{\theta} \right\|_0. \tag{A.30}$$

Now take

$$\mathcal{C} = \{b_\theta \in \mathcal{B}_q : \theta \in \mathcal{D}\}. \tag{A.31}$$

We then have that

$$R_n(h, \mathcal{B}_q) \geq \frac{h}{2} \inf_{\hat{\theta} \in \mathcal{D}} \sup_{\theta \in \mathcal{D}} E_{F_{b_\theta}} [\|b_\theta - b_{\hat{\theta}}\|_1] \quad (\text{A.32})$$

$$\geq \frac{hc_l}{2} \inf_{\hat{\theta} \in \mathcal{D}} \sup_{\theta \in \mathcal{D}} E_{F_{b_\theta}} [\|\theta - \hat{\theta}\|_E] \quad (\text{A.33})$$

$$\geq \frac{hc_l \epsilon_n}{2} \inf_{\hat{\theta} \in \mathcal{D}} \sup_{\theta \in \mathcal{D}} E_{F_{b_\theta}} \left[\sqrt{\|\theta - \hat{\theta}\|_0} \right] \quad (\text{A.34})$$

$$\geq \frac{hc_l \epsilon_n \sqrt{q}}{2\sqrt{2}} \inf_{\hat{\theta} \in \mathcal{D}} \sup_{\theta \in \mathcal{D}} P_{F_{b_\theta}} (\hat{\theta} \neq \theta) \quad (\text{A.35})$$

$$\geq \frac{hc_l \epsilon_n \sqrt{q}}{2\sqrt{2}} \inf_{\hat{\theta} \in \mathcal{D}} \left[1 - \inf_{\theta \in \mathcal{D}} P_{F_{b_\theta}} (\hat{\theta} = \theta) \right], \quad (\text{A.36})$$

where the infimum in (A.32) is taken over the set of all estimators $\hat{\theta}$ taking values in \mathcal{D} ; (A.33) follows from (A.25), (A.31) and (3.25); (A.34) follows from (A.30); (A.35) follows from (A.28).

By Lemma 8 of Massart and Nédélec (2006), we have that, for a given point $\tilde{\theta} \in \mathcal{D}$,

$$\inf_{\hat{\theta} \in \mathcal{D}} P_{F_{b_\theta}} (\hat{\theta} = \theta) \leq \phi \vee \frac{\bar{\mathcal{K}}}{\ln(|\mathcal{D}|)}, \quad (\text{A.37})$$

where

$$\bar{\mathcal{K}} = \frac{n}{|\mathcal{D}| - 1} \sum_{\theta \in \mathcal{D}, \theta \neq \tilde{\theta}} \mathcal{K}(F_{b_\theta}, F_{b_{\tilde{\theta}}}),$$

and $\mathcal{K}(F_{b_\theta}, F_{b_{\tilde{\theta}}})$ is the Kullback-Leibler information between F_{b_θ} and $F_{b_{\tilde{\theta}}}$. For $\tilde{\theta}, \theta \in \mathcal{D}$ and $\theta \neq \tilde{\theta}$, using Lemma 7 of Massart and Nédélec (2006), we have that, for $h < 1$,

$$\begin{aligned} \mathcal{K}(F_{b_\theta}, F_{b_{\tilde{\theta}}}) &= h \ln \left(\frac{1+h}{1-h} \right) \|b_\theta - b_{\tilde{\theta}}\|_1 \\ &\leq \frac{2c_u h^2}{1-h} \|\theta - \tilde{\theta}\|_E \\ &\leq \frac{4c_u h^2 \epsilon_n}{1-h} \sqrt{\|\theta - \tilde{\theta}\|_0} \\ &\leq \frac{4c_u h^2 \epsilon_n \sqrt{2q}}{1-h}, \end{aligned}$$

where the last inequality follows since $\|\theta - \tilde{\theta}\|_0 \leq 2q$ for all $\tilde{\theta}, \theta \in \mathcal{D}$. Hence, we have that

$$\bar{\mathcal{K}} \leq \frac{4nc_u h^2 \epsilon_n \sqrt{2q}}{1-h}. \quad (\text{A.38})$$

Putting together (A.36), (A.37) and (A.38), we have that

$$R_n(h, \mathcal{B}_q) \geq \frac{h\epsilon_n c_l \sqrt{q}}{2\sqrt{2}} (1 - \phi) \quad (\text{A.39})$$

provided that

$$\frac{4nc_u h^2 \epsilon_n \sqrt{2q}}{(1-h) \ln(|\mathcal{D}|)} \leq \phi. \quad (\text{A.40})$$

By (A.27) and (A.29), condition (A.40) holds whenever

$$\frac{8nc_u h^2 \epsilon_n \sqrt{2q}}{(1-h) q \ln\left(\frac{p-q}{q/2}\right)} \leq \phi. \quad (\text{A.41})$$

By (3.26), (A.29) and (A.31), we have that $\epsilon_n \leq \kappa$. Therefore, we can get the result (A.24) by setting

$$\epsilon_n = \frac{\phi \sqrt{q} (1-h) \ln\left(\frac{p-q}{q/2}\right)}{8nc_u h^2 \sqrt{2}}$$

provided that this choice of ϵ_n also satisfies that $\epsilon_n \leq \kappa$, which can be easily seen to hold under the condition (3.28) for the lower bound on the value of h . ■

References

- ABREVAYA, J. (2000): “Rank estimation of a generalized fixed-effects regression model,” *Journal of Econometrics*, 95(1), 1–23.
- ABREVAYA, J., AND J. HUANG (2005): “On the bootstrap of the maximum score estimator,” *Econometrica*, 73(4), 1175–1204.
- BENOIT, D. F., AND D. VAN DEN POEL (2012): “Binary quantile regression: a Bayesian approach based on the asymmetric Laplace distribution,” *Journal of Applied Econometrics*, 27(7), 1174–1188.
- BERTSIMAS, D., A. KING, AND R. MAZUMDER (2016): “Best subset selection via a modern optimization lens,” *Annals of Statistics*, 44(2), 813–852.
- BLEVINS, J. R. (2015): “Non-Standard Rates of Convergence of Criterion-Function-Based Set Estimators,” *Econometrics Journal*, 18, 172–199.
- BLEVINS, J. R., AND S. KHAN (2013a): “Distribution-Free Estimation of Heteroskedastic Binary Response Models in Stata,” *Stata Journal*, 13, 588–602.
- (2013b): “Local NLLS Estimation of Semiparametric Binary Choice Models,” *Econometrics Journal*, 16, 135–160.

- CHARLIER, E., B. MELENBERG, AND A. H. O. VAN SOEST (1995): “A smoothed maximum score estimator for the binary choice panel data model with an application to labour force participation,” *Statistica Neerlandica*, 49(3), 324–342.
- CHEN, L.-Y., AND S. LEE (2015): “Breaking the curse of dimensionality in conditional moment inequalities for discrete choice models,” Cemmap Working Paper CWP26/15.
- CHEN, L.-Y., S. LEE, AND M. J. SUNG (2014): “Maximum score estimation with nonparametrically generated regressors,” *Econometrics Journal*, 17(3), 271–300.
- CHEN, S. (2010): “An integrated maximum score estimator for a generalized censored quantile regression model,” *Journal of Econometrics*, 155(1), 90–98.
- CHEN, S., AND H. ZHANG (2015): “Binary quantile regression with local polynomial smoothing,” *Journal of Econometrics*, 189(1), 24–40.
- CHEN, X. (2007): “Large Sample Sieve Estimation of Semi-Nonparametric Models (Chapter 76),” in *Handbook of Econometrics, Volume 6*, pp. 5549–5632. Elsevier.
- CONFORTI, M., G. CORNUÉJOLS, AND G. ZAMBELLI (2014): *Integer Programming*. Springer.
- DANILOV, D., AND J. R. MAGNUS (2004): “On the harm that ignoring pretesting can cause,” *Journal of Econometrics*, 122(1), 27–46.
- DE JONG, R., AND T. WOUTERSEN (2011): “Dynamic Time Series Binary Choice,” *Econometric Theory*, pp. 1–30.
- DELGADO, M. A., J. M. RODRIGUEZ-POO, AND M. WOLF (2001): “Subsampling inference in cube root asymptotics with an application to Manski’s maximum score estimator,” *Economics Letters*, 73(2), 241–250.
- DEVROYE, L., L. GYÖRFI, AND G. LUGOSI (1996): *Probabilistic Theory of Pattern Recognition*. Springer.
- ELLIOTT, G., AND R. LIELI (2013): “Predicting binary outcomes,” *Journal of Econometrics*, 174(1), 15–26.
- FLORIOS, K., AND S. SKOURAS (2008): “Exact computation of max weighted score estimators,” *Journal of Econometrics*, 146(1), 86–91.
- FOX, J. T. (2007): “Semiparametric estimation of multinomial discrete-choice models using a subset of choices,” *RAND Journal of Economics*, pp. 1002–1019.
- FRIEDMAN, J., T. HASTIE, AND R. TIBSHIRANI (2010): “Regularization paths for generalized linear models via coordinate descent,” *Journal of statistical software*, 33(1), 1–22.
- GRAHAM, B. S. (2016): “Homophily and transitivity in dynamic network formation,” Cemmap Working Paper CWP16/16.

- GREENSHTEIN, E. (2006): “Best subset selection, persistence in high-dimensional statistical learning and optimization under L_1 constraint,” *Annals of Statistics*, 34(5), 2367–2386.
- GREENSHTEIN, E., AND Y. RITOV (2004): “Persistence in High-Dimensional Linear Predictor Selection and the Virtue of Overparametrization,” *Bernoulli*, 10(6), 971–988.
- GUERRE, E., AND H. MOON (2006): “A study of a semiparametric binary choice model with integrated covariates,” *Econometric Theory*, 22(4), 721–742.
- HOROWITZ, J. L. (1992): “A Smoothed Maximum Score Estimator for the Binary Response Model,” *Econometrica*, 60(3), 505–531.
- (1993): “Semiparametric estimation of a work-trip mode choice model,” *Journal of Econometrics*, 58(1), 49–70.
- (1998): *Semiparametric Methods in Econometrics*. Springer.
- (2002): “Bootstrap critical values for tests based on the smoothed maximum score estimator,” *Journal of Econometrics*, 111(2), 141–167.
- JIANG, W., AND M. A. TANNER (2010): “Risk Minimization for Time Series Binary Choice with Variable Selection,” *Econometric Theory*, 26(5), 1437–1452.
- JOHNSON, D., AND F. PREPARATA (1978): “The densest hemisphere problem,” *Theoretical Computer Science*, 6(1), 93–107.
- JUN, S. J., J. PINKSE, AND Y. WAN (2015): “Classical Laplace estimation for $\sqrt[3]{n}$ -consistent estimators: improved convergence rates and rate-adaptive inference,” *Journal of Econometrics*, 187(1), 201–216.
- (2017): “Integrated Score Estimation,” *Econometric Theory*, forthcoming.
- KHAN, S. (2013): “Distribution free estimation of heteroskedastic binary response models using Probit/Logit criterion functions,” *Journal of Econometrics*, 172(1), 168–182.
- KIM, J., AND D. POLLARD (1990): “Cube Root Asymptotics,” *Annals of Statistics*, 18(1), 191–219.
- KITAGAWA, T., AND A. TETENOV (2015): “Who should be Treated? Empirical Welfare Maximization Methods for Treatment Choice,” Cemmap Working Paper CWP10/15.
- KOMAROVA, T. (2013): “Binary choice models with discrete regressors: Identification and misspecification,” *Journal of Econometrics*, 177(1), 14–33.
- KOSOROK, M. R. (2008): *Introduction to Empirical Processes and Semiparametric Inference*. Springer.

- LEE, S., AND M. H. SEO (2008): “Semiparametric estimation of a binary response model with a change-point due to a covariate threshold,” *Journal of Econometrics*, 144(2), 492–499.
- LEE, S., M. H. SEO, AND Y. SHIN (2011): “Testing for Threshold Effects in Regression Models,” *Journal of the American Statistical Association*, 106(493), 220–231.
- LEE, S. M. S., AND M. C. PUN (2006): “On m out of n Bootstrapping for Non-standard M-Estimation With Nuisance Parameters,” *Journal of the American Statistical Association*, 101(475), 1185–1197.
- LIELI, R. P., AND A. NIETO-BARTHABURU (2010): “Optimal binary prediction for group decision making,” *Journal of Business & Economic Statistics*, 28(2), 308–319.
- LIELI, R. P., AND M. SPRINGBORN (2013): “Closing the gap between risk estimation and decision making: efficient management of trade-related invasive species risk,” *Review of Economics and Statistics*, 95(2), 632–645.
- LIELI, R. P., AND H. WHITE (2010): “The construction of empirical credit scoring rules based on maximization principles,” *Journal of Econometrics*, 157(1), 110–119.
- LUGOSI, G. (2002): “Pattern Classification and Learning Theory,” in *Principles of Nonparametric Learning*, ed. by L. Györfi, pp. 1–56. Springer.
- MAGNAC, T., AND E. MAURIN (2008): “Partial identification in monotone binary models: discrete regressors and interval data,” *The Review of Economic Studies*, 75(3), 835–864.
- MAGNUS, J. R., AND J. DURBIN (1999): “Estimation of Regression Coefficients of Interest when Other Regression Coefficients are of no Interest,” *Econometrica*, 67(3), 639–643.
- MAMMEN, E., AND A. B. TSYBAKOV (1999): “Smooth discrimination analysis,” *Annals of Statistics*, 27(6), 1808–1829.
- MANSKI, C. F. (1975): “Maximum score estimation of the stochastic utility model of choice,” *Journal of Econometrics*, 3(3), 205–228.
- (1985): “Semiparametric analysis of discrete response. Asymptotic properties of the maximum score estimator,” *Journal of Econometrics*, 27(3), 313–333.
- (1987): “Semiparametric Analysis of Random Effects Linear Models from Binary Panel Data,” *Econometrica*, 55(2), 357–362.
- (1988): “Identification of binary response models,” *Journal of the American Statistical Association*, 83(403), 729–738.

- MANSKI, C. F., AND E. TAMER (2002): “Inference on regressions with interval data on a regressor or outcome,” *Econometrica*, 70(2), 519–546.
- MANSKI, C. F., AND T. S. THOMPSON (1986): “Operational characteristics of maximum score estimation,” *Journal of Econometrics*, 32(1), 85–108.
- (1989): “Estimation of best predictors of binary response,” *Journal of Econometrics*, 40(1), 97–123.
- MASSART, P., AND E. NÉDÉLEC (2006): “Risk bounds for statistical learning,” *Annals of Statistics*, 34(5), 2326–2366.
- MATZKIN, R. L. (1993): “Nonparametric identification and estimation of polychotomous choice models,” *Journal of Econometrics*, 58(1-2), 137–168.
- MOON, H. (2004): “Maximum score estimation of a nonstationary binary choice model,” *Journal of Econometrics*, 122(2), 385–403.
- NATARAJAN, B. K. (1995): “Sparse Approximate Solutions to Linear Systems,” *SIAM Journal on Computing*, 24(2), 227–234.
- PATRA, R. K., E. SEIJO, AND B. SEN (2015): “A Consistent Bootstrap Procedure for the Maximum Score Estimator,” arXiv:1105.1976.
- PINKSE, C. (1993): “On the computation of semiparametric estimates in limited dependent variable models,” *Journal of Econometrics*, 58(1), 185–205.
- RASKUTTI, G., M. J. WAINWRIGHT, AND B. YU (2011): “Minimax rates of estimation for high-dimensional linear regression over l_q -balls,” *IEEE Transactions on Information Theory*, 57(10), 6976–6994.
- SEO, M. H., AND T. OTSU (2017): “Local M-estimation with Discontinuous Criterion for Dependent and Limited Observations,” *Annals of Statistics*, forthcoming.
- TALAGRAND, M. (1994): “Sharper Bounds for Gaussian and Empirical Processes,” *Annals of Probability*, 22(1), 28–76.
- TSYBAKOV, A. B. (2004): “Optimal aggregation of classifiers in statistical learning,” *Annals of Statistics*, 32(1), 135–166.

Online Appendices (NOT FOR PUBLICATION)

B The Branch-and-Bound Method for Solving MIO Problems

For completeness of the paper and for readers who are unfamiliar with MIO, we will present briefly the branch-and-bound method for solving the MIO problem. For further details, see, e.g., Conforti, Cornuéjols, and Zambelli (2014) for a recent and comprehensive study on the MIO theory and solution methods.

We take the formulation (4.5) as an expositional example and explain how the branch-and-bound method can be used to solve this MIO problem. The maximization problem (4.5) consists of $n + p$ binary control variables. Let $v = (d_1, \dots, d_n, e_1, \dots, e_p)$ denote the $(n + p)$ dimensional vector collecting all these binary controls. Let $\tilde{S}_n(\beta, \gamma, v)$ denote the objective function of (4.5). We may maximize \tilde{S}_n over v by enumerating all possible values of v , which amounts to exhaustively searching over a binary tree that has 2^{n+p} leaf nodes. This naive method is inefficient and becomes practically infeasible for large scale problems. The branch-and-bound method improves the search efficiency by avoid visiting those tree nodes which can be fathomed not to constitute the optimum.

Let Γ_0 denote the space of the controls (β, γ, v) defined by all the constraints stated in the MIO problem (4.5). Let $\bar{\Gamma}_0$ be an enlargement of Γ_0 , which is defined analogously to Γ_0 with the dichotomization constraints $v \in \{0, 1\}^{n+p}$ being replaced by the constraints $v \in [0, 1]^{n+p}$. Optimizing the objective function \tilde{S}_n over $(\beta, \gamma, v) \in \bar{\Gamma}_0$ reduces to a simple linear programming (LP) problem. Clearly, the maximized objective value of this LP relaxation problem forms an upper bound on the function $\tilde{S}_n(\beta, \gamma, v)$ defined on the original domain Γ_0 . Moreover, if the solution for v in the LP relaxation problem turns out to be a vector of binary values, we can deduce that the LP relaxation solution for (β, γ, v) is also the solution to the MIO problem (4.5).

When the LP relaxation solution for v contains fractional-valued elements, we choose a fractional-valued element v_j and then construct the two LP sub-problems, denoted as LP_1 and LP'_1 , which correspond to maximizing $\tilde{S}_n(\beta, \gamma, v)$ over the subspaces $\bar{\Gamma}_1 \equiv \bar{\Gamma}_0 \cap \{(\beta, \gamma, v) : v_j = 0\}$ and $\bar{\Gamma}'_1 \equiv \bar{\Gamma}_0 \cap \{(\beta, \gamma, v) : v_j = 1\}$, respectively. Consider the problem LP_1 and note that the treatment of LP'_1 is similar. There are four possible cases for LP_1 : (i) $\bar{\Gamma}_1$ is empty and hence LP_1 is infeasible. (ii) $\bar{\Gamma}_1$ is non-empty and the maximized objective value of LP_1 is not larger than the best known lower bound on the objective value of (4.5). (iii) $\bar{\Gamma}_1$ is

non-empty, the maximized objective value of LP_1 is larger than the best known lower bound on the objective value of (4.5), and the solution for v of LP_1 is in $\{0, 1\}^{n+p}$. (iv) $\bar{\Gamma}_1$ is non-empty, the maximized objective value of LP_1 is larger than the best known lower bound on the objective value of (4.5), and the solution for v of LP_1 contains fractional-valued elements.

For cases (i) and (ii), we can bypass further sub-problems of LP_1 since these will not yield a solution to the MIO problem (4.5). In other words, all nodes of the binary search tree along the branch implied by LP_1 can be pruned and need not be further considered. For case (iii), we can update the best known feasible solution to the MIO problem (4.5) as the optimal solution to the problem LP_1 . For case (iv), the sub-domain $\Gamma_0 \cap \{(\beta, \gamma, v) : v_j = 0\}$ may still contain an optimal solution. Therefore, in case (iv), we branch on a fraction-valued component of the LP_1 solution for v to create further two sub-problems and then repeat this process as described above.

C Simulation Results for the Heteroskedastic Error Design

In this part of the appendix, we report the simulation results under the setup of DGP(ii) of Section 5, where the distribution of the error term ε admits a heteroskedasticity component. Table 7 gives the MIO computation time statistics for solving the PRESCIENCE(q) problem under DGP(ii). Compared to the results of Table 1, the PRESCIENCE problem appeared to be more computationally difficult for the high dimensional setup in the DGP(ii) design where the maximum computation time could exceed 2.5 hours. However, the mean and median computation time remained well capped below 6 minutes across all cases in Table 7. In fact, the case of the MIO computation lasting over one hour appeared in only 3 out of the 100 repetitions for the PRESCIENCE(3) simulations in the setup of $p = 60$.

We compare in Tables 8 and 9 the predictive and variable selection performance results for the various prediction methods given in (5.2). The overall picture revealed in these two tables is similar to that given in Tables 2 and 3. In particular, the PRESCIENCE_CV approach still had the best overall performance among the six prediction approaches. We also note that the PRESCIENCE_CV approach could outperform the logit_lasso and probit_lasso approaches by a large margin in both the in-sample and out-of-sample predictive performances for the DGP(ii) design within a high dimensional variable selection setup. It is well known in

Table 7: PRESCIENCE(q) computation time (CPU seconds) under DGP(ii)

q	$p = 10$			$p = 60$		
	1	2	3	1	2	3
mean	0.39	1.55	0.73	3.47	68.53	350.4
min	0.05	0.04	0.03	0.21	0.03	0.04
median	0.40	0.76	0.38	2.77	23.07	50.13
max	1.02	20.34	9.97	16.56	417.2	8552

Table 8: Comparison of prediction methods under DGP(ii) with $p = 10$

method	PRESCIENCE(q)			PRESCIENCE_CV	logit_lasso	probit_lasso
	$q = 1$	$q = 2$	$q = 3$			
<i>Corr_sel</i>	0.86	0.95	0.95	0.91	0.91	0.88
<i>Orac_sel</i>	0.86	0.01	0	0.53	0.09	0.09
<i>Num_irrel</i>	0.14	1.04	2.02	0.67	2.93	2.49
<i>in_Score</i>	0.834	0.871	0.894	0.860	0.787	0.776
<i>in_RS</i>	1.095	1.144	1.175	1.131	1.031	1.016
<i>out_Score</i>	0.724	0.711	0.696	0.716	0.673	0.668
<i>out_RS</i>	0.948	0.930	0.910	0.937	0.881	0.874

Table 9: Comparison of prediction methods under DGP(ii) with $p = 60$

method	PRESCIENCE(q)			PRESCIENCE_CV	logit_lasso	probit_lasso
	$q = 1$	$q = 2$	$q = 3$			
<i>Corr_sel</i>	0.76	0.82	0.88	0.82	0.64	0.63
<i>Orac_sel</i>	0.76	0	0	0.41	0	0
<i>Num_irrel</i>	0.24	1.18	2.12	0.88	4.53	4.17
<i>in_Score</i>	0.842	0.894	0.927	0.880	0.766	0.759
<i>in_RS</i>	1.103	1.171	1.216	1.153	1.000	0.990
<i>out_Score</i>	0.713	0.693	0.673	0.700	0.608	0.608
<i>out_RS</i>	0.934	0.907	0.881	0.917	0.797	0.796

the binary prediction literature (see e.g. Elliott and Lieli, 2013) that the optimal prediction rule in terms of score maximization does not hinge on knowing the true distribution of Y given W , and binary prediction based on the MLE approach with a misspecified likelihood can yield poor predictive performance. Noting that the logit and probit likelihoods are both misspecified under DGP(ii), our simulation results thus suggest that this known fact may also carry well through to the context with variable selection.

D Additional Simulations on the Performance of the Warm-Start MIO Approaches to the PRESCIENCE Problem

In Appendix D, we conduct a simulation study on the performance of adopting the warm-start strategy of Section 4.3 in solving the MIO formulations (4.5) and (4.14). We used the setup of DGP(ii) of Section 5 for the simulation design. For all simulation experiments in this section, we used a training sample of $n = 100$ observations over which we computed the exact solutions to all the MIO problems. We set p , the dimension of the vector of auxiliary covariates, to be 10. We used the space (5.1) as the parameter space Θ for the cold-start MIO solution approach to the PRESCIENCE problem. For the warm-start MIO formulations, we set $\tau = 1.5$ and constructed $(\hat{P}_i)_{i=1}^n$ using the fitted choice probabilities from the logit regression of Y on the entire covariate vector W to derive the space $\hat{\Theta}(\tau)$ as a refinement of the initial parameter space Θ . The number of simulation repetition was set to be 100.

We now present the simulation results. Table 10 gives the summary statistics of the MIO computation time in CPU seconds and the average of the maximized scores over all the simulation repetitions. From this table, we note that, fixing the start method, the average of maximized scores under the MIO formulation (4.5) was identical to that under the MIO formulation (4.14). In fact, the maximized score values computed under the two MIO formulations were also identical across all the simulation repetitions. This matched the mathematical equivalence between the MIO problems (4.5) and (4.14). Across the start method, we find that the warm start approach, which is based on a smaller parameter space, could miss the global optimum. However, the difference of the cold and warm-start based maximized objective values was very small. In fact, in no more than 7% of the simulation repetitions did we observe the occurrence of such differences among which the maximal difference was about 0.02.

On the other hand, we observe significant reduction of computation time from employing the warm start strategy across nearly all cases in Table 10. However, the cold start approach might sometimes be less computationally costly than its warm start version. Yet this happened only in the simulation with the MIO formulation (4.5) and its occurrence was rare (in only 1 out of 100 repetitions). Therefore, we believe that the warm start method may be a useful heuristic device to improve the computational efficiency for the MIO based computation of PRESCIENCE. Finally, concerning the performance comparison of the MIO formulations (4.5) and (4.14), Table 10 reveals that the formulation (4.5) tended to outperform the formulation (4.14) in terms of computation time in the cold-start setting. This tendency appeared to be reversed in the warm-start setting.

Table 10: Performance comparison of cold and warm-start MIO formulations

	cold start			warm start		
q	1	2	3	1	2	3
MIO formulation (4.5)						
maximized score	0.812	0.837	0.855	0.811	0.836	0.854
MIO computation time						
mean	1.44	40.5	149	0.49	4.30	43.4
min	0.64	1.68	0.57	0.07	0.11	0.18
median	1.33	33.5	94.5	0.36	2.11	6.54
max	3.18	223	1102	1.31	59.7	3074
MIO formulation (4.14)						
maximized score	0.812	0.837	0.855	0.811	0.836	0.854
MIO computation time						
mean	7.78	79.2	401	0.55	4.32	19.4
min	0.55	1.09	0.42	0.05	0.09	0.11
median	4.66	63.2	242	0.40	1.92	5.05
max	31.3	868	3153	4.28	35.7	367

E Empirical Illustration for Best Subset Selection Using Linear Specification

In this section of the appendix, we report empirical results using the linear specification of covariates. Specifically, we use the same focused covariates as constructed in the main text; for auxiliary covariates Z , we consider the simple specification where $Z = (CARS, DOVTT, DIVTT)$. Under this setup, we have that $k = 1$ and $p = 3$. Following Florios and Skouras (2008), we set all unknown parameters to be

within the range $[-10, 10]$ and took the parameter space Θ to be $[-10, 10]^4$. The refined space $\widehat{\Theta}(\tau)$ was computed accordingly as described in Section 4.3. Table 11 presents the resulting refined parameter bounds derived from $\widehat{\Theta}(\tau)$.

Table 11: Refined parameter bounds ($\tau = 1.5$)

Covariate specification: $k = 1, p = 3$

Variable	lower bound	upper bound
Intercept	-7.8275	7.8275
<i>CARS</i>	-5.4143	5.4143
<i>DOVTT</i>	-1.9229	1.9229
<i>DIVTT</i>	-0.7884	0.7884

We can clearly see from Table 11 that using the refined space $\widehat{\Theta}(\tau)$ helped to reduce the parameter search space in both of the MIO formulations (4.5) and (4.14). The extent of this reduction could be quite large even when the enlargement parameter τ was set to be 1.5. In fact, the size of $\widehat{\Theta}(\tau)$ was merely about 0.64% of that of Θ . Therefore, we can anticipate considerably computational efficiency gain from using the warm-start MIO formulations.

In Table 12, we present comparative results of the warm-start and cold-start approaches for the MIO formulations in (4.5) specified with different values of the cardinality bound q . Since there are only 4 unknown parameters in this simple setup, we set $\varepsilon = 0$ in (4.12) and solved for the exact PRESCIENCE.

To interpret the results of Table 12, first note that $MIO_gap = 0$ for all MIO problems in this table. Thus all these MIO solutions were exact; moreover, both the cold-start and warm-start MIO approaches yielded the same maximized objective values and the parameter estimates for auxiliary covariates indeed respected the ℓ_0 -norm constraint specified in (2.4). The parameter estimates solved by these two different approaches were very similar for the cases of $q \in \{1, 3\}$. Since the maximum score objective function is a step function, it is not surprising to have multiple solutions and thus the PRESCIENCE for a given value of q need not be unique. This can be clearly seen from the case of $q = 2$, where the results of the cold-start and warm-start methods differed in the covariate to be excluded from the corresponding PRESCIENCE.

We now assess the computational efficiency of the warm-start and cold-start approaches. From Table 12, we can see that both approaches performed very well. Most of the MIO cases considered in this table were solved in few minutes and the case taking the longest time was also solved in about half an hour. We also notice

Table 12: Implementation Using MIO formulation (4.5)

Covariate specification: $k = 1, p = 3$

q	1		2		3	
MIO start method	warm	cold	warm	cold	warm	cold
focus covariates						
<i>DCOST</i>	1	1	1	1	1	1
Intercept	3.4654	3.4258	3.2447	5.2013	4.7953	4.9267
auxiliary covariates						
<i>CARS</i>	2.7420	2.7191	2.7470	4.3828	3.6571	3.8179
<i>DOVTT</i>	0	0	0	0.9278	0.7951	0.7952
<i>DIVTT</i>	0	0	0.6653	0	0.3830	0.4978
in-sample performance						
maximized score	756	756	763	763	765	765
maximized average score	0.8979	0.8979	0.9062	0.9062	0.9086	0.9086
MIO solver output						
<i>MIO_gap</i>	0	0	0	0	0	0
CPU time (in seconds)	15	13	112	903	253	1887
branch-and-bound nodes	6321	6120	62630	304708	129947	887349

from these results that, except for the case of $q = 1$ under which both approaches were comparable, the cold-start formulation was clearly outperformed by its corresponding warm-start version. The difference in computational efficiency can be sizable: the warm-start approach just took about 12% (13%) of the time and 20% (14%) of the branch-and-bound nodes used by the cold-start approach to solve the $q = 2$ ($q = 3$) case.

In this empirical application, the computational merit of using the refined parameter space is also evident for the PRESCIENCE implementation using the MIO formulation (4.14). The results for this formulation are summarized in Table 13.

As in Table 12, the parameter estimates from both the warm-start and cold-start approaches in Table 13 are very similar for the cases $q \in \{1, 3\}$. For these two cases, the parameter estimates as displayed in Tables 12 and 13 are also qualitatively and quantitatively similar across the MIO formulations (4.5) and (4.14). For the case of $q = 2$, in contrast to the results for the formulation (4.5), the variable being excluded is the same for both warm-start and cold-start versions of the formulation (4.14). We notice that the case of $q = 3$ reduces to the maximum score estimation problem using all covariates; for this case, our MIO estimates are quite similar to those computed by Florios and Skouras (2008) and both our and

Table 13: Implementation Using MIO formulation (4.14)

Covariate specification: $k = 1, p = 3$

q	1		2		3	
MIO start method	warm	cold	warm	cold	warm	cold
focus covariates						
<i>DCOST</i>	1	1	1	1	1	1
Intercept	3.2803	3.2803	4.6701	5.2013	4.8270	4.9267
auxiliary covariates						
<i>CARS</i>	2.4666	2.4667	3.4212	4.3828	3.7000	3.8179
<i>DOVTT</i>	0	0	0.9542	0.9278	0.8141	0.7952
<i>DIVTT</i>	0	0	0	0	0.3211	0.4978
in-sample performance						
maximized score	756	756	763	763	765	765
maximized average score	0.8979	0.8979	0.9062	0.9062	0.9086	0.9086
MIO solver output						
<i>MIO_gap</i>	0	0	0	0	0	0
CPU time (in seconds)	62	470	81	6428	144	2052
branch-and-bound nodes	38163	139833	89977	4867501	154390	1539205

their maximized objective values are identical.

It can be noticed that, across all cases in Tables 12 and 13, the variable *CARS* was always selected and its parameter estimate was of a much larger magnitude than those of other selected auxiliary covariates. This indicates that *CARS* is the most important variable among the three auxiliary covariates. Moreover, by comparing the maximized average scores derived under difference cases of q , there is very little loss in the goodness of fit from adopting the parsimonious specification using only *CARS* as the auxiliary covariate.

Regarding the computational efficiency, we can clearly see from the results across all cases of q in Table 13 that, for the MIO formulation (4.14), there was huge performance gain from using the warm-start approach in terms of reduction of the CPU time and branch-and-bound nodes. Putting together the results from both Tables 2 and 3, we thus find that, regardless of the MIO formulations (4.5) and (4.14), it is generally far more computationally efficient to adopt the warm-start strategy in the implementation.

We now compare computational performance across the formulations (4.5) and (4.14). From Tables 12 and 13, we can see that, for all three cases of q , it took far fewer branch-and-bound nodes to solve the formulation (4.5) than to solve the formulation (4.14). However, in terms of usage of the CPU time, the

former formulation did not completely dominate the latter. This can be intuitively explained as follows.

There are $2n$ inequalities stated in (4.6) of the formulation (4.5). By contrast, the number of inequalities in (4.15) of the formulation (4.14) is only half of that amount. Hence, the corresponding LP relaxation problems in the branch-and-bound solution procedure for the formulation (4.5) are likely to be tighter than those for the formulation (4.14). This would help to reduce the number of branching steps required to reach the optimum. On the other hand, for the formulation (4.5), it may take much longer to solve at each node the corresponding LP relaxation problem which contains a massive amount of inequality constraints. Thus, there is a tradeoff between the computational cost per node and the total number of required nodes in the solution procedure. This tradeoff may depend on the sample size, the support of the data and the variable selection bound q . Therefore, we find that each of these two MIO formulations has its strength and hence both complement each other for solving the PRESCIENCE problems.