

Rosen, Adam M.; Ura, Takuya

Working Paper

Finite sample inference for the maximum score estimand

cemmap working paper, No. CWP22/20

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Rosen, Adam M.; Ura, Takuya (2020) : Finite sample inference for the maximum score estimand, cemmap working paper, No. CWP22/20, Centre for Microdata Methods and Practice (cemmap), London,
<https://doi.org/10.1920/wp.cem.2020.2220>

This Version is available at:

<https://hdl.handle.net/10419/241897>

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.

Finite Sample Inference for the Maximum Score Estimand

Adam M. Rosen
Takuya Ura

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP22/20



Economic
and Social
Research Council

Finite Sample Inference for the Maximum Score Estimand*

Adam M. Rosen

Takuya Ura

Duke University and CeMMAP[†]

University of California, Davis[‡]

May 8, 2020

Abstract

We provide a finite sample inference method for the structural parameters of a semiparametric binary response model under a conditional median restriction originally studied by Manski (1975, 1985). Our inference method is valid for any sample size and irrespective of whether the structural parameters are point identified or partially identified, for example due to the lack of a continuously distributed covariate with large support. Our inference approach exploits distributional properties of observable outcomes conditional on the observed sequence of exogenous variables. Moment inequalities conditional on this size n sequence of exogenous covariates are constructed, and the test statistic is a monotone function of violations of sample moment inequalities. The critical value used for inference is provided by the appropriate quantile of a known function of n independent Rademacher random variables. Simulation studies compare the performance of the test to two alternative tests using an infeasible likelihood ratio statistic and Horowitz's (1992) smoothed maximum score estimator.

JEL classification: C12 C14.

Keywords: Finite sample inference, Maximum score estimation, Moment inequalities, Partial identification.

*We thank Tim Armstrong, Federico Bugni, Bryan Graham, Michal Kolesár, Sokbae Lee, Matt Masten, Francesca Molinari, Ulrich Müller, Whitney Newey, Pepe Montiel Olea, Chris Sims, and seminar and conference participants at Columbia, Duke, Penn State, Princeton, Vanderbilt, Yale, the 2019 Triangle Econometrics Conference, the 2019 Southern Economic Association Conference, the 2019 California Econometrics Conference, the 2019 Young Econometricians Conference, the Cemmap/WISE Workshop on Advances in Econometrics, and the University of Tokyo Workshop on Advances in Econometrics for insightful comments and discussion. Kuong (Lucas) Do and Cheuk Fai Ng provided excellent research assistance. Financial support from the Economic and Social Research Council ESRC Large Research Grant ES/P008909/1 to the Centre for Microdata Methods and Practice is gratefully acknowledged.

[†]Address: Adam Rosen, Department of Economics, Duke University, 213 Social Sciences Box 90097, Durham, NC 27708; Email: adam.rosen@duke.edu

[‡]Address: Takuya Ura, Department of Economics, University of California, Davis, One Shields Avenue, Davis, CA 95616; Email: takura@ucdavis.edu

1 Introduction

In Chapter 41 of Volume 4 of the *Handbook of Econometrics* on the estimation of semiparametric models, Powell (1994) on page 2488 cites Manski (1975) as the earliest example of semiparametric analysis of limited dependent variable models. Manski (1985) provided further analysis for the binary outcome version of the model, in which the outcome is determined by the linear index threshold-crossing specification

$$Y = 1\{X\beta + U \geq 0\},$$

for observable variables $Y \in \{0, 1\}$ and X a row vector in \mathbb{R}^K , where the unobservable variable U is restricted to satisfy the zero conditional median restriction

$$\text{median}(U \mid X) = 0.$$

This semiparametric model is thus distribution-free with regard to unobservable U .¹ Full stochastic independence between U and X is not required, allowing for the conditional distribution of U given $X = x$ to vary with the conditioning value x , and thus accommodating general forms of heteroskedasticity. Under both a rank condition and a large support condition on a continuous regressor Manski (1985) established point identification of β as well as the large deviations convergence rate of the maximum score estimator.

Several further analyses of the maximum score and similar estimators for this and closely related semiparametric binary response models have since been provided, and the literature on the *asymptotic* properties of the maximum score estimator is now vast. Kim and Pollard (1990) showed that the convergence rate of the maximum score estimator is cube-root and established its nonstandard asymptotic distribution after appropriate centering and scaling. Horowitz (1992) developed a smoothed maximum score estimator that converges faster than the $n^{-1/3}$ rate and is asymptotically normal under some additional smoothness assumptions. Additional papers that study large sample estimation and inference applicable in the maximum score context include Manski and Thompson (1986), Delgado, Rodríguez-Poo, and Wolf (2001), Abrevaya and Huang (2005), Léger and MacGibbon (2006), Komarova (2013), Blevins (2015), Jun, Pinkse, and Wan (2015, 2017), Chen and Lee (2018, 2019), Patra, Seijo, and Sen (2018), Seo and Otsu (2018), Cattaneo, Jansson, and Nagasawa (2019), and Mukherjee, Banerjee, and Ritov (2019).

In contrast to prior approaches for inference on β that employ asymptotic distributional approximations, in this paper we develop a method for conducting *finite sample* inference on β . To do this we employ a conditional moment inequality characterization of the observable implications of the binary response model in the finite sample. Moment inequality characterizations of the model's implications have been previously used by Komarova (2013), Blevins (2015), and Chen and Lee (2019), but none of these papers proposed an inference method which is valid in the finite sample. As was the case in the analysis provided in these papers, we do not require that β is point identified. For instance, we do not require that any component of X is continuously distributed, much less with large support.

In fact, even if β is point identified, and regardless of the support of X in the population, the set of observed values of X in any finite sample is discrete. Indeed, Manski (1985, page 320) defines, “the maximum score estimate \hat{B}_n to be the *set* of solutions to the problem $\max_{b \in \mathcal{B}} S_n(b)$ ” where \mathcal{B} is the parameter space and $S_n(\cdot)$ denotes the sample score function.² He shows that if the given sufficient conditions for point

¹As noted in Manski (1985), his analysis easily generalizes to cover the restriction that the conditional τ th quantile of U given X is 0, where $\tau \in (0, 1)$ is known. The analysis in this paper can be similarly generalized.

²Manski (1985) used B to denote the parameter space and upper case N to denote sample size, which we have changed to

identification hold, then the distance between \hat{B}_n and β converges almost surely to zero, implying consistency of any sequence of $\hat{\beta}_n \in \hat{B}_n$ for β . Intuitively, the set of possible maximum score point estimators shrinks to a point as $n \rightarrow \infty$. Given that our aim in this paper is to conduct finite sample inference, we must own up to the fact that even if β is point identified, there is a proper set of values b to which β is observationally equivalent on the basis of only values of X observed in the finite sample.

We thus introduce the concept of the *finite sample identified set* as the set of parameter vectors $b \in \mathcal{B}$ that satisfy the observable implications of the binary response model *conditional* on a size n sequence of observable covariate vectors $\mathcal{X}_n \equiv (X_1, \dots, X_n)$. This set can be thought of as what the population identified set would be if the support of X were in fact \mathcal{X}_n . It differs from the set estimator \hat{B}_n because it is not a function of the outcomes Y_1, \dots, Y_n , instead only incorporating the observable implications of the model for the distribution of outcomes conditional on \mathcal{X}_n . This newly defined finite sample identified set is useful because our finite sample inference approach is driven only by observable implications regarding Y_1, \dots, Y_n conditional on \mathcal{X}_n , and will be explicit in not being able to detect violations of conditional moment inequalities that condition on values of X not observed in the sample.

The approach taken here exploits the implication of the binary response model that conditional on \mathcal{X}_n , each outcome Y_i is distributed Bernoulli with parameter $p(X_i, \beta) \equiv \mathbb{P}(U_i \geq -X_i\beta \mid \mathcal{X}_n)$. In practice the Bernoulli probabilities $p(X_i, \beta)$ are unknown. Nonetheless, conditional on \mathcal{X}_n , each $p(X_i, \beta)$ is bounded from above or below by $1/2$ according to the sign of $X_i\beta$. Consequently, for any known nonnegative-valued function $g(\cdot) : \mathcal{X}_n \rightarrow \mathbb{R}$, the finite sample distributions of $\omega_{ui}(\beta, g) \equiv (2Y_i - 1)1\{X_i\beta \geq 0\}g(X_i)$ and $\omega_{li}(\beta, g) \equiv (1 - 2Y_i)1\{X_i\beta \leq 0\}g(X_i)$ conditional on \mathcal{X}_n can be bounded from below. The test statistic $T_n(b)$ that we use to implement our test of the null hypothesis $H_0 : \beta = b$ is a supremum of weighted sample averages of $-\omega_{ui}(\beta, g)$ and $-\omega_{li}(\beta, g)$, where the supremum is taken over particular collections of functions $g(\cdot)$. The test statistic $T_n(b)$ is shown to be bounded from above by a function $T_n^*(b)$ of n independent Rademacher random variables, such that the finite sample distribution of $T_n^*(b)$ given \mathcal{X}_n is known. Then, under the null hypothesis $\beta = b$, we have

$$\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n) \leq \alpha,$$

where $q_{1-\alpha}$ is the $1 - \alpha$ quantile of $T_n^*(b)$ given \mathcal{X}_n . We establish that if particular functions $g(\cdot)$ are used, the moment functions which $T_n(b)$ incorporates preserve all the identifying information contained in the finite sample identified set. We further establish a power result for alternatives that lie outside this set.

For the sake of comparison we also consider likelihood ratio tests. Optimality of the likelihood ratio test for the null hypothesis of $\beta = b$ depends on the specification of the alternative hypothesis. The likelihood ratio test is a most powerful test when the alternative hypothesis is simple, meaning that it specifies a unique conditional distribution $\mathbb{P}(Y_1, \dots, Y_n \mid \mathcal{X}_n)$. However, when the alternative hypothesis is of the form $\beta \neq b$ or $\beta = \tilde{b}$, many possible distributions of $\mathbb{P}(Y_1, \dots, Y_n \mid \mathcal{X}_n)$ are permitted under the conditional median restriction. These alternatives are therefore composite. The likelihood ratio test of the null hypothesis $\beta = b$ against either of these alternatives corresponds to a coin flip independent of the data, and has only trivial power against any distribution under either of these composite alternative hypotheses. We establish a lower bound on the power of our test, and we show that this power bound is increasing in the degree to which a hypothesized parameter vector b violates the inequalities that characterize the finite sample identified set. Thus values of b that are sufficiently far from the finite sample identified set by this measure are guaranteed

\mathcal{B} and lower case n to match our notation.

to be rejected with probability exceeding the size of the test.

Among the aforementioned papers from the literature on maximum score, the most closely related is that of Chen and Lee (2019), who also cast the implications of Manski’s (1985) model as conditional moment inequalities for the sake of delivering a new insight, albeit one that is entirely different from ours. Chen and Lee (2019) expand on the conditional moment inequalities used by Komarova (2013) and Blevins (2015) to develop a novel conditional moment inequality characterization of the identified set which involves conditioning on two linear indices instead of on the entire exogenous covariate vector. They apply intersection bound inference from Chernozhukov, Lee, and Rosen (2013) to this conditional moment inequality characterization to achieve asymptotically valid inference. This cleverly exploits the model’s semiparametric linear index restriction in order to side step the curse of dimensionality. Although a good deal of focus is given to Manski’s (1985) binary response model, their method can also be applied to other semiparametric models.

To the best of our knowledge, this paper is the first to propose a finite sample inference method for β in Manski’s (1985) semiparametric binary response model. The test proposed is shown to control size for any sample size n . This paper is also the first to introduce the concept of a finite sample identified set, both within the context of the semiparametric binary response model and more broadly, explicitly defining the set of model parameters logically consistent with the modeling restrictions and only information that can be gathered from observable implications conditional on realizations of exogenous variables observed in the finite sample. There are however a handful of precedents for employing finite sample inference with other partially identifying models. Manski (2007) considers the problem of predicting choice probabilities for the choices individuals would make if subjected to counterfactual variation in their choice sets. In the absence of the structure afforded by commonly used random utility models, he shows that counterfactual choice probabilities are partially identified, and proposes a procedure for inference using results from Clopper and Pearson (1934). Chernozhukov, Hansen, and Jansson (2009) propose a finite sample inference method in quantile regression models in which the outcome is continuously distributed. Their approach exploits a “conditionally pivotal property” to bound the finite sample distribution of a GMM criterion incorporating moment equalities, but which does not require point identification for its validity. Syrgkanis, Tamer, and Ziani (2018) conduct inference on partially identified parameters of interest in auction models imposing weak assumptions on bidders’ information. They propose a method to conduct finite sample inference on moments of functions of the underlying valuation distribution using concentration inequalities. Armstrong and Kolesár (2018) provide methods for optimal inference on average treatments that are finite sample valid in the special case in which regression errors are normal, and asymptotically valid more generally. Their conditions cover cases where identification may fail due to lack of overlap of the support of conditioning variables. The approach taken in this paper for finite sample inference in the context of Manski’s (1985) binary response model is different from all of these.

The rest of this paper is organized as follows. Section 2 formally sets out the testing problem and the moment inequality representation of the finite sample identified set. Section 3 lays out the construction of the test statistic and corresponding critical value, and establishes the finite sample validity of the test. It also establishes a finite sample (lower) power bound for our test. Section 4 considers the likelihood ratio test. Section 5 demonstrates the performance of our approach by reporting results from Monte Carlo simulations with comparison to inference using an infeasible likelihood ratio test and Horowitz’s (1992) smoothed maximum score estimator. Section 6 concludes and discusses avenues for future research. All proofs are in the Appendix. Unless otherwise stated, our analysis throughout this paper should be read as conditional on observable covariate vectors $\mathcal{X}_n \equiv (X_1, \dots, X_n)$.

2 Model and Moment Restrictions

This section is divided into three subsections, the first of which formally presents the modeling restrictions imposed. The second subsection describes the observable implications of the binary response model *conditional* on a size n sequence of covariate vectors, \mathcal{X}_n , in contrast to those observable implications obtainable from knowledge of the population distribution of observable variables. Based on these observable implications, this second subsection introduces our definition of the finite sample identified set. It clarifies what violations of our model's implications the proposed test can feasibly detect, which is useful for power considerations. The third subsection uses the special case in which covariates X_i are bivariate as an illustrative example. The developments of Subsections 2.2 and 2.3 are however not essential for establishing size control of the test presented in Section 3.

2.1 Model

The following assumption formalizes the restrictions of the semiparametric binary response model under study and the requirements on the sampling process. We maintain these assumptions throughout this paper.

Assumption 1. (i) Random vectors $\{(Y_i, X_i, U_i) : i = 1, \dots, n\}$ reside on a probability space $(\Omega, \mathfrak{F}, \mathbb{P})$, where \mathfrak{F} contains the Borel sets on Ω . (ii) Variables $\{(Y_i, X_i) : i = 1, \dots, n\}$ are observed. (iii) There is a column vector $\beta \in \mathbb{R}^K$ such that $\mathbb{P}(Y_i = 1\{X_i\beta + U_i \geq 0\} \mid \mathcal{X}_n) = 1$ and $\mathbb{P}(U_i \geq 0 \mid \mathcal{X}_n) = 1/2$ for every $i = 1, \dots, n$, where $\mathcal{X}_n \equiv (X_1, \dots, X_n)$. (iv) There is a known set $\mathcal{B} \subseteq \mathbb{R}^K$ to which β belongs. (v) The unobservable variables $(1\{U_1 \geq 0\}, \dots, 1\{U_n \geq 0\})$ are mutually independent conditional on \mathcal{X}_n .

The requirements of Assumption 1 are slightly weaker than the assumptions used in the existing literature (e.g., Manski, 1975, 1985). Parts (i), (ii) and (iv) are standard. Although it is not necessary in this paper because we employ partial identification analysis, the parameter space \mathcal{B} can be restricted by imposing one of the usual scale normalizations from the literature, such as $|b_1| = 1$ for all $b \in \mathcal{B}$. Part (iii) imposes the binary response structure and the requirement that $\mathbb{P}(U_i \geq 0 \mid \mathcal{X}_n) = 1/2$ for each i . Binary response models typically require that U_i is continuously distributed in a neighborhood of zero, in which case this is implied by the usual conditional median restriction. Strictly speaking, we do not need to impose that each U_i is continuously distributed at zero, and hence we replace the median restriction with this requirement. Part (v) holds if (Y_i, X_i, U_i) are independent and identically distributed, but is much more general. The observations $\{(Y_i, X_i) : i = 1, \dots, n\}$ can be non-independent and non-identically distributed. Indeed, even if $\{(Y_i, X_i) : i = 1, \dots, n\}$ were restricted to be i.i.d., it would remain the case that $\{Y_i : i = 1, \dots, n\}$ would not be i.i.d. conditional on \mathcal{X}_n . Throughout the text, $\mathbb{E}[\cdot]$ is used to denote population expectations taken with respect to \mathbb{P} , and $\mathbb{E}_n[\cdot] \equiv n^{-1} \sum_{i=1}^n [\cdot]$.

The power result presented in Theorem 5 in Section 3 and the results of Section 4 on the likelihood ratio test additionally invoke the following assumption.

Assumption 2. The unobservable variables (U_1, \dots, U_n) are mutually independent conditional on \mathcal{X}_n .

This restriction is common in the prior literature on maximum score estimation, but is not required for many of the results in this paper, and in particular is not necessary to establish size control for our test. The assumption is satisfied for example in models that restrict (X_i, U_i) to be i.i.d.³ Note however that Assumption 2 does not require U_1, \dots, U_n to be i.i.d. given \mathcal{X}_n , but only mutually independent.

³For instance this is implied by Assumption 3 of Manski (1985).

2.2 Observable Implications Conditional on \mathcal{X}_n

To conduct finite sample inference, we focus solely on the implications obtainable from a sequence of n draws of (Y, X) in a sample $\{(Y_i, X_i) : i = 1, \dots, n\}$ and not on features of the population distribution of these variables that could only be obtained on the basis of infinitely many observations. Consequently our focus is not on the identified set that could be obtained from knowledge of the population distribution of (Y, X) in an infinitely large population, but rather on the set obtainable solely from knowledge of a size n sample of observations in accord with Assumption 1. By definition, this is the set of parameter vectors $b \in \mathcal{B}$ such that the distribution of Y_1, \dots, Y_n conditional on \mathcal{X}_n matches that of $\tilde{Y}_i \equiv 1\{X_i b + \tilde{U}_i \geq 0\}$ for a sequence of random variables $\tilde{U}_1, \dots, \tilde{U}_n$ that satisfy the restrictions placed on the conditional distribution of U_1, \dots, U_n in Assumption 1. We refer to this set as the finite sample identified set and denote it as \mathcal{B}_n^* .

Definition 1. *The **finite sample identified set** for β , denoted \mathcal{B}_n^* , is the set of $b \in \mathcal{B}$ for which there exist random variables $\{\tilde{Y}_i : i = 1, \dots, n\}$ and $\{\tilde{U}_i : i = 1, \dots, n\}$ such that:*

- (i): $\mathbb{P}(\tilde{Y}_i = 1\{X_i b + \tilde{U}_i \geq 0\} \mid \mathcal{X}_n) = 1$,
- (ii): $(\tilde{Y}_1, \dots, \tilde{Y}_n)$ and (Y_1, \dots, Y_n) have the same distribution conditional on \mathcal{X}_n ,
- (iii): $\mathbb{P}(\tilde{U}_i \geq 0 \mid \mathcal{X}_n) = 1/2$ for every $i = 1, \dots, n$,
- (iv): $\{1\{\tilde{U}_i \geq 0\} : i = 1, \dots, n\}$ are mutually independent given \mathcal{X}_n .

Our next task is to express \mathcal{B}_n^* with a moment inequality representation useful for inference. The following lemma sets out two observable implications that will be useful for this purpose.

Lemma 1. *Under Assumption 1,*

$$X_i \beta \geq 0 \implies \mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] \geq 0, \quad (2.1)$$

$$X_i \beta \leq 0 \implies \mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] \leq 0. \quad (2.2)$$

From the inequalities of the lemma, it further follows that if $X_i \beta = 0$ then $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] = 0$. Moreover (2.1) and (2.2) and the implications of them described above hold with β replaced by any b that is an element of the finite sample identified set \mathcal{B}_n^* . This can be proven by following precisely the same steps as in the proof of the lemma with \tilde{U}_i from Definition 1 replacing U_i .

With Lemma 1 in hand, the following theorem provides a moment inequality characterization of the finite sample identified set.

Theorem 1. *Under Assumption 1, the finite sample identified set for β is*

$$\mathcal{B}_n^* = \{b \in \mathcal{B} : \mathbb{E}[(2Y_i - 1)1\{X_i b \geq 0\} \mid \mathcal{X}_n] \geq 0 \geq \mathbb{E}[(2Y_i - 1)1\{X_i b \leq 0\} \mid \mathcal{X}_n] \text{ for every } i = 1, \dots, n\}.$$

The conditional moment inequalities characterizing \mathcal{B}_n^* in Theorem 1 are equivalent to (2.1) and (2.2) for all $i = 1, \dots, n$. However, using this conditional moment inequality representation to conduct inference on β is complicated by the fact that in a sample of n observations the distribution of Y_i given \mathcal{X}_n can vary across i , even if $(Y_i, X_i) : i = 1, \dots, n$ are identically distributed, and there is only one observation of (Y_i, X_i) for each i .

Some level of aggregation of these implications across i is therefore required. One way to do this is to interact the expressions inside the conditional expectation operators in Theorem 1 with nonnegative-valued

functions of the exogenous covariates and take sample averages. For this purpose, define $\{g_u(\cdot, v) : v \in \mathcal{V}_u\}$ and $\{g_l(\cdot, v) : v \in \mathcal{V}_l\}$ to be collections of such non-negative instrument functions indexed by $v \in \mathcal{V}_u$ and $v \in \mathcal{V}_l$, respectively. That is, for any $v \in \mathcal{V}_u$, $g_u(\cdot, v) : \mathbb{R}^K \rightarrow \mathbb{R}_+$ is an instrument function mapping from K dimensional Euclidean space (on which each X_i resides) to the nonnegative reals. Likewise, for each $v \in \mathcal{V}_l$, the function $g_l(\cdot, v)$ also maps to the nonnegative reals. Shortly, particular collections of such functions and corresponding index sets \mathcal{V}_u and \mathcal{V}_l will be defined for construction of our proposed test statistic.

Since the instrument functions are nonnegative-valued, interacting them with the expressions inside the conditional expectations in Theorem 1 and averaging across i will preserve the sign of the conditional expectation. Specifically, consider for any nonnegative valued instrument function $g(\cdot)$ the conditional moment inequalities

$$\mathbb{E}[\mathbb{E}_n[(2Y - 1)1\{Xb \geq 0\}g(X)] \mid \mathcal{X}_n] \geq 0, \quad (2.3)$$

$$\mathbb{E}[\mathbb{E}_n[(1 - 2Y)1\{Xb \leq 0\}g(X)] \mid \mathcal{X}_n] \geq 0. \quad (2.4)$$

These inequalities are valid for all $b \in \mathcal{B}_n^*$ because they are implications of the conditional moment inequalities characterizing \mathcal{B}_n^* in Theorem 1. Indeed, they are both valid for any *collection* of nonnegative-valued instrument functions. A potential drawback to aggregation of the conditional moments however is that (2.3) and (2.4) for just any such collection of positive instrument functions need not in general fully characterize \mathcal{B}_n^* , so that using the latter inequalities can result in a loss of identification power.

Particular collections $\{g_u(\cdot, v) : v \in \mathcal{V}_u\}$ and $\{g_l(\cdot, v) : v \in \mathcal{V}_l\}$ are now defined so that imposing (2.3) and (2.4) ensures preservation of the full identifying power of the conditional-on- \mathcal{X}_n moment inequalities in Theorem 1. These collections differ from those used by Andrews and Shi (2013) for translating the identifying power of conditional moment inequalities to a collection of unconditional moment inequalities. In the present setting, there is no issue of converting inequalities conditional on continuous variables to unconditional ones, because the conditioning set in the inequalities characterizing \mathcal{B}_n^* is finite. Instead, the problem to be addressed is how best to aggregate these implications across observations i given the non-i.i.d. nature of Y_i conditional on \mathcal{X}_n . In constructing our collection of information-preserving instrument functions, we exploit two features specific to the task at hand, namely first that our focus is on finite sample inference conditional on \mathcal{X}_n and second that whether or not $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] \geq 0$ (≤ 0) depends only on whether the linear index $X_i\beta$ is at least (at most) zero. In the existing literature, Pinkse (1993, Section 3.3) has used the second feature for exact computation of the maximum score estimator.

Toward this end, consider the following two sequences of binary indicators:

$$\begin{aligned} r_u(b) &\equiv (1\{X_1b \geq 0\}, \dots, 1\{X_nb \geq 0\}), \\ r_l(b) &\equiv (1\{X_1b \leq 0\}, \dots, 1\{X_nb \leq 0\}). \end{aligned}$$

Irrespective of how $g(X)$ is defined, $\mathbb{E}_n[(2Y - 1)1\{Xb \geq 0\}g(X)] = \mathbb{E}_n[(2Y - 1)1\{Xb' \geq 0\}g(X)]$ whenever $r_u(b) = r_u(b')$ and $\mathbb{E}_n[(1 - 2Y)1\{Xb \leq 0\}g(X)] = \mathbb{E}_n[(1 - 2Y)1\{Xb' \leq 0\}g(X)]$ whenever $r_l(b) = r_l(b')$. Using the functions $r_u(\cdot)$ and $r_l(\cdot)$, we denote by \mathbb{V}_u the coimage of the function $r_u(\cdot)$ on \mathcal{B} , and by \mathbb{V}_l that of the function $r_l(\cdot)$.⁴ That is, taking for instance \mathbb{V}_u , the sequences of inequalities defining unique values for $r_u(\cdot)$, namely $X_ib \geq 0$ and $X_ib < 0$ for $i = 1, \dots, n$, partition \mathbb{R}^K into the collections of sets \mathbb{V}_u . Similarly, the partition \mathbb{V}_l comprises regions on which the function $r_l(\cdot)$ takes the same sequence of ones

⁴The coimage of a function f is defined as the quotient set of the equivalence relation defined by f .

and zeros according to the satisfaction of inequalities $X_i b \leq 0$ and $X_i b > 0$, $i = 1, \dots, n$. The partitions \mathbf{V}_u and \mathbf{V}_l coincide with each other, with the exception of how points at which $X_i b = 0$ are assigned within each partition. This ensures that when $X_i b = 0$, the joint implication of the inequalities (2.1) and (2.2) are captured by our testing procedure.

Thus the partitions \mathbf{V}_u and \mathbf{V}_l are collections of sets defined by whether or not they satisfy sequences of linear inequalities $X_i b \geq 0$ and $X_i b \leq 0$, $i = 1, \dots, n$, respectively. Equivalently, each such set comprises an intersection of n hyperplanes in \mathbb{R}^K . Such partitions are referred to as hyperplane arrangements in the computational geometry literature. Algorithms for their enumeration and computation have been developed, and indeed put to good use recently in econometrics by Gu and Koenker (2018) for the purpose of computing nonparametric maximum likelihood estimators for binary response models. The model studied by Gu and Koenker (2018) differs from the one in this paper, as it allows for random coefficients but requires independence of covariates and unobservable variables, and its focus is on computation rather than finite sample inference. Nonetheless, the same developments from the computational geometry literature on the enumeration of hyperplane arrangements can be employed for efficient computation here.

Since n is finite, there are only finitely many elements of each of \mathbf{V}_u and \mathbf{V}_l . With regard to the inequalities (2.3) and (2.4), all members v of any set in \mathbf{V}_u and all members v of any set in \mathbf{V}_l produce the same values of $\mathbb{E}[\mathbb{E}_n[(2Y - 1)1\{Xb \geq 0\}g_u(X, v)] \mid \mathcal{X}_n]$ and $\mathbb{E}[\mathbb{E}_n[(1 - 2Y)1\{Xb \leq 0\}g_l(X, v)] \mid \mathcal{X}_n]$ for any $b \in \mathcal{B}$. Thus it will suffice for us to work with a single representative from each set, i.e., the full identifying power is preserved as long as \mathcal{V}_u and \mathcal{V}_l have a representative from each element of \mathbf{V}_u and \mathbf{V}_l , respectively. This is formalized in Theorem 2 below.

In order to assess how large the partitions \mathbf{V}_u and \mathbf{V}_l can be, it is helpful to note that in our context each separating hyperplane is of the form $\{v \in \mathbb{R}^K : X_i v = 0\}$, inducing a *homogeneously linearly separable dichotomy*, as described by Cover (1965). From Theorem 1 of that paper, the number of such dichotomies that can be constructed from n points in \mathbb{R}^K is bounded from above by

$$C(n, K) \equiv 2 \sum_{j=0}^{K-1} \binom{n-1}{j}.$$

This is equivalently the upper bound on the number of elements in each of the partitions \mathbf{V}_u and \mathbf{V}_l , and is attained when every subset of K points from the n points X_1, \dots, X_n are linearly independent.⁵

This upper bound on the cardinality of \mathbf{V}_u and \mathbf{V}_l , and thus on the number of representative points in the sets \mathcal{V}_u and \mathcal{V}_l , can be quite large. Fortunately their computation can be carried out by making use of the aforementioned results from computational geometry on computing an exhaustive set of hyperplane arrangements, also referred to as the vertex enumeration problem. Examples of available algorithms include Avis and Fukuda (1996), Sleumer (1998), and Rada and Černý (2018), as well as a novel computational method proposed by Gu and Koenker (2018) building on Rada and Černý (2018).⁶

With partitions \mathbf{V}_u and \mathbf{V}_l now defined, the following theorem establishes a representation of the finite sample identified set \mathcal{B}_n^* given in Theorem 1 that takes the form of a finite collection of inequalities of the form (2.3) and (2.4).

⁵In this case X_1, \dots, X_n are said to be in general position.

⁶To investigate the performance of our inference method, Monte Carlo experiments are presented in Section 5. Like most of the prior literature, we use designs with two covariates. This constitutes the simplest setting in which to investigate the performance of our inference method, and affords computational tractability. In such cases the use of sophisticated vertex enumeration algorithms is unnecessary, for reasons explained in Section 2.3. We thus do not investigate the use of these algorithms here, but we note their availability for higher dimensional settings.

Theorem 2. *Let Assumption 1 hold and let $b \in \mathcal{B}$. If $b \in \mathcal{B}_n^*$ then*

$$\forall v \in \mathcal{V}_u : \mathbb{E} [\mathbb{E}_n [(2Y - 1)1\{Xb \geq 0\}1\{Xv < 0\}] \mid \mathcal{X}_n] \geq 0, \quad (2.5)$$

and

$$\forall v \in \mathcal{V}_l : \mathbb{E} [\mathbb{E}_n [(1 - 2Y)1\{Xb \leq 0\}1\{Xv > 0\}] \mid \mathcal{X}_n] \geq 0. \quad (2.6)$$

Moreover, if \mathcal{V}_u and \mathcal{V}_l have at least one element from each member of \mathbf{V}_u and \mathbf{V}_l , respectively, then (2.5) and (2.6) imply that $b \in \mathcal{B}_n^*$, so that

$$\mathcal{B}_n^* = \{b \in \mathcal{B} : (2.5) \text{ and } (2.6) \text{ hold}\}.$$

The moment inequalities (2.5) and (2.6) are conditional on \mathcal{X}_n and are thus different from those employed previously in the literature. Our representation is perhaps most closely related to that of Chen and Lee (2019) for the identified set in the underlying population. Their characterization uses inequalities that condition on the values of two linear indices in X : Xb and $X\gamma$, leading to significant dimension reduction when estimating conditional moments employed for asymptotic inference. In this paper our goal is finite sample inference, made operational by conditioning on \mathcal{X}_n . Our construction leading to (2.5) and (2.6) exploits the finite nature of \mathcal{X}_n . This is done by establishing that given \mathcal{X}_n , one can partition the parameter space \mathcal{B} into equivalence classes \mathbf{V}_u and \mathbf{V}_l whose members comprise elements that all produce the same values of the moment functions appearing in (2.5) and (2.6), respectively. Then, using our proposed instrument functions the moment inequalities (2.5) and (2.6) aggregate values of $2Y_i - 1$ and $1 - 2Y_i$ according to whether pairs of indices Xb and Xv in each of the two inequalities disagree in particular directions.

In Section 3 we provide a test statistic that combines the extent to which these moment inequalities are violated when evaluated at any conjectured parameter vector $b \in \mathbb{R}^K$. A confidence set for β can then be constructed by way of test inversion. Note however that for test inversion parameter vectors b and \tilde{b} need only be considered if they do not lie in the same elements of both partitions \mathbf{V}_u and \mathbf{V}_l . Any two such vectors that both reside in the same partitions will produce identical values for all moment inequalities (2.5) and (2.6), and consequently the same value of the test statistic proposed in Section 3.


2.3 A Simple Example: Computation with Two Covariates

When the covariates X_i have only two components the characterization of the relevant hyperplane arrangements is greatly simplified. This makes it an ideal case in which to showcase and demonstrate the inference approach. Because the distribution of the unobservables is nonparametrically specified, a scale normalization may be imposed, e.g., restricting the first component of β to have absolute value of one. Due to the scale normalization, the circumstance in which there are only two components of X_i is the simplest non-trivial case in which to study the semiparametric binary response model of this paper. For this reason, the case in which $\dim(X_i) = 2$ has appeared prominently in simulation studies of maximum score type estimators, and this is also the setting on which we focus in our Monte Carlo investigations in Section 5.

The panels of Figure 1 illustrate the construction of partitions \mathbf{V}_u and \mathbf{V}_l for bivariate X_i in a simple example in which $n = 3$. Panel (a) depicts X_1 , X_2 , and X_3 , into which panel (b) additionally incorporates the values of v_i for which $X_i v_i = 0$. Panel (c) drops the covariate vectors X_i and panel (d) uses different colors to depict the interiors of the elements of \mathbf{V}_u and \mathbf{V}_l . Note that the partitions \mathbf{V}_u and \mathbf{V}_l only differ on their boundaries. Theorem 2 indicates that in order to preserve the full identifying power of the finite

sample identified set \mathcal{B}_n^* , it suffices to employ moment inequalities of the form (2.5) and (2.6) with sets \mathcal{V}_u and \mathcal{V}_l each containing one element from each of these six different colored regions.

Figure 1: A schematic illustration of partitions \mathcal{V}_u and \mathcal{V}_l for $n = 3$.



To understand the further simplification afforded by the scale normalization, first note that for any v with $v_1 \neq 0$ we can normalize the first component of v such that $|v_1| = 1$ in the instrument functions that appear in (2.5) and (2.6) for the same reason that one can normalize the first component of the parameter vector β .⁷ Specifically for any $v \in \mathbb{R}^K$ with $v_1 \neq 0$,

$$X_i v \geq 0 \iff X_{i1} \frac{v_1}{|v_1|} + X_{i2} \frac{v_2}{|v_1|} \geq 0. \quad (2.7)$$

Consequently it suffices in the inequalities of (2.5) and (2.6) to use only values of v with $|v_1| = 1$, and we impose the common scale normalization on the parameter space that \mathcal{B} comprises a compact subset of \mathbb{R}^2 such that for all $b \in \mathcal{B}$, $|b_1| = 1$. Geometrically, this means that we can select the required six representatives v to form \mathcal{V}_u and \mathcal{V}_l from the regions illustrated in Figure 1 panel (d) by focusing solely on values in which the first component is either -1 or 1 .

⁷The alternative normalization that $\|v\| = 1$ could also be used.

For the sake of actually computing such sets of representative values of v , possibly with much larger n , we may proceed making use of the normalization $|v_1| = 1$ imposed as follows. Values of v that are on the boundary of satisfying the inequalities constituent in $r_u(v)$ and $r_l(v)$, namely those satisfying $X_i v = 0$, can be separated into two cases:

$$\begin{aligned} \text{if } v_1 = +1 : \quad X_i v = 0 &\iff -\frac{X_{i1}}{X_{i2}} = v_2 \\ \text{if } v_1 = -1 : \quad X_i v = 0 &\iff \frac{X_{i1}}{X_{i2}} = v_2. \end{aligned}$$

Thus the sequences of indicators, $r_u(v)$ and $r_l(v)$, depend only on where v_2 lies relative to the ordered sequence of values of $-Z_i$ if $v_1 = +1$ and those of Z_i if $v_1 = -1$, where $Z_i \equiv \frac{X_{i1}}{X_{i2}}$.⁸

Consequently, a set of instrument values \mathcal{V}_u that contains one element from each member of \mathbf{V}_u can be obtained by dividing the real line into intervals according to ordered sequences of values of $-Z_i$ and Z_i , $i = 1, \dots, n$, and then collecting all pairs $v = (1, v_2)$ such that v_2 lies in the interior of the first sequence of intervals, and all pairs $v = (-1, v_2)$ such that v_2 lies in the interior of the second sequence of intervals. Specifically, let $\vartheta_1 \leq \dots \leq \vartheta_n$ denote the order statistics of Z_i , so that $\vartheta_1 \equiv \min_i Z_i$ and $\vartheta_n \equiv \max_i Z_i$ and consider the following ordered sequences of intervals:

$$\mathbf{l}_u \equiv \{(-\infty, -\vartheta_n), (-\vartheta_n, -\vartheta_{n-1}), \dots, (-\vartheta_2, -\vartheta_1), (-\vartheta_1, \infty)\}, \quad (2.8)$$

and

$$\mathbf{l}_l \equiv \{(-\infty, \vartheta_1), (\vartheta_1, \vartheta_2), \dots, (\vartheta_{n-1}, \vartheta_n), (\vartheta_n, \infty)\}, \quad (2.9)$$

such that \mathbf{l}_u and \mathbf{l}_l each comprise $n + 1$ open non-overlapping intervals on \mathbb{R} .

Recall that the inequalities in $r_u(v)$ and $r_l(v)$ depend only on where v_2 lies relative to the endpoints of these intervals. Consider a set of values of v , say \mathcal{V} , comprising a pair $(1, v_2)$ for each $v_2 \in \mathcal{I}_u$ and a pair $(-1, v_2)$ for each $v_2 \in \mathcal{I}_l$.⁹ Then this set \mathcal{V} will feature one representative from each member of the partition \mathbf{V}_u . Moreover, because the elements of the partition \mathbf{V}_l differ from those of \mathbf{V}_u only up their boundaries, such a set \mathcal{V} also has one representative from each member of \mathbf{V}_l . Thus, using any such \mathcal{V} as the sets \mathcal{V}_u and \mathcal{V}_l specified in (2.5) will yield a moment inequality characterization of the finite sample identified set \mathcal{B}_n^* , i.e. with no loss of identifying power.

Such a construction for \mathcal{V} is quite useful in cases where X_i has two components. A set \mathcal{V} can be constructed simply by ordering the observed values of $Z_i (= X_{i1}/X_{i2})$, selecting one element from each of the intervals in (2.8) and (2.9), and pairing them with ± 1 accordingly. Computation of the test statistic proposed in the following section will require taking the supremum of a function of v over the set \mathcal{V} , which turns out to be easy by computing \mathcal{V} and taking the maximum by brute force in the case of two dimensional X_i . When X_i has more elements, the resulting hyperplane arrangements are not as straightforward to characterize. Nonetheless, it appears that advances from the computational geometry literature can be used to achieve computational tractability in such cases.

⁸Here it is to be understood that when $X_{i2} = 0$, Z_i is defined to be $\pm\infty$ according to the sign of X_{i1} if $X_{i1} \neq 0$ and $Z_i = 0$ if $X_{i1} = 0$.

⁹More precisely, the condition for \mathcal{V} is (i) for every interval $I_u \in \mathcal{I}_u$, the intersection $\mathcal{V} \cap (\{1\} \times I_u)$ is nonempty, and (ii) for every interval $I_l \in \mathcal{I}_l$, the intersection $\mathcal{V} \cap (\{-1\} \times I_l)$ is nonempty.

3 Inference Based on Moment Inequalities

For a given value $b \in \mathcal{B}$, we consider the hypothesis test

$$H_0 : \beta = b \quad \text{versus} \quad H_1 : \beta \neq b, \quad (3.1)$$

on the basis of n observations $\{(Y_i, X_i) : i = 1, \dots, n\}$ following the restrictions of the semiparametric binary response model given by Assumption 1. If one wishes to construct a confidence set for β , the set of b for which H_0 is not rejected by a size α test will provide a confidence set guaranteed to contain β with probability at least $1 - \alpha$. As noted in the introduction, our method does not require point identification of β , and thus we do not assume sufficient conditions for point identification. Most notably, the existence of a continuous covariate – much less one with full support on \mathbb{R} – is not required.

To perform inference based on moment inequalities in Theorem 2, we incorporate sample analogs of the moments appearing in (2.5) and (2.6), which are

$$\begin{aligned} \hat{m}_u(b, v) &\equiv \mathbb{E}_n [(2Y - 1)1\{Xb \geq 0, Xv < 0\}], \quad v \in \mathcal{V}_u, \\ \hat{m}_l(b, v) &\equiv \mathbb{E}_n [(1 - 2Y)1\{Xb \leq 0, Xv > 0\}], \quad v \in \mathcal{V}_l, \end{aligned}$$

into our test statistic

$$T_n(b) \equiv \max\{0, \hat{T}_u(b, \mathcal{V}_u), \hat{T}_l(b, \mathcal{V}_l)\}, \quad (3.2)$$

where

$$\hat{\sigma}_u^2(b, v) \equiv \mathbb{E}_n [1\{Xb \geq 0 > Xv\}] - \hat{m}_u(b, v)^2, \quad (3.3)$$

$$\hat{\sigma}_l^2(b, v) \equiv \mathbb{E}_n [1\{Xb \leq 0 < Xv\}] - \hat{m}_l(b, v)^2, \quad (3.4)$$

$$\hat{T}_c(b, \mathcal{V}) \equiv \sup_{v \in \mathcal{V}} \sqrt{n} \frac{-\hat{m}_c(b, v)}{\max\{\hat{\sigma}_c(b, v), \epsilon\}}, \quad c \in \{u, l\}, \quad (3.5)$$

and ϵ is an arbitrarily small positive number taken to ensure a non-zero denominator.¹⁰

Instead of deriving the finite sample distribution of $T_n(b)$ under H_0 , which is unknown, we construct a random variable $T_n^*(b)$ which has a known finite sample distribution given \mathcal{X}_n and which satisfies

$$T_n(b) \leq T_n^*(b) \quad \text{under } H_0 : \beta = b. \quad (3.6)$$

To this purpose define Y_1^*, \dots, Y_n^* by

$$Y_i^* = 1\{U_i \geq 0\}, \quad i = 1, \dots, n,$$

and define $T_n^*(b) \equiv \max\{0, T_u^*(b, \mathcal{V}_u), T_l^*(b, \mathcal{V}_l)\}$ analogously to (3.3) – (3.4) but with

$$\begin{aligned} \hat{m}_u^*(b, v) &\equiv \mathbb{E}_n [(2Y^* - 1)1\{Xb \geq 0, Xv < 0\}] \\ \hat{m}_l^*(b, v) &\equiv \mathbb{E}_n [(1 - 2Y^*)1\{Xb \leq 0, Xv > 0\}], \end{aligned}$$

¹⁰In cases where either $\hat{\sigma}_u^2(b, v)$ or $\hat{\sigma}_l^2(b, v)$ is zero, the ratio in the definition of $\hat{T}_c(b, \mathcal{V})$ can simply be set to $\pm\infty$ according to the sign of the numerator, or zero if the numerator is zero. The ratio that results from the use of arbitrarily small ϵ here simply serves as a placeholder for $\pm\infty$ in computations without any substantive effect.

replacing $\hat{m}_u(b, v)$ and $\hat{m}_l(b, v)$, respectively. The random variable $T_n^*(b)$ itself is not observed because Y_1^*, \dots, Y_n^* are not observed, but the finite sample distribution of $T_n^*(b)$ given \mathcal{X}_n is known since $(2Y_1^* - 1, \dots, 2Y_n^* - 1)$ are independent Rademacher random variables conditional on \mathcal{X}_n .

Thus, for a given level $\alpha \in (0, 1)$, the critical value used for our test is the conditional $1 - \alpha$ quantile of $T_n^*(b)$ given \mathcal{X}_n , namely

$$q_{1-\alpha} \equiv \inf\{c \in \mathbb{R} : \mathbb{P}(T_n^*(b) \leq c \mid \mathcal{X}_n) \geq 1 - \alpha\}.$$

This critical value can be computed up to arbitrary accuracy by drawing a large number of simulations, each of which comprises a sequence of n independent Rademacher random variables.

The relationship between $T_n(b)$ and $T_n^*(b)$ in (3.6) implies Theorem 3, establishing finite sample size control of the proposed test. As is the case with all formal mathematical results stated in the paper, the proofs of inequality (3.6) and Theorem 3 are in the Appendix.

Theorem 3. *Let Assumption 1 hold. Under the null $H_0 : \beta = b$, $\mathbb{P}(T_n(b) \leq q_{1-\alpha} \mid \mathcal{X}_n) \geq 1 - \alpha$.*

Theorem 3 establishes finite sample size control of the rejection rule $1\{T_n(b) > q_{1-\alpha}\}$ for hypothesis test (3.1). While it is possible that $\mathbb{P}(T_n(b) \leq q_{1-\alpha} \mid \mathcal{X}_n)$ strictly exceeds $1 - \alpha$, the following theorem shows that a test with a smaller critical value cannot achieve size control if the critical value is a deterministic function of \mathcal{X}_n . It should however be noted that Theorem 4 is silent with regard to critical values that are a function of both X_1, \dots, X_n and Y_1, \dots, Y_n .

Theorem 4. *Let Assumption 1 hold. The proposed critical value is not conservative in the sense that, for every function cv of \mathcal{X}_n with $cv < q_{1-\alpha}$, there is a distribution of (U_1, \dots, U_n) given \mathcal{X}_n under which $\mathbb{P}(T_n(b) \leq cv \mid \mathcal{X}_n) < 1 - \alpha$ under the null $H_0 : \beta = b$.*

Theorem 4 implies that the use of a more stringent critical value that is a function of \mathcal{X}_n for testing using $T_n(b)$ is not possible without losing size control.

We can establish a power result for our test as a function of a measure of the violation of moment inequalities that define the finite sample identified set. Specifically, Hoeffding's inequality is used to establish a lower bound on finite sample power for certain violations of the inequalities (2.5) and (2.6) from Theorem 2. The result is given in the following theorem.

Theorem 5. *Let Assumptions 1 and 2 hold, and let ρ be any number in $(0, 1)$. If there is $v \in \mathcal{V}_u$ such that*

$$\begin{aligned} & \mathbb{E}[\mathbb{E}_n[(2Y - 1)1\{Xb \geq 0, Xv < 0\}] \mid \mathcal{X}_n] \\ & \leq -\frac{1}{\sqrt{n}} \left(q_{1-\alpha} \max \left\{ \epsilon, \sqrt{\frac{\mathbb{E}_n[1\{Xb \geq 0, Xv < 0\}]}{1 + q_{1-\alpha}^2/n}} \right\} + \sqrt{2 \log(1/\rho) \mathbb{E}_n[1\{Xb \geq 0, Xv < 0\}]} \right), \end{aligned} \quad (3.7)$$

or there is $v \in \mathcal{V}_l$ such that

$$\begin{aligned} & \mathbb{E}[\mathbb{E}_n[(1 - 2Y)1\{Xb \leq 0, Xv > 0\}] \mid \mathcal{X}_n] \\ & \leq -\frac{1}{\sqrt{n}} \left(q_{1-\alpha} \max \left\{ \epsilon, \sqrt{\frac{\mathbb{E}_n[1\{Xb \leq 0, Xv > 0\}]}{1 + q_{1-\alpha}^2/n}} \right\} + \sqrt{2 \log(1/\rho) \mathbb{E}_n[1\{Xb \leq 0, Xv > 0\}]} \right), \end{aligned} \quad (3.8)$$

then the rejection probability is at least $1 - \rho$, i.e., $\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n) \geq 1 - \rho$.

Furthermore, inversion of the conditions in Theorem 5 provides a power guarantee for any parameter value $b \notin \mathcal{B}_n^*$.

Corollary 1. *Let Assumptions 1 and 2 hold. For any $b \notin \mathcal{B}_n^*$, the rejection probability $\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n)$ is at least the maximum of the following two expressions:*

$$\max_{v \in \mathcal{V}_u} \left(1 - \exp \left(-\frac{1}{2} \left(\max \left\{ 0, \sqrt{n} \tilde{\zeta}_u(b, v) - q_{1-\alpha} \max \left\{ \tilde{\epsilon}_u(b, v), (1 + q_{1-\alpha}^2/n)^{-1/2} \right\} \right\} \right)^2 \right) \right), \quad (3.9)$$

$$\max_{v \in \mathcal{V}_l} \left(1 - \exp \left(-\frac{1}{2} \left(\max \left\{ 0, \sqrt{n} \tilde{\zeta}_l(b, v) - q_{1-\alpha} \max \left\{ \tilde{\epsilon}_l(b, v), (1 + q_{1-\alpha}^2/n)^{-1/2} \right\} \right\} \right)^2 \right) \right), \quad (3.10)$$

where the quantities in the above expressions are defined as

$$\tilde{\zeta}_u(b, v) \equiv \frac{-\mathbb{E} [\mathbb{E}_n [(2Y - 1)1\{Xb \geq 0, Xv < 0\}] \mid \mathcal{X}_n]}{\sqrt{\mathbb{E}_n [1\{Xb \geq 0, Xv < 0\}]}}$$

$$\tilde{\zeta}_l(b, v) \equiv \frac{-\mathbb{E} [\mathbb{E}_n [(1 - 2Y)1\{Xb \leq 0, Xv > 0\}] \mid \mathcal{X}_n]}{\sqrt{\mathbb{E}_n [1\{Xb \leq 0, Xv > 0\}]}}$$

$$\tilde{\epsilon}_u(b, v) \equiv \frac{\epsilon}{\sqrt{\mathbb{E}_n [1\{Xb \geq 0, Xv < 0\}]}}$$

$$\tilde{\epsilon}_l(b, v) \equiv \frac{\epsilon}{\sqrt{\mathbb{E}_n [1\{Xb \leq 0, Xv > 0\}]}}.$$

The power bound provided by Theorem 5 depends on the degree to which the inequalities (2.5) and (2.6) that characterize \mathcal{B}_n^* are violated relative to $\sqrt{\mathbb{E}_n [1\{Xb \geq 0, Xv < 0\}]}$ and $\sqrt{\mathbb{E}_n [1\{Xb \leq 0, Xv > 0\}]}$. However, Theorem 5 further implies an explicit mapping between (i) the extent to which a given parameter vector b violates the inequalities that define the finite sample identified set and (ii) a lower bound on the finite sample power of our test for $\beta = b$ that does not depend on sample quantities. To see this, define

$$Q(b) \equiv \max\{Q_u(b), Q_l(b)\},$$

where

$$Q_u(b) \equiv -\min\{0, \min_{v \in \mathcal{V}_u} \mathbb{E} [\mathbb{E}_n [(2Y - 1)1\{Xv < 0 \leq Xb\}] \mid \mathcal{X}_n]\},$$

$$Q_l(b) \equiv -\min\{0, \min_{v \in \mathcal{V}_l} \mathbb{E} [\mathbb{E}_n [(1 - 2Y)1\{Xb \leq 0 < Xv\}] \mid \mathcal{X}_n]\}.$$

The values of $Q_u(b)$ and $Q_l(b)$ denote the maximal violation exhibited by b of the inequalities (2.5) and (2.6) that characterize \mathcal{B}_n^* in Theorem 2. Theorem 5 implies that for $\gamma = 1 - \rho$, our test is guaranteed to have power at least γ whenever the measure of violation $Q(b)$ is at least $C(\gamma)$, defined by

$$C(\gamma) \equiv \frac{1}{\sqrt{n}} \left(q_{1-\alpha} \max \left\{ \epsilon, (1 + q_{1-\alpha}^2/n)^{-1/2} \right\} + \sqrt{-2 \log(1 - \gamma)} \right). \quad (3.11)$$

Inversion of this relation also provides an explicit power guarantee as a function of $Q(b)$. The following corollary to Theorem 5 gives the formal results.

Corollary 2. *Let Assumptions 1 and 2 hold and let $\gamma \in (0, 1)$.*

1. *If $Q(b) \geq C(\gamma)$, then $\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n) \geq \gamma$.*

2. For any $b \notin \mathcal{B}_n^*$, the rejection probability $\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n)$ is at least

$$1 - \exp \left(-\frac{1}{2} \left(\max \left\{ 0, \sqrt{n}Q(b) - q_{1-\alpha} \max \left\{ \epsilon, (1 + q_{1-\alpha}^2/n)^{-1/2} \right\} \right\} \right)^2 \right).$$

In particular, Corollary 2 can be used to indicate how big $Q(b)$ must be in order for Theorem 5 to guarantee our test has power at least α against a parameter value \tilde{b} irrespective of sample quantities. As we show in the next section, the likelihood ratio test of $\beta = b$ against $\beta = \tilde{b}$ can only achieve power α against any \tilde{b} .

Finally, it should be noted that the power bounds delivered by Theorem 5 provide power guarantees, but may not be sharp. That is, the test may achieve higher power than the bounds guarantee.

4 Likelihood Ratio Tests

In this section, we consider the use of likelihood ratio tests under Assumptions 1 and 2. The analysis is divided according to the type of null and alternative hypotheses considered, the key distinctions being whether each hypothesis is simple or composite. When considering the likelihood ratio test, the relevant consideration is the set of conditional distributions for Y_1, \dots, Y_n given \mathcal{X}_n allowed under each hypothesis. A point null hypothesis for β is a composite hypothesis because it does not specify the distribution of U_1, \dots, U_n conditional on \mathcal{X}_n , and therefore admits many different possible conditional distributions for Y_1, \dots, Y_n even under the maintained assumptions that $\mathbb{P}(U_i \geq 0 \mid \mathcal{X}_n) = 1/2$ and that $\mathbb{P}(U_1 \leq t_1, \dots, U_n \leq t_n \mid \mathcal{X}_n) = \prod_{i=1}^n \mathbb{P}(U_i \leq t_i \mid \mathcal{X}_n)$ for any real numbers t_1, \dots, t_n .

The likelihood ratio test is most powerful as long as the alternative hypothesis is simple. However, as shown below, when the alternative hypothesis is composite, the likelihood ratio test based on the least favorable pair has only trivial power (i.e. its power is equal to its size, α) because the least favorable null and alternative distributions coincide. Unlike the likelihood ratio test, our proposed test has power greater than its size α against alternatives \tilde{b} that are sufficiently in violation of the moment inequalities that characterize the finite sample identified set, as can be seen from Theorem 5.

Let $G_{U|\mathcal{X}_n}(\cdot)$ denote the conditional distribution of U_1, \dots, U_n given \mathcal{X}_n . Under Assumptions 1 and 2, the likelihood of $(Y_1, \dots, Y_n) = (y_1, \dots, y_n)$ conditional on \mathcal{X}_n is defined by the product form

$$P(y_1, \dots, y_n; \beta, G_{U|\mathcal{X}_n}) \equiv \prod_{i=1}^n (y_i \cdot G_i([-X_i\beta, \infty)) + (1 - y_i) \cdot G_i((-\infty, -X_i\beta))),$$

where G_i denotes the marginal distribution of U_i conditional on \mathcal{X}_n induced by $G_{U|\mathcal{X}_n}$, i.e., $G_i(\mathcal{T}) \equiv G_{U|\mathcal{X}_n}(\{s \in \mathbb{R}^n : s_i \in \mathcal{T}\})$ for every measurable set $\mathcal{T} \subset \mathbb{R}$.

Below we begin by considering the simplest case, in which the null and alternative hypothesis both specify each of β and $G_{U|\mathcal{X}_n}$, and hence uniquely determine the likelihood $P(y_1, \dots, y_n; \beta, G_{U|\mathcal{X}_n})$, before considering composite hypotheses.

Simple Null Hypotheses and Simple Alternative Hypotheses

Consider the following test of two simple hypotheses.

$$H_0 : (\beta, G_{U|\mathcal{X}_n}) = (b, G) \quad \text{versus} \quad H_1 : (\beta, G_{U|\mathcal{X}_n}) = (\tilde{b}, \tilde{G}). \quad (4.1)$$

The likelihood ratio test is the test that rejects H_0 in favor of H_1 with probability $\phi(y_1, \dots, y_n)$ defined by

$$\phi(y_1, \dots, y_n) = \begin{cases} 1 & \text{if } P(y_1, \dots, y_n; \tilde{b}, \tilde{G}) > kP(y_1, \dots, y_n; b, G) \\ \xi & \text{if } P(y_1, \dots, y_n; \tilde{b}, \tilde{G}) = kP(y_1, \dots, y_n; b, G) \\ 0 & \text{otherwise,} \end{cases}$$

where k and $\xi \in [0, 1]$ are chosen such that

$$\sum_{(y_1, \dots, y_n)} \phi(y_1, \dots, y_n) P(y_1, \dots, y_n; b, G) = \alpha, \quad (4.2)$$

where the summation is to be understood as taken over all $(y_1, \dots, y_n) \in \{0, 1\}^n$.¹¹ Defining

$$p_i \equiv G_i([-X_i b, \infty) \mid \mathcal{X}_n), \quad (4.5)$$

and

$$\tilde{p}_i \equiv \tilde{G}_i([-X_i \tilde{b}, \infty) \mid \mathcal{X}_n), \quad (4.6)$$

The test $\phi(Y_1, \dots, Y_n)$ can then be written as

$$\phi(y_1, \dots, y_n) = \begin{cases} 1 & \text{if } \prod_{i=1}^n \tilde{p}_i^{y_i} (1 - \tilde{p}_i)^{1-y_i} > k \prod_{i=1}^n p_i^{y_i} (1 - p_i)^{1-y_i}, \\ \xi & \text{if } \prod_{i=1}^n \tilde{p}_i^{y_i} (1 - \tilde{p}_i)^{1-y_i} = k \prod_{i=1}^n p_i^{y_i} (1 - p_i)^{1-y_i}, \\ 0 & \text{otherwise,} \end{cases}$$

where k and $\xi \in [0, 1]$ are chosen such that

$$\sum_{(y_1, \dots, y_n) \in \{0, 1\}^n} \phi(y_1, \dots, y_n) \prod_{i=1}^n p_i^{y_i} (1 - p_i)^{1-y_i} = \alpha.$$

When the null and alternative distributions of a hypothesis test are simple, then by the Neyman-Pearson Lemma (e.g., Chapter 3.2 of Lehmann and Romano (2005)), the likelihood ratio test is the uniformly most powerful test. However, the hypothesis test of interest, (3.1), differs because neither the null or alternative specify $G_{U \mid \mathcal{X}_n}$.

¹¹This test uses randomization in the event that $P(y_1, \dots, y_n; b', G') = kP(y_1, \dots, y_n; b, G)$. If a non-randomized implementation of the LRT is preferred that can be done in the usual way by making the modification

$$\phi(y_1, \dots, y_n) = \begin{cases} 1 & \text{if } P(y_1, \dots, y_n; b', G') > kP(y_1, \dots, y_n; b, G), \\ 0 & \text{otherwise,} \end{cases} \quad (4.3)$$

with k chosen as small as possible subject to the constraint that

$$\sum_{(y_1, \dots, y_n)} \phi(y_1, \dots, y_n) P(y_1, \dots, y_n; b, G) \leq \alpha. \quad (4.4)$$

The conclusion of Theorem 6 would then be that the LRT is a most powerful non-randomized test of (4.7) subject to achieving size control. Subsequent results could be similarly modified without substantive change of conclusions.

Composite Null Hypothesis and Simple Alternative Hypothesis

In a step closer to the two-sided hypothesis test (3.1), consider first the hypothesis test

$$H_0 : \beta = b \quad \text{versus} \quad H_1 : (\beta, G_{U|\mathcal{X}_n}) = (\tilde{b}, \tilde{G}), \quad (4.7)$$

which comprises the same null hypothesis of (3.1) and the alternative hypothesis of (4.1). This test features a composite null hypothesis and a simple alternative hypothesis for $(\beta, G_{U|\mathcal{X}_n})$.

Using Theorem 3.8.1 of Lehmann and Romano (2005), we can reduce H_0 to a simple hypothesis by finding the least favorable distribution of (Y_1, \dots, Y_n) given \mathcal{X}_n . This entails pairing the null value $\beta = b$ with the distribution $G_{U|\mathcal{X}_n}$ among those satisfying the restrictions maintained in Assumptions 1 and 2 for which the likelihood ratio test has minimal power against H_1 . Using the independence restriction of Assumption 2, the admitted distributions of (Y_1, \dots, Y_n) given \mathcal{X}_n are given by the marginal Bernoulli probabilities for each i , so that the hypotheses can be characterized by the implied collections of probabilities (p_1, \dots, p_n) and $(\tilde{p}_1, \dots, \tilde{p}_n)$ coincident with (4.5) under the null and alternative, respectively. Thus the null hypothesis in (4.7) can be written as

$$H_0 : (\beta, G_{U|\mathcal{X}_n}) = \left(b, \prod_{i=1}^n G_i \right) \text{ for some } (G_1, \dots, G_n).$$

For each observation i consider the probabilities p_i and \tilde{p}_i . Note that under Assumptions 1 and 2, $p_i - 1/2$ is nonpositive (nonnegative) if $X_i b$ is nonpositive (nonnegative), and likewise $\tilde{p}_i - 1/2$ is nonpositive (nonnegative) if $X_i \tilde{b}$ is nonpositive (nonnegative). When $X_i b$ and $\tilde{p}_i - 1/2$ have the same sign, then there exists a distribution of unobservable U_i conditional on \mathcal{X}_n such that p_i can be made equal to \tilde{p}_i . For such i , the least favorable distribution will thus have $p_i = \tilde{p}_i$. When instead $X_i b$ and $\tilde{p}_i - 1/2$ have the opposite sign, then in general p_i cannot be equal to \tilde{p}_i under the null due to the requirement in Assumption 1(iii) that $\mathbb{P}(U_i \geq 0 \mid \mathcal{X}_n) = 1/2$. The probability p_i will however be made as close as possible to \tilde{p}_i while adhering to this assumption if the conditional distribution of U_i allocates all mass to regions in which $|U_i| > |X_i b|$, so that $p_i = 1/2$.

We now formally define \bar{p}_i to denote the null probabilities $p_i = \mathbb{P}\{Y_i = 1 \mid \mathcal{X}_n\}$ in keeping with the above intuition for constructing the least favorable distribution under the composite null $\beta = b$ as follows:

$$\bar{p}_i = \begin{cases} 1/2 & \text{if } X_i b(\tilde{p}_i - 1/2) < 0 \text{ or } (X_i b = 0 \text{ and } \tilde{p}_i < 1/2), \\ \tilde{p}_i & \text{otherwise.} \end{cases} \quad (4.8)$$

These probabilities lie within the null set in (4.7) because $p_i = 1/2$ is possible for any value of b , and $p_i = \tilde{p}_i$ is possible when $\tilde{p}_i - 1/2$ and $X_i b$ are either both positive or negative, as well as when $X_i b = 0$ and $\tilde{p}_i \geq 1/2$. By the Neyman-Pearson Lemma, the likelihood ratio test for the simple null of $\bar{p}_1, \dots, \bar{p}_n$ against $\tilde{p}_1, \dots, \tilde{p}_n$ is most powerful. Theorem 6 verifies that the likelihood ratio test for this simple hypothesis controls size under the composite null $\beta = b$, thus enabling application of Theorem 3.8.1 of Lehmann and Romano (2005) and establishing that these null probabilities are least favorable.

To proceed consider the corresponding test $\bar{\phi}(Y_1, \dots, Y_n)$ for (4.7) that makes use of this configuration

of implied probabilities under the null.

$$\bar{\phi}(y_1, \dots, y_n) = \begin{cases} 1 & \text{if } \prod_{i=1}^n \tilde{p}_i^{y_i} (1 - \tilde{p}_i)^{1-y_i} > k \prod_{i=1}^n \bar{p}_i^{y_i} (1 - \bar{p}_i)^{1-y_i}, \\ \xi & \text{if } \prod_{i=1}^n \tilde{p}_i^{y_i} (1 - \tilde{p}_i)^{1-y_i} = k \prod_{i=1}^n \bar{p}_i^{y_i} (1 - \bar{p}_i)^{1-y_i}, \\ 0 & \text{otherwise,} \end{cases} \quad (4.9)$$

where k and $\xi \in [0, 1]$ are chosen such that

$$\sum_{(y_1, \dots, y_n) \in \{0, 1\}^n} \bar{\phi}(y_1, \dots, y_n) \prod_{i=1}^n \bar{p}_i^{y_i} (1 - \bar{p}_i)^{1-y_i} = \alpha. \quad (4.10)$$

There is the following result, leveraging Theorem 3.8.1 of Lehmann and Romano (2005).

Theorem 6. *Let Assumptions 1 and 2 hold. Consider the null and alternative hypotheses stated in (4.7) and the test $\bar{\phi}(y_1, \dots, y_n)$ defined in (4.9) and (4.10). If the null hypothesis is true then the rejection probability is no greater than α , that is, $\mathbb{E} [\bar{\phi}(Y_1, \dots, Y_n) \mid \mathcal{X}_n] \leq \alpha$. Moreover,*

$$P(y_1, \dots, y_n; b, G_{U \mid \mathcal{X}_n}) = \prod_{i=1}^n \bar{p}_i^{y_i} (1 - \bar{p}_i)^{1-y_i} \quad (4.11)$$

is the least favorable distribution of (Y_1, \dots, Y_n) given \mathcal{X}_n under H_0 against H_1 , and the test $\bar{\phi}(Y_1, \dots, Y_n)$ is a most powerful test of H_0 against H_1 .

Composite Null Hypothesis and Composite Alternative Hypothesis

When the researcher's goal is to test $\beta = b$ against the simple alternative hypothesis in (4.7) that completely specifies the conditional distribution of Y_1, \dots, Y_n given \mathcal{X}_n , then the likelihood ratio test implemented by adopting the rejection probability specified in (4.8) – (4.10) is most powerful. When instead the researcher wishes to control power against a composite alternative, such as that of

$$H_0 : \beta = b \quad \text{versus} \quad H_1 : \beta = \tilde{b}, \quad (4.12)$$

Theorem 6 is silent because each conditional distribution of Y_1, \dots, Y_n allowed under the alternative hypothesis will result in a different likelihood ratio test. However, following arguments in Chapter 8.1 of Lehmann and Romano (2005), it is straightforward to construct the least favorable pair of distributions for this test. Note that for *any* \tilde{b} hypothesized under the alternative, distributions \tilde{G}_i can be specified such that for all i :

$$\tilde{p}_i \equiv \tilde{G}_i([-X_i \tilde{b}, \infty) \mid \mathcal{X}_n) = 1/2. \quad (4.13)$$

In particular, this is achieved by distributions \tilde{G}_i that allocate all mass to regions on which $|U_i| \geq |X_i \tilde{b}|$, while obeying the constraint $\mathbb{P}(U_i \geq 0 \mid \mathcal{X}_n) = 1/2$.¹² Such a combination of (\tilde{b}, \tilde{G}) under the alternative yields $\tilde{p}_i = 1/2$ for all i . This conclusion holds irrespective of the hypothesized value of \tilde{b} in (4.7). Indeed it also holds for the value of $\beta = b$ hypothesized under the null. Correspondingly the least favorable $\{\bar{p}_i : i = 1, \dots, n\}$ under the null given by (4.8) when all $\tilde{p}_i = 1/2$ is given by $\bar{p}_i = 1/2$ for all i . Thus from Theorem 6 we have

¹²If the distribution of U_i were restricted to have positive density on \mathbb{R} , then G_i could be specified to allocate probability $1 - \epsilon$ to $\{u_i : |u_i| \geq |X_i b|\}$ for any small $\epsilon > 0$.

the following implications for the two-sided test of $\beta = b$ against $\beta = \tilde{b}$.

Corollary 3. *Let Assumptions 1 and 2 hold. The least favorable pair of $\{\bar{p}_i : i = 1, \dots, n\}$ and $\{\tilde{p}_i : i = 1, \dots, n\}$ for the test (4.7) is given by $\bar{p}_i = 1/2$ for all i and $\tilde{p}_i = 1/2$ for all i . Moreover, since the conclusion holds for any b and \tilde{b} , this is also the least favorable pair for the two-sided hypothesis test of $\beta = b$ in (3.1).*

This corollary is a direct implication of Theorem 6. While it may be conceptually appealing to construct the likelihood ratio test for the two-sided test (4.12) based on the least favorable pair, we see from (4.9) and (4.10) that this results in a test which rejects the null hypothesis with probability α irrespective of the data. This is because the hypothesis $\beta = b$ always includes the distribution under which $p_i = 1/2$ for all i , for any hypothesized value of b . Thus the sets of feasible conditional distributions for (Y_1, \dots, Y_n) given \mathcal{X}_n compatible with each of the two hypotheses, $\beta = b$ and $\beta = \tilde{b}$ overlap. Therefore, in the language of Chesher and Rosen (2017) and Kaido and Zhang (2019), even if \tilde{b} lies outside the finite sample identified set, parameter values b and \tilde{b} are potentially observationally equivalent.

Against certain alternatives, the likelihood ratio test using the least favorable pair can achieve higher power than the inequality test using $T_n(b)$, but the reverse is also true, so that neither test dominates. Unlike the likelihood ratio test for the composite hypothesis 3.1, the inequality test we propose in this paper based on $T_n(b)$ does not simply reject with probability α irrespective of the data. Moreover, Theorem 5 and Corollaries 2 and 3 provide lower power envelopes as a function of the violation of the inequalities characterizing the finite sample identified set. In particular, these results can be used to characterize values of \tilde{b} against which the inequality test achieves nontrivial power by setting $\gamma > \alpha$. For such alternatives, the inequality test achieves a higher power than the likelihood ratio test.

5 Monte Carlo Experiments

In this section, we present Monte Carlo results illustrating the relative performance of our test compared to that of a test using the smoothed maximum score estimator from Horowitz (1992) with a bootstrap critical value following Horowitz (2002), and a likelihood ratio test against a simple alternative hypothesis.¹³ We consider examples in which there are two covariates. In each design, the first component of β is normalized to one, and we report the results of conducting tests of the null hypothesis $\beta = (1, \theta)'$ against the alternative hypothesis $\beta \neq (1, \theta)'$ for θ ranging from $\theta - 3$ to $\theta + 3$. In the population design for the actual data generation process, the true parameter value is $\beta = (1, 1)'$.

For the likelihood ratio test, the alternative hypothesis consists of the simple hypothesis corresponding to the true data distribution. This is known in the Monte Carlo simulation, but would be unknown in practice. This likelihood ratio test is therefore the (infeasible) optimal test of the hypothesis $\beta = b$ against $\beta = (1, 1)'$ paired with $G_{U|\mathcal{X}_n}$, producing Bernoulli probabilities for $\tilde{p}_1, \dots, \tilde{p}_n$ for the actual DGP employed in the Monte Carlo simulations. This is thus a test of the composite null $\beta = b$ against a point alternative as described by (4.7).

In the Monte Carlo experiments, we use one random draw of the n -tuple (X_1, \dots, X_n) , and 500 independent draws of the n -tuple (Y_1, \dots, Y_n) , where the sample size is $n = 100$. In this simulation design the finite sample identified set is thus fixed across experiments.

To implement the three inference methods, we use the following tuning parameters. The variable ϵ was set to MATLAB's `eps` value of approximately $2.2 \cdot 10^{-16}$. To compute the critical value, we use 500 random

¹³For the likelihood ratio test, as with our own test, we use a non-randomized test, so that the size of the test is as close as possible to α without exceeding it, as described in Footnote 11.

draws of n Rademacher random variables and 500 samples for the bootstrap procedure described in Horowitz (2002). In all cases we considered tests with size $\alpha = 0.10$.

Following Horowitz (1992, Section 3), $X = (X_1, X_2)$ is a bivariate normal random vector with $\mathbb{E}[X_1] = 0$, $\mathbb{E}[X_2] = 1$, $Var(X_1) = Var(X_2) = 1$ and $Cov(X_1, X_2) = 0$. The distribution of unobservable U in each design is as follows, with U independent of X in designs 1-3 and V independent of X in design 4.

- Design 1: U is distributed according to the logistic distribution with mean zero and variance one.
- Design 2: U is uniformly distributed on $[-\sqrt{12}/2, \sqrt{12}/2]$.
- Design 3: U is distributed according to the Student's t distribution normalized to have variance one.
- Design 4: $U = 0.25(1 + 2Z^2 + Z^4)V$ where $Z = X_1 + X_2$ and V is distributed according to the logistic distribution with mean zero and variance one.

Figure 2 presents non-rejection frequencies for θ ranging from $\theta - 3$ to $\theta + 3$ for $\alpha = 0.10$. For the non-rejection frequencies, we conduct inference using the smoothed maximum score bootstrap implementation described in Horowitz (2002) (Horowitz-Bootstrap), our proposed method (Rosen-Ura), and the (infeasible) likelihood ratio test (LRT). For reference, the finite sample identified set (FSID) is also illustrated with height of $1 - \alpha$ on the vertical axis by a dashed blue line. The qualitative comparison among the three inference methods is similar for all four designs. The non-rejection frequency for our method is maximized around the population value of $\theta = 1$ and exceeds 0.90 in a range spanning from about 0.6 to 1.7. Our method is always less powerful than the infeasible likelihood ratio test, which results from the fact that the alternative hypothesis for our method is the composite hypothesis $\beta \neq b$, while the likelihood ratio test is based on using the actual population distribution as the simple alternative hypothesis. This is what makes the likelihood ratio test considered here infeasible; in an application, the population distribution of Y_1, \dots, Y_n conditional on \mathcal{X}_n is unknown.

Compared to inference using the bootstrap implementation of the smoothed maximum score estimator, we find an asymmetry in the simulation results. For values smaller than the true parameter value, our method is more powerful than the smoothed maximum score bootstrap procedure, although the performance is almost the same in Design 4. For values larger than the true parameter value, our method is less powerful than the smoothed maximum score bootstrap procedure, but the smoothed maximum score bootstrap procedure is more powerful than the optimal test (the likelihood ratio test) for some values larger than the true parameter value. This is likely because the formal asymptotic theory for the smoothed maximum score estimator in Horowitz (1992) and its use for inference by way of the bootstrap implementation described in Horowitz (2002) invoke stronger assumptions than those of Assumptions 1 and 2 in this paper, and these additional assumptions hold in the data generating processes used in these illustrations. Further, we see in Designs 1 and 2 that inference using the smoothed maximum score bootstrap results in rejection frequencies in excess of $\alpha = 0.1$ for certain parameter values that lie inside the finite sample identified set. This is not surprising because its performance is guaranteed to achieve size control asymptotically under conditions whereby the identified set is in fact a singleton.

In Appendix B, we provide additional illustrations of simulation results for a different population parameter value $\theta_0 = 2$ and/or a larger sample size $n = 250$. The results are qualitatively similar to those shown in Figure 2.

6 Conclusion

In this paper we have proposed an approach to conduct finite sample inference on the parameters of Manski's (1985) semiparametric binary response model, for which the maximum score estimator has been shown to be cube-root consistent with a non-normal asymptotic distribution when there is point identification. Our finite sample inference approach circumvents the need to accommodate the complicated asymptotic behavior of this point estimator. Since our goal was finite sample inference, we considered the problem of making inference conditional on the n covariate vectors observable in a finite sample. With covariates taking on only a finite number of observed values, the parameter vector β is not point identified. We therefore employed moment inequality implications for β for the sake of constructing our test statistic for inference, as the moment inequalities are valid no matter whether β is point identified. In order to exposit what observable implications can be distilled on only the basis of exogenous variables observed in the finite sample, we defined the notion of a finite sample identified set. We showed how to make use of the full set of observable implications conditional on the size n sequence of exogenous variables in our construction of a test statistic $T_n(b)$. Finite sample valid critical values were established, and were shown to be easily computed by making use of many simulations of size n sequences of independent Rademacher variables. A finite sample power (lower) bound was also presented and the results of some Monte Carlo experiments were reported, illustrating the performance of the test.

Several interesting directions for future research are possible. The maximum score estimator is one of several estimators in the econometrics literature that consistently estimate a model parameter that may only be identified under support conditions that can never be satisfied by an empirical distribution based on a finite sample. Some such estimators, like the maximum score estimator, exhibit slower than $n^{-1/2}$ convergence rates. Other such estimators, such as the maximum rank correlation estimator of Han (1987) achieve the parametric rate. In this paper we have exploited the particular structure of the semi-parametric binary response model, but there may nonetheless be potential to extend ideas in this paper to such settings to alleviate dependence on conditions not satisfied by empirical distributions that result from finite data.

Other possible avenues of investigation pertain to optimal testing. One direction could be to exploit the likelihood ratio test analysis in this paper to consider minimax testing rates under additional assumptions on the distribution of U_i .¹⁴ Minimax optimal estimation has recently been investigated in a setting in which covariates have high-dimension by Mukherjee, Banerjee, and Ritov (2019) in an asymptotic framework under sufficient conditions for point identification. Investigation of minimax optimal estimators and tests could be interesting to consider when conditions for point identification are not guaranteed to hold. More generally, in future work we aim to continue to explore the interplay between partial identification and testability, and in particular the implications of not having point identification based on an underlying discrete data distribution, as one always has when using the empirical distribution obtained in a finite data set.

¹⁴This line of research was suggested by Tim Armstrong.




Figure 2.a: Design 1.




Figure 2.b: Design 2.




Figure 2.c: Design 3.




Figure 2.d: Design 4.

Figure 2: Non-rejection frequencies with $1 - \alpha = 90\%$ with true $\theta_0 = 1$ and $n = 100$.

References

- ABREVAYA, J., AND J. HUANG (2005): “On the Bootstrap of the Maximum Score Estimator,” *Econometrica*, 73(4), 1175–1204.
- ANDREWS, D. W. K., AND X. SHI (2013): “Inference Based on Conditional Moment Inequalities,” *Econometrica*, 81(2), 609–666.
- ARMSTRONG, T., AND M. KOLESÁR (2018): “Finite-Sample Optimal Estimation and Inference on Average Treatment Effects Under Unconfoundedness,” arXiv:1712.04594v2.
- AVIS, D., AND K. FUKUDA (1996): “Reverse Search for Enumeration,” *Discrete Applied Mathematics*, 65(1-3), 21–46.
- BLEVINS, J. R. (2015): “Non-standard Rates of Convergence of Criterion-Function-Based Set Estimators for Binary Response Models,” *The Econometrics Journal*, 18(2), 172–199.
- CATTANEO, M. D., M. JANSSON, AND K. NAGASAWA (2019): “Bootstrap-Based Inference for Cube Root Consistent Estimators,” Working paper, arXiv:1704.08066.
- CHEN, L.-Y., AND S. LEE (2018): “Best Subset Binary Prediction,” *Journal of Econometrics*, 206(1), 39 – 56.
- (2019): “Breaking the Curse of Dimensionality in Conditional Moment Inequalities for Discrete Choice Models,” *Journal of Econometrics*, 210(2), 482–497.
- CHERNOZHUKOV, V., C. HANSEN, AND M. JANSSON (2009): “Finite Sample Inference for Quantile Regression Models,” *Journal of Econometrics*, 152(2), 93 – 103.
- CHERNOZHUKOV, V., S. LEE, AND A. M. ROSEN (2013): “Intersection Bounds: Estimation and Inference,” *Econometrica*, 81(2), 667–737.
- CHESHER, A., AND A. M. ROSEN (2017): “Generalized Instrumental Variable Models,” *Econometrica*, 85(3), 959–989.
- CLOPPER, C., AND E. S. PEARSON (1934): “The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial,” *Biometrika*, 26(4), 404–413.
- COVER, T. M. (1965): “Geometrical and Statistical Properties of Systems of Linear Inequalities with Applications in Pattern Recognition,” *IEEE Transactions on Electrical Computers*, EC-14(3), 326–334.
- DELGADO, M. A., J. M. RODRÍGUEZ-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.
- GU, J., AND R. KOENKER (2018): “Nonparametric Maximum Likelihood Methods for Binary Response Models with Random Coefficients,” arXiv:1811.03329v1.
- HAN, A. K. (1987): “Non-parametric Analysis of a Generalized Regression Model: The Maximum Rank Correlation Estimator,” *Journal of Econometrics*, 35, 303–316.
- HOEFFDING, W. (1963): “Probability Inequalities for Sums of Bounded Random Variables,” *Journal of the American Statistical Association*, 58(301), 13–30.
- HOROWITZ, J. L. (1992): “A Smoothed Maximum Score Estimator for the Binary Response Model,” *Econometrica*, 60(3), 505–531.
- (2002): “Bootstrap Critical Values for Tests Based on the Smoothed Maximum Score Estimator,” *Journal of Econometrics*, 141-167(2), 505–531.

- JUN, S. J., J. PINKSE, AND Y. WAN (2015a): “Classical Laplace estimation for n^3 -consistent estimators: Improved convergence rates and rate-adaptive inference,” *Journal of Econometrics*, 187(1), 201–216.
- (2015b): “Integrated Score Estimation,” *Econometric Theory*, 33(6), 1418–1456.
- KAIDO, H., AND Y. ZHANG (2019): “Robust Likelihood Ratio Test for Incomplete Economic Models,” Working paper, arXiv:1910.04610.
- KIM, J., AND D. POLLARD (1990): “Cube Root Asymptotics,” *Annals of Statistics*, 18(1), 191–219.
- KOMAROVA, T. (2013): “Binary Choice Models with Discrete Regressors: Identification and Misspecification,” *Journal of Econometrics*, 177(1), 14 – 33.
- LÉGER, C., AND B. MACGIBBON (2006): “On the Bootstrap in Cube Root Asymptotics,” *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, 34(1), 29–44.
- LEHMANN, E. L., AND J. P. ROMANO (2005): *Testing Statistical Hypotheses, Third Edition*. Springer.
- 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.
- (2007): “Partial Identification of Counterfactual Choice Probabilities,” *International Economic Review*, 48(4), 1393–1410.
- MANSKI, C. F., AND T. S. THOMPSON (1986): “Operational Characteristics of Maximum Score Estimation,” *Journal of Econometrics*, 32(1), 85–108.
- MUKHERJEE, D., M. BANERJEE, AND Y. RITOV (2019): “Non-Standard Asymptotics in High Dimensions: Manski’s Maximum Score Estimator Revisited,” Working paper, arXiv:1903.10063.
- PATRA, R. K., E. SEIJO, AND B. SEN (2018): “A Consistent Bootstrap Procedure for the Maximum Score Estimator,” *Journal of Econometrics*, 205(2), 488–507.
- PINKSE, C. (1993): “On the Computation of Semiparametric Estimates in Limited Dependent Variable Models,” *Journal of Econometrics*, 58(1), 185 – 205.
- POWELL, J. L. (1994): “Estimation of Semiparametric Models,” in *The Handbook of Econometrics*, ed. by R. F. Engle, and D. L. McFadden, vol. 4. North-Holland.
- RADA, M., AND M. ČERNÝ (2018): “A New Algorithm for Enumeration of Cells of Hyperplane Arrangements and a Comparison with Avis and Fukuda’s Reverse Search,” *Siam Journal of Discrete Mathematics*, 32(1), 455 – 473.
- SEO, M. H., AND T. OTSU (2018): “Local M-estimation with discontinuous criterion for dependent and limited observations,” *Ann. Statist.*, 46(1), 344–369.
- SLEUMER, N. (1998): “Output-sensitive Cell Enumeration in Hyperplane Arrangements,” in *Algorithm Theory — SWAT’98*, ed. by S. Arnborg, and L. Ivansson, pp. 300–309. Springer-Verlag Berlin Heidelberg.
- SYRGKANIS, V., E. TAMER, AND J. ZIANI (2018): “Inference on Auctions with Weak Assumptions on Information,” Working paper, arXiv:1710.03830.

A Proofs

Proof of Lemma 1. If $X_i\beta \geq 0$, then $Y_i = 1\{X_i\beta + U_i \geq 0\} \geq 1\{U_i \geq 0\}$ and therefore $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] \geq 2\mathbb{P}(U_i \geq 0 \mid \mathcal{X}_n) - 1 = 0$. If $X_i\beta \leq 0$, then $Y_i = 1\{X_i\beta + U_i \geq 0\} \leq 1\{U_i \geq 0\}$ and therefore $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] \leq 2\mathbb{P}(U_i \geq 0 \mid \mathcal{X}_n) - 1 = 0$. \square

Proof of Theorem 1. Directly from Lemma 1, $b \in \mathcal{B}_n^*$ implies

$$\mathbb{E}[(2Y_i - 1)1\{X_i b \geq 0\} \mid \mathcal{X}_n] \geq 0 \text{ and } \mathbb{E}[(1 - 2Y_i)1\{X_i b \leq 0\} \mid \mathcal{X}_n] \geq 0.$$

To demonstrate the other direction, let b be any element of \mathcal{B} such that $\mathbb{E}[(2Y_i - 1)1\{X_i b \geq 0\} \mid \mathcal{X}_n] \geq 0$ and $\mathbb{E}[(1 - 2Y_i)1\{X_i b \leq 0\} \mid \mathcal{X}_n] \geq 0$ for every $i = 1, \dots, n$. Then

$$X_i b \geq 0 \implies \mathbb{P}(Y_i = 1 \mid \mathcal{X}_n) \geq 1/2, \quad (\text{A.1})$$

$$X_i b \leq 0 \implies \mathbb{P}(Y_i = 1 \mid \mathcal{X}_n) \leq 1/2. \quad (\text{A.2})$$

To show that b is in \mathcal{B}_n^* as defined in Definition 1, we now construct a collection of random variables $\{\tilde{U}_i : i = 1, \dots, n\}$ such that for all $i = 1, \dots, n$: (i) $\mathbb{P}(Y_i = 1\{X_i b + \tilde{U}_i \geq 0\} \mid \mathcal{X}_n) = 1$, and (ii) $\mathbb{P}(\forall i, 1\{U_i \geq 0\} = 1\{\tilde{U}_i \geq 0\} \mid \mathcal{X}_n) = 1$. To do so, let $\kappa_i : i = 1, \dots, n$ be n positive random variables defined on $(\Omega, \mathfrak{F}, \mathbb{P})$ and consider each of the cases $X_i\beta < 0$, $X_i\beta = 0$, and $X_i\beta > 0$ in turn as follows.

Case 1: $X_i\beta < 0$. By Lemma 1,

$$\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] \leq 0. \quad (\text{A.3})$$

Let

$$\tilde{U}_i \equiv \begin{cases} \max\{-X_i b, 0\} + \kappa_i & \text{if } U_i \geq -X_i\beta \\ 0 & \text{if } 0 \leq U_i < -X_i\beta \\ \min\{-X_i b, 0\} - \kappa_i & \text{if } U_i < 0. \end{cases}$$

Then $1\{\tilde{U}_i \geq 0\} = 1\{U_i \geq 0\}$, which verifies (ii). To verify (i), note that:

$$1\{X_i b + \tilde{U}_i \geq 0\} = 1\{X_i\beta + U_i \geq 0\} + 1\{0 \leq U_i < -X_i\beta, X_i b \geq 0\} = Y_i + 1\{0 \leq U_i < -X_i\beta, X_i b \geq 0\},$$

because

$$X_i b + \tilde{U}_i = \begin{cases} \max\{X_i b, 0\} + \kappa_i & \text{if } U_i \geq -X_i\beta \\ X_i b & \text{if } 0 \leq U_i < -X_i\beta \\ \min\{X_i b, 0\} - \kappa_i & \text{if } U_i < 0. \end{cases}$$

Therefore,

$$\begin{aligned} \mathbb{P}(Y_i = 1\{X_i b + \tilde{U}_i \geq 0\} \mid \mathcal{X}_n) &= \mathbb{P}(1\{0 \leq U_i < -X_i\beta, X_i b \geq 0\} = 0 \mid \mathcal{X}_n) \\ &\geq \mathbb{P}(1\{0 \leq U_i < -X_i\beta\} = 0 \mid X_i b \geq 0, \mathcal{X}_n) \\ &= \mathbb{P}(U_i \geq -X_i\beta \mid X_i b \geq 0, \mathcal{X}_n) - \mathbb{P}(U_i \geq 0 \mid X_i b \geq 0, \mathcal{X}_n) \\ &= \mathbb{P}(Y_i = 1 \mid X_i b \geq 0, \mathcal{X}_n) - 1/2. \end{aligned}$$

Since (A.1) and (A.3) imply that $\mathbb{P}(Y_i = 1 \mid X_i b \geq 0, \mathcal{X}_n) = 1/2$, we have $\mathbb{P}(Y_i = 1\{X_i b + \tilde{U}_i \geq 0\} \mid \mathcal{X}_n) = 1$, which verifies (i).

Case 2: $X_i\beta = 0$. Let

$$\tilde{U}_i \equiv \begin{cases} \max\{-X_i b, 0\} + \kappa_i & \text{if } U_i \geq 0 \\ \min\{-X_i b, 0\} - \kappa_i & \text{if } U_i < 0. \end{cases} \quad (\text{A.4})$$

Then $1\{\tilde{U}_i \geq 0\} = 1\{U_i \geq 0\}$, which verifies (ii). It further follows from (A.4) that

$$X_i b + \tilde{U}_i = \begin{cases} \max\{0, X_i b\} + \kappa_i & \text{if } U_i \geq 0 \\ \min\{0, X_i b\} - \kappa_i & \text{if } U_i < 0. \end{cases}$$

Consequently since $X_i\beta = 0$, $X_i b + \tilde{U}_i \geq 0$ if and only if $X_i\beta + U_i \geq 0$, verifying (i).

Case 3: $X_i\beta > 0$. The proof is similar to Case 1 with

$$\tilde{U}_i \equiv \begin{cases} \max\{-X_i b, 0\} + \kappa_i & \text{if } U_i \geq 0 \\ -X_i b & \text{if } -X_i\beta \leq U_i < 0 \\ \min\{-X_i b, 0\} - \kappa_i & \text{if } U_i < -X_i b. \end{cases}$$

□

Proof of Theorem 2. That $b \in \mathcal{B}_n^*$ implies (2.5) and (2.6) is immediate. To demonstrate the other direction, we are going to show that (2.5) and (2.6) imply that

$$\mathbb{E}[(2Y_i - 1)1\{X_i b \geq 0\} \mid \mathcal{X}_n] \geq 0 \text{ and } \mathbb{E}[(1 - 2Y_i)1\{X_i b \leq 0\} \mid \mathcal{X}_n] \geq 0 \quad (\text{A.5})$$

for every $i = 1, \dots, n$. To show the result, let $b \in \mathcal{B}$ such that (2.5) and (2.6) hold. Let $v_u \in \mathcal{V}_u$ and $v_l \in \mathcal{V}_l$ such that $r_u(v_u) = r_u(\beta)$ and $r_l(v_l) = r_l(\beta)$. Note that under Assumption 1 (iii), such v_u and v_l exist. Lemma 1 implies

$$\begin{aligned} -\mathbb{E}_n[\mathbb{E}[2Y - 1 \mid \mathcal{X}_n]1\{Xb \geq 0, X\beta < 0\}] &= \mathbb{E}_n[\mathbb{E}[2Y - 1 \mid \mathcal{X}_n]1\{Xb \geq 0, X\beta < 0\}] \\ &= \mathbb{E}_n[\mathbb{E}[2Y - 1 \mid \mathcal{X}_n]1\{Xb \geq 0, Xv_u < 0\}] \\ &= \mathbb{E}[\mathbb{E}_n[(2Y - 1)1\{Xb \geq 0, Xv_u < 0\}] \mid \mathcal{X}_n] \\ &\geq 0 \end{aligned} \quad (\text{A.6})$$

$$\begin{aligned} -\mathbb{E}_n[\mathbb{E}[2Y - 1 \mid \mathcal{X}_n]1\{Xb \leq 0, X\beta > 0\}] &= \mathbb{E}_n[\mathbb{E}[1 - 2Y \mid \mathcal{X}_n]1\{Xb \leq 0, X\beta > 0\}] \\ &= \mathbb{E}_n[\mathbb{E}[1 - 2Y \mid \mathcal{X}_n]1\{Xb \leq 0, Xv_l > 0\}] \\ &= \mathbb{E}[\mathbb{E}_n[(1 - 2Y)1\{Xb \leq 0, Xv_l > 0\}] \mid \mathcal{X}_n] \\ &\geq 0. \end{aligned} \quad (\text{A.7})$$

Moreover, since both $-\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n]1\{X_i b \geq 0, X_i\beta < 0\}$ and $-\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n]1\{X_i b \leq 0, X_i\beta > 0\}$ must be non-positive for every $i = 1, \dots, n$, we have

$$\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n]1\{X_i b \geq 0, X_i\beta < 0\} = \mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n]1\{X_i b \leq 0, X_i\beta > 0\} = 0 \quad (\text{A.8})$$

for every $i = 1, \dots, n$. We can demonstrate Eq. (A.5) for every $i = 1, \dots, n$ by considering the following three cases: $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] = 0$, $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] > 0$, and $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] < 0$. For every i with $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] = 0$, Eq. (A.5) holds with equality. For every i with $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] > 0$, we have $X_i\beta > 0$ from Lemma 1, and

therefore Eq. (A.8) implies $X_i b > 0$, which in turn implies Eq. (A.5). For every i with $\mathbb{E}[2Y_i - 1 \mid \mathcal{X}_n] < 0$, we have $X_i \beta < 0$ from Lemma 1, and therefore Eq. (A.8) implies $X_i b < 0$, which in turn implies Eq. (A.5). \square

Proof of Theorem 3. If Eq. (3.6) holds under H_0 , then $\mathbb{P}(T_n(\beta) \leq q_{1-\alpha} \mid \mathcal{X}_n) \geq \mathbb{P}(T_n^*(\beta) \leq q_{1-\alpha} \mid \mathcal{X}_n) \geq 1 - \alpha$. For the rest of the proof, we are going to show inequality (3.6) under H_0 . Since

$$\begin{cases} Y_i \geq Y_i^* & \text{if } X_i \beta \geq 0 \\ Y_i \leq Y_i^* & \text{if } X_i \beta \leq 0, \end{cases}$$

for every $i = 1, \dots, n$, we have

$$\mathbb{E}_n[(2Y - 1)1\{X\beta \geq 0, Xv < 0\}] \geq \mathbb{E}_n[(2Y^* - 1)1\{X\beta \geq 0 > Xv\}], \forall v \in \mathcal{V}_u$$

and

$$\mathbb{E}_n[(1 - 2Y)1\{X\beta \leq 0, Xv > 0\}] \geq \mathbb{E}_n[(1 - 2Y^*)1\{X\beta \leq 0 < Xv\}], \forall v \in \mathcal{V}_l.$$

By the construction of $T^*(\beta)$ and $T_n(\beta)$, it suffices to show that, for every $v \in \mathbb{R}^K$, the two functions,

$$t \mapsto \max \left\{ 0, \sqrt{n} \frac{-t}{\max\{\epsilon, \sqrt{\mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t^2}\}} \right\}$$

and

$$t \mapsto \max \left\{ 0, \sqrt{n} \frac{-t}{\max\{\epsilon, \sqrt{\mathbb{E}_n[1\{X\beta \leq 0, Xv > 0\}] - t^2}\}} \right\}$$

are weakly decreasing. For the rest of the proof, we focus on the first function

$$f(t) \equiv \max \left\{ 0, \sqrt{n} \frac{-t}{\max\{\epsilon, \sqrt{\mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t^2}\}} \right\}.$$

Consider t_1 and t_2 with $t_1 < t_2$. If $t_2 \geq 0$, we have $f(t_1) \geq 0 = f(t_2)$. For the rest of the proof, therefore, we are going to show $f(t_1) \geq f(t_2)$ when $t_1 < t_2 < 0$. Since $t_2^2 < t_1^2$, we have

$$\mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t_1^2 < \mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t_2^2,$$

so

$$0 < \max \left\{ \epsilon, \sqrt{\mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t_1^2} \right\} \leq \max \left\{ \epsilon, \sqrt{\mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t_2^2} \right\}.$$

Since $-t_1 > -t_2 > 0$, we have

$$\frac{-t_1}{\max\{\epsilon, \sqrt{\mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t_1^2}\}} > \frac{-t_2}{\max\{\epsilon, \sqrt{\mathbb{E}_n[1\{X\beta \geq 0, Xv < 0\}] - t_2^2}\}}.$$

Therefore, $f(t_1) > f(t_2)$. \square

Proof of Theorem 4. By the choice of cv and $q_{1-\alpha}$, we have $\mathbb{P}(T_n^*(\beta) \leq cv \mid \mathcal{X}_n) < 1 - \alpha - \eta$ for some positive number η . Consider a distribution of (U_1, \dots, U_n) given \mathcal{X}_n under which U_1, \dots, U_n are independent given

\mathcal{X}_n and $U_i \mid \mathcal{X}_n \sim N(0, \sigma^2)$ for every $i = 1, \dots, n$, where σ is chosen to satisfy

$$1 - \prod_{i=1}^n \left(\Phi \left(\frac{\min\{-X_i\beta, 0\}}{\sigma} \right) + 1 - \Phi \left(\frac{\max\{-X_i\beta, 0\}}{\sigma} \right) \right) \leq \eta.$$

(Such σ exists because the left hand side converges to 0 as $\sigma \rightarrow \infty$.) Note that if $Y_i = Y_i^*$ for every $i = 1, \dots, n$, then $T_n(\beta) = T_n^*(\beta)$. Then

$$\begin{aligned} \mathbb{P}(T_n(\beta) \leq cv \mid \mathcal{X}_n) &\leq \mathbb{P}(T_n(\beta) \leq cv, Y_i = Y_i^* \forall i = 1, \dots, n \mid \mathcal{X}_n) + 1 - \mathbb{P}(Y_i = Y_i^* \forall i = 1, \dots, n \mid \mathcal{X}_n) \\ &\leq \mathbb{P}(T_n^*(\beta) \leq cv \mid \mathcal{X}_n) + 1 - \mathbb{P}(Y_i = Y_i^* \forall i = 1, \dots, n \mid \mathcal{X}_n) \\ &< 1 - \alpha - \eta + 1 - \mathbb{P}(Y_i = Y_i^* \forall i = 1, \dots, n \mid \mathcal{X}_n). \end{aligned}$$

To show the statement of this theorem, the rest of the proof is going to show $\mathbb{P}(Y_i = Y_i^* \forall i = 1, \dots, n \mid \mathcal{X}_n) \geq 1 - \eta$. Since U_1, \dots, U_n are independent given \mathcal{X}_n , the events $\{Y_1 = Y_1^*\}, \dots, \{Y_n = Y_n^*\}$ are also independent given \mathcal{X}_n , and then $1 - \mathbb{P}(Y_i = Y_i^* \forall i = 1, \dots, n \mid \mathcal{X}_n) = 1 - \prod_{i=1}^n \mathbb{P}(Y_i \neq Y_i^* \mid \mathcal{X}_n)$. Since the distribution of U_i given \mathcal{X}_n satisfies

$$\begin{aligned} \mathbb{P}(Y_i = Y_i^* \mid \mathcal{X}_n) &= \mathbb{P}(1\{X_i\beta + U_i \geq 0\} = 1\{U_i \geq 0\} \mid \mathcal{X}_n) \\ &= \mathbb{P}(U_i < \min\{-X_i\beta, 0\} \text{ or } U_i \geq \max\{-X_i\beta, 0\} \mid \mathcal{X}_n) \\ &= \Phi \left(\frac{\min\{-X_i\beta, 0\}}{\sigma} \right) + 1 - \Phi \left(\frac{\max\{-X_i\beta, 0\}}{\sigma} \right), \end{aligned}$$

we have

$$1 - \mathbb{P}(Y_i = Y_i^* \forall i = 1, \dots, n \mid \mathcal{X}_n) = 1 - \prod_{i=1}^n \left(\Phi \left(\frac{\min\{-X_i\beta, 0\}}{\sigma} \right) + 1 - \Phi \left(\frac{\max\{-X_i\beta, 0\}}{\sigma} \right) \right) \leq \eta.$$

□

Proof of Theorem 5. In this proof, we focus on Eq. (3.7). Define $W = (2Y - 1)1\{Xb \geq 0, Xv < 0\}$. First, we are going to show that

$$\sqrt{n}\mathbb{E}_n[W] < -q_{1-\alpha} \max \left\{ \epsilon, \sqrt{\frac{\mathbb{E}_n[W^2]}{1 + q_{1-\alpha}^2/n}} \right\} \implies T_n(b) > q_{1-\alpha}. \quad (\text{A.9})$$

Suppose $\sqrt{n}\mathbb{E}_n[W] < -q_{1-\alpha} \max \left\{ \epsilon, \sqrt{\frac{\mathbb{E}_n[W^2]}{1 + q_{1-\alpha}^2/n}} \right\}$. Note that

$$\mathbb{E}_n[W] < 0 \quad (\text{A.10})$$

and

$$n\mathbb{E}_n[W]^2 > \max \left\{ \epsilon^2 q_{1-\alpha}^2, \mathbb{E}_n[W^2] \frac{q_{1-\alpha}^2}{1 + q_{1-\alpha}^2/n} \right\}.$$

The second inequality implies $n\mathbb{E}_n[W]^2 > \max\{\epsilon^2 q_{1-\alpha}^2, \mathbb{E}_n[W^2] q_{1-\alpha}^2 - \mathbb{E}_n[W]^2 q_{1-\alpha}^2\}$. Using Eq. (A.10), $-\sqrt{n}\mathbb{E}_n[W] > q_{1-\alpha} \max\{\epsilon, \sqrt{\mathbb{E}_n[W^2] - \mathbb{E}_n[W]^2}\}$ and then

$$\sqrt{n} \frac{-\mathbb{E}_n[W]}{\max\{\epsilon, \sqrt{\mathbb{E}_n[W^2] - \mathbb{E}_n[W]^2}\}} > q_{1-\alpha}$$

which implies $T_n(b) > q_{1-\alpha}$.

Then, we are going to show $\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n) \geq 1 - \rho$. Using Eq. (A.9), we have

$$\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n) \geq \mathbb{P}\left(\sqrt{n}\mathbb{E}_n[W] < -q_{1-\alpha} \max\left\{\epsilon, \sqrt{\frac{\mathbb{E}_n[W^2]}{1 + q_{1-\alpha}^2/n}}\right\} \mid \mathcal{X}_n\right).$$

Eq. (3.7) implies

$$\begin{aligned} \mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n) &\geq \mathbb{P}\left(\mathbb{E}_n[W] < \mathbb{E}[\mathbb{E}_n[W] \mid \mathcal{X}_n] + \sqrt{\frac{2\log(1/\rho)\mathbb{E}_n[1\{Xb \geq 0, Xv < 0\}]}{n}} \mid \mathcal{X}_n\right) \\ &= 1 - \mathbb{P}\left(\mathbb{E}_n[W] \geq \mathbb{E}[\mathbb{E}_n[W] \mid \mathcal{X}_n] + \sqrt{\frac{2\log(1/\rho)\mathbb{E}_n[1\{Xb \geq 0, Xv < 0\}]}{n}} \mid \mathcal{X}_n\right). \end{aligned}$$

Since $-1\{Xb \geq 0, Xv < 0\} \leq W_i \leq 1\{Xb \geq 0, Xv < 0\}$ for every $i = 1, \dots, n$, Hoeffding (1963)'s inequality implies

$$\mathbb{P}(T_n(b) > q_{1-\alpha} \mid \mathcal{X}_n) \geq 1 - \exp\left(-\frac{2n^2 \left(\sqrt{\frac{2\log(1/\rho)\mathbb{E}_n[1\{Xb \geq 0, Xv < 0\}]}{n}}\right)^2}{4 \sum_{i=1}^n 1\{X_i b \geq 0, X_i v < 0\}}\right) = 1 - \rho.$$

□

Proof of Corollary 1. Let $1 - \rho$ denote the maximum of the power bounds (3.9) and (3.10) in the statement of the corollary. If the expressions $\max\{0, \cdot\}$ in (3.9) and (3.10) are both zero for all $v \in \mathcal{V}_u$ and $v \in \mathcal{V}_l$, then the implication of the corollary is trivially satisfied. Thus suppose instead that the maximum of (3.9) and (3.10) is greater than zero. It follows that there is either a $v \in \mathcal{V}_u$ such that

$$1 - \rho = 1 - \exp\left(-\frac{1}{2} \left(\sqrt{n}\tilde{\zeta}_u(b, v) - q_{1-\alpha} \max\left\{\tilde{\epsilon}_u(b, v), (1 + q_{1-\alpha}^2/n)^{-1/2}\right\}\right)^2\right),$$

or a $v \in \mathcal{V}_l$ such that

$$1 - \rho = 1 - \exp\left(-\frac{1}{2} \left(\sqrt{n}\tilde{\zeta}_l(b, v) - q_{1-\alpha} \max\left\{\tilde{\epsilon}_l(b, v), (1 + q_{1-\alpha}^2/n)^{-1/2}\right\}\right)^2\right),$$

implying (3.9) in the former case and (3.10) in the latter. The conclusion of the corollary then follows from Theorem 5. □

Proof of Corollary 2. The first part of the corollary can be shown by first noting that $Q(b) \geq C(\gamma)$ implies that at least one of the following inequalities hold:

$$\begin{aligned} \mathbb{E}[\mathbb{E}_n[(2Y - 1)1\{Xb \geq 0, Xv < 0\}] \mid \mathcal{X}_n] &\leq -C(\gamma), \\ \mathbb{E}[\mathbb{E}_n[(1 - 2Y)1\{Xb \leq 0, Xv \geq 0\}] \mid \mathcal{X}_n] &\leq -C(\gamma). \end{aligned}$$

Then because $-C(1 - \rho)$ is less than or equal to each of the expressions on the right hand side of inequalities (3.7) and (3.8), at least one of (3.7) and (3.8) is true and Theorem 5 delivers the result.

For the second claim of the corollary, consider first that if $\sqrt{n}Q(b) \leq q_{1-\alpha} \max\{\epsilon, (1 + q_{1-\alpha}^2/n)^{-1/2}\}$ the

result holds trivially. Thus, suppose instead that $\sqrt{n}Q(b) > q_{1-\alpha} \max \{ \epsilon, (1 + q_{1-\alpha}^2/n)^{-1/2} \}$ and consider

$$\gamma = 1 - \exp \left(-\frac{1}{2} \left(\sqrt{n}Q(b) - q_{1-\alpha} \max \{ \epsilon, (1 + q_{1-\alpha}^2/n)^{-1/2} \} \right)^2 \right).$$

Then $Q(b) = C(\gamma)$ and the desired implication follows from the first part of the corollary. \square

Proof of Theorem 6. Assume w.l.o.g. that $\{i : \bar{p}_i \neq \tilde{p}_i\} = \{1, \dots, \bar{n}\}$, which can be achieved by rearranging i 's. Then $(\bar{p}_1, \dots, \bar{p}_{\bar{n}}) = (1/2, \dots, 1/2)$ and the test in (4.9) simplifies as

$$\bar{\phi}(y_1, \dots, y_{\bar{n}}) = \begin{cases} 1 & \text{if } \prod_{i=1}^{\bar{n}} \tilde{p}_i^{y_i} (1 - \tilde{p}_i)^{1-y_i} > k2^{-\bar{n}}, \\ \xi & \text{if } \prod_{i=1}^{\bar{n}} \tilde{p}_i^{y_i} (1 - \tilde{p}_i)^{1-y_i} = k2^{-\bar{n}}, \\ 0 & \text{otherwise.} \end{cases}$$

Define

$$\text{RP}(p_1, \dots, p_{\bar{n}}) \equiv \sum_{(y_1, \dots, y_{\bar{n}}) \in \{0,1\}^{\bar{n}}} \bar{\phi}(y_1, \dots, y_{\bar{n}}) \prod_{i=1}^{\bar{n}} p_i^{y_i} (1 - p_i)^{1-y_i}.$$

For the rest of the proof, we first show that $\text{RP}(p_1, \dots, p_{\bar{n}})$ is the probability that $\bar{\phi}(y_1, \dots, y_{\bar{n}})$ rejects the null hypothesis when $\mathbb{P}\{Y_i = 1 \mid \mathcal{X}_n\} = p_i$ for all i , that is,

$$\text{RP}(p_1, \dots, p_{\bar{n}}) = \sum_{(y_1, \dots, y_n) \in \{0,1\}^n} \bar{\phi}(y_1, \dots, y_n) \prod_{i=1}^n p_i^{y_i} (1 - p_i)^{1-y_i}, \quad (\text{A.11})$$

where the right-hand side of the above equation is the sum of 2^n terms instead of $2^{\bar{n}}$. From this, it follows from (4.10) that $\text{RP}(\bar{p}_1, \dots, \bar{p}_{\bar{n}}) = \alpha$. Subsequently we show that $(\bar{p}_1, \dots, \bar{p}_{\bar{n}})$ is the constrained maximizer of $\text{RP}(p_1, \dots, p_{\bar{n}})$ with respect to $p_1, \dots, p_{\bar{n}}$, subject to $p_1, \dots, p_{\bar{n}}$ being compatible with the null hypothesis. Thus (4.9) achieves size control, i.e., $\sum_{(y_1, \dots, y_n) \in \{0,1\}^n} \bar{\phi}(y_1, \dots, y_n) \prod_{i=1}^n p_i^{y_i} (1 - p_i)^{1-y_i} \leq \alpha$ for any sequence (p_1, \dots, p_n) under the null hypothesis. By Theorem 3.8.1 of Lehmann and Romano (2005) the test (4.9) is a most powerful test of the composite null $\beta = b$ against the simple alternative H_1 , and (4.11) is the least favorable distribution of (Y_1, \dots, Y_n) given \mathcal{X}_n .

First, to establish (A.11), we have

$$\begin{aligned} \text{RP}(p_1, \dots, p_{\bar{n}}) &= \sum_{(y_1, \dots, y_{\bar{n}}) \in \{0,1\}^{\bar{n}}} \bar{\phi}(y_1, \dots, y_{\bar{n}}) \prod_{i=1}^{\bar{n}} p_i^{y_i} (1 - p_i)^{1-y_i} \\ &= \sum_{(y_1, \dots, y_{\bar{n}}) \in \{0,1\}^{\bar{n}}} \bar{\phi}(y_1, \dots, y_{\bar{n}}) \sum_{(y_{\bar{n}+1}, \dots, y_n) \in \{0,1\}^{n-\bar{n}}} \prod_{i=\bar{n}+1}^n p_i^{y_i} (1 - p_i)^{1-y_i} \\ &= \sum_{(y_1, \dots, y_n) \in \{0,1\}^n} \bar{\phi}(y_1, \dots, y_{\bar{n}}) \prod_{i=1}^{\bar{n}} p_i^{y_i} (1 - p_i)^{1-y_i}, \end{aligned}$$

where the second equality holds because $\sum_{(y_{\bar{n}+1}, \dots, y_n) \in \{0,1\}^{n-\bar{n}}} \prod_{i=\bar{n}+1}^n p_i^{y_i} (1 - p_i)^{1-y_i} = 1$ and the third equality follows because $\bar{\phi}(y_1, \dots, y_{\bar{n}})$ does not depend on $y_{\bar{n}+1}, \dots, y_n$.

For the next step of the proof, notice that (4.8) implies that, under the null hypothesis, $(\mathbb{P}(Y_i = 1 \mid \mathcal{X}_n) - 1/2)(\tilde{p}_i - 1/2) \leq 0$ for every $i = 1, \dots, \bar{n}$. Thus we next show that $(\bar{p}_1, \dots, \bar{p}_{\bar{n}})$ is the constrained maximizer of $\text{RP}(p_1, \dots, p_{\bar{n}})$ subject to $(p_i - 1/2)(\tilde{p}_i - 1/2) \leq 0$ for all $i = 1, \dots, \bar{n}$. Let $\bar{\phi}(d, y_{-j})$ be the shorthand for

$\bar{\phi}(y_1, \dots, y_{j-1}, d, y_{j+1}, \dots, y_n)$, where $d \in \{0, 1\}$. Note that

$$\begin{cases} \bar{\phi}(1, y_{-j}) - \bar{\phi}(0, y_{-j}) \geq 0 & \text{if } \tilde{p}_j > 1/2, \\ \bar{\phi}(1, y_{-j}) - \bar{\phi}(0, y_{-j}) \leq 0 & \text{if } \tilde{p}_j < 1/2, \end{cases}$$

for every $j = 1, \dots, \bar{n}$, and that $\tilde{p}_j \neq 1/2$ because $\bar{p}_j \neq \tilde{p}_j$ for all $j = 1, \dots, \bar{n}$. Since

$$\begin{aligned} \frac{\partial}{\partial p_j} \text{RP}(p_1, \dots, p_{\bar{n}}) &= \sum_{(y_1, \dots, y_{\bar{n}}) \in \{0, 1\}^{\bar{n}}} (2y_j - 1) \bar{\phi}(y_1, \dots, y_{\bar{n}}) \prod_{i \neq j} p_i^{y_i} (1 - p_i)^{1-y_i} \\ &= \sum_{y_{-j}} \bar{\phi}(1, y_{-j}) \prod_{i \neq j} p_i^{y_i} (1 - p_i)^{1-y_i} - \sum_{y_{-j}} \bar{\phi}(0, y_{-j}) \prod_{i \neq j} p_i^{y_i} (1 - p_i)^{1-y_i} \\ &= \sum_{y_{-j}} (\bar{\phi}(1, y_{-j}) - \bar{\phi}(0, y_{-j})) \prod_{i \neq j} p_i^{y_i} (1 - p_i)^{1-y_i}, \end{aligned}$$

we have

$$\begin{cases} \frac{\partial}{\partial p_j} \text{RP}(p_1, \dots, p_{\bar{n}}) \geq 0 & \text{if } \tilde{p}_j > 1/2, \\ \frac{\partial}{\partial p_j} \text{RP}(p_1, \dots, p_{\bar{n}}) \leq 0 & \text{if } \tilde{p}_j < 1/2. \end{cases}$$

Thus $(1/2, \dots, 1/2)$ maximizes $\text{RP}(p_1, \dots, p_{\bar{n}})$ with respect to $(p_1, \dots, p_{\bar{n}})$ subject to $(p_i - 1/2)(\tilde{p}_i - 1/2) \leq 0$ for all $i = 1, \dots, \bar{n}$, completing the proof. \square

Proof of Corollary 3. The corollary follows directly from Theorem 6 and the reasoning given in the text. \square

B Additional Simulation Results




Figure 3.a: Design 1.




Figure 3.b: Design 2.




Figure 3.c: Design 3.




Figure 3.d: Design 4.

Figure 3: Non-rejection frequencies with $1 - \alpha = 90\%$ with true $\theta_0 = 1$ and $n = 250$.




Figure 4.a: Design 1.




Figure 4.b: Design 2.




Figure 4.c: Design 3.




Figure 4.d: Design 4.

Figure 4: Non-rejection frequencies with $1 - \alpha = 90\%$ with true $\theta_0 = 2$ and $n = 100$.




Figure 5.a: Design 1.




Figure 5.b: Design 2.




Figure 5.c: Design 3.




Figure 5.d: Design 4.

Figure 5: Non-rejection frequencies with $1 - \alpha = 90\%$ with true $\theta_0 = 2$ and $n = 250$.