

SIMPLE SUBVECTOR INFERENCE ON SHARP IDENTIFIED SET IN AFFINE MODELS ¹BULAT GAFAROV ²

This paper studies a regularized support function estimator for bounds on components of the parameter vector in the case in which the identified set is a polygon. The proposed regularized estimator has three important properties: (i) it has a uniform asymptotic Gaussian limit in the presence of flat faces in the absence of redundant (or overidentifying) constraints (or vice versa); (ii) the bias from regularization does not enter the first-order limiting distribution; (iii) the estimator remains consistent for sharp identified set for the individual components even in the non-regular case. These properties are used to construct *uniformly valid* confidence sets for an element θ_1 of a parameter vector $\theta \in \mathbb{R}^d$ that is partially identified by affine moment equality and inequality conditions. The proposed confidence sets can be computed as a solution to a small number of linear and convex quadratic programs, which leads to a substantial decrease in computation time and guarantees a global optimum. As a result, the method provides uniformly valid inference in applications in which the dimension of the parameter space, d , and the number of inequalities, k , were previously computationally unfeasible ($d, k = 100$). The proposed approach can be extended to construct confidence sets for intersection bounds, to construct joint polygon-shaped confidence sets for multiple components of θ , and to find the set of solutions to a linear program. Inference for coefficients in the linear IV regression model with an interval outcome is used as an illustrative example.

KEYWORDS: affine-moment inequalities, asymptotic linear representation, higher-order analysis, delta method, interval data, intersection bounds, partial identification, regularization, strong approximation, stochastic programming, subvector inference, uniform inference.

University of California, Davis, Department of Agricultural and Resource Economics

¹This version is December 1, 2023. Earlier versions of the paper was circulated with the title “Inference on Scalar Parameters in Set-Identified Affine Models” and “Inference in high-dimensional Set-Identified Affine Models”. The work is based on a chapter in my PhD dissertation (Gafarov, 2017). First-draft date: November 10, 2015.

²I am extremely grateful to Joris Pinkse and Patrik Guggenberger for their very helpful and detailed comments. I would like to thank (in alphabetical order) Donald Andrews, Andres Aradillas-Lopez, Christian Bontemps, Ivan Canay, Peng Ding, Graham Elliott, Zheng Fang, Dalia Ghanem, Joachim Fryberger, Ronald Gallant, Michael Gechter, Marc Henry, Keisuke Hirano, Sung Jae Jun, Nail Kashaev, Francesca Molinari, Demian Pouzo, Adam Rosen, Thomas Russell, Andres Santos, Xiaoxia Shi, Jing Tao, and Alexander Torgovitsky for their comments and suggestions.

1. INTRODUCTION

Strong econometric assumptions can lead to poor estimates. Sometimes, moment inequalities can provide alternative estimates under weaker assumptions. Linear models with interval-valued outcome data are a good example.¹ It is common practice to replace income-bracket data with the corresponding midpoints when estimating the returns to schooling (Trostel et al. (2002)). However, the conventional approach is applicable only under strong assumptions about the distribution of the residual term.² The affine moment inequality approach to interval-valued data proposed by Manski and Tamer (2002) can set-identify the return to schooling without such strong assumptions.

Multiple methods can be used to construct confidence sets (CS) for parameters defined by moment inequalities. The pioneering procedures of Chernozhukov et al. (2007) and Andrews and Soares (2010) (AS) and their subsequent refinements by Bugni et al. (2016) (BCS) and Kaido et al. (2015) (KMS) are powerful statistical methods that solve this inference problem in the small-dimensional case. Some applications, such as panel or semiparametric regression models with interval-measured outcome variables, have a high-dimension parameter space, which poses a computational challenge for the existing procedures.³

I propose a novel regularized support function estimator for the lower and upper extremes of the identified set for an element θ_1 of an unknown parameter vector $\theta \in \mathbb{R}^d$ in models defined by affine moment equalities and inequalities. In the example of returns to schooling, θ_1 corresponds to the returns to schooling and $\theta \in \mathbb{R}^d$ to the full vector of the regression coefficients. The novel estimator has a closed-form asymptotic Gaussian distribution, which I use to construct uniformly valid confidence bounds and confidence intervals for θ_1 . The proposed sets has valid asymptotic coverage probability uniformly over a class of data-generating processes (DGP). The uniformity in DGP is a desirable property, as it results in better coverage-probability control in small samples when compared with point-wise analogs in *nonregular statistical models* such as the affine moment inequality model.

The regularized support function proposed in this paper is a solution to a convex quadratic program that minimizes the sum of θ_1 and a penalty $\mu_n \|\theta\|^2$ with $\mu_n \rightarrow 0$, subject to the sample moment restrictions. If the set of optima for $\mu = 0$ is not a singleton, this additional convex term selects the optimum with the minimal norm as n increases. The standard errors are computed using the sample variance of the weighted moment conditions at the unique optima. To correct the asymptotic bias resulting from the regularization exactly, I suggest using the argmin of the regularized program with a larger tuning parameter $\kappa_n \rightarrow 0$. If $\kappa_n/\mu_n \rightarrow \infty$ as $n \rightarrow \infty$, then the bias correction does not affect the asymptotic distribution of the estimator. To achieve a uniformly valid confidence interval (CI), I replace the exact correction with an upper bound on the maximum of $\mu_n \|\theta\|^2$ over the argmin set of the nonregularized program.

The proposed CIs have several attractive statistical and computational properties that make them viable in high-dimensional affine moment inequality models.

First, the estimator of the regularized support function has an asymptotically linear (Bahadur-Kiefer) representation that provides easy-to-compute asymptotic standard error. Consequently, this paper is the first to propose a closed-form estimator of the bounds on θ_1 convex moment inequality models

¹Other examples of affine-moment inequalities include monotone instrumental variables (Manski and Pepper (2000), Freyberger and Horowitz (2015)) and models with missing data (Manski (2003)).

²Another common approach is to assume Gaussian distribution for the residuals and apply the maximum likelihood method (Stewart (1983)).

³For example, Trostel et al. (2002) consider a panel regression with more than sixty variables including country fixed effects, time effects, exogenous demographic control variables, and their interactions.

with asymptotic Gaussian distribution in the nonregular case (in the absence of strict convexity). In contrast, the estimator of the ordinary support function estimator used in the existing literature (Beresteanu and Molinari (2008), Kaido and Santos (2014), Freyberger and Horowitz (2015, FH), Gafarov et al. (2018), among others) can have a non-Gaussian asymptotic distribution, which complicates uniform inference. To establish the uniform remainder bound in the Bahadur-Kiefer expansion, I developed a novel second-order directional envelope theorem, which is a theoretical result of independent interest.

Second, the proposed approach requires only a fraction of the computational time of the existing uniform procedures if θ has a large number of dimensions.⁴ The computational cost is low since it involves only four quadratic programs, it does not require any resampling, and it depends on covariance of the moment conditions at two points.

The computation time for my procedure increases slowly in the dimension of $\theta \in \mathbb{R}^d$ and takes only 0.1 second for $d = 20$ and $k = 40$ moment inequalities and 2.5 seconds for $d = 100$ and $k = 200$. As a result, the proposed method can address parameter θ with a large dimension and a large number of moment conditions. In contrast, the existing uniform-inference methods for moment inequalities proposed by AS, KMS, and BCS are based on costly nonconvex optimization. Indeed, despite the fact that the moment inequalities are convex (linear) the standardized moment conditions are not convex which can potentially result in creation of multiple local optima which complicate computation for these statistical procedures (see Appendix Sections C.1 and C.2).

I provide an example of an affine moment inequality model by showing that the number of local optimal solutions in the existing uniform procedures (AS, BCS, and KMS) can grow exponentially with the dimension d . As a result, the procedures take more computational time and can produce misleadingly short CIs if the optimization routine fails to find the global optimum. It takes 100 seconds to compute the CI of AS in an affine model with $d = 20$ and $d = 40$ moment inequalities which is 1000 times longer than the newly proposed method.

The simulation evidence suggests that the computational-speed gains come without a substantial loss of statistical power. In the nonregular cases, the proposed uniform CIs have length properties that are not worse than those of the existing uniform methods. (The proposed *uniform* CI has length within simulation error from the projection CI of AS in the Monte Carlo (MC) design considered in this paper.) In the regular cases, the novel confidence bounds attain the efficiency bound of the usual support function estimators (as shown in Kaido and Santos (2014)).

The proposed idea of regularizing support functions for linear moment inequality inference got a subsequent development in a recent work Cho and Russell (2023).⁵ The authors argue that if one is satisfied with an inference on an *enlarged* identified set, one can simply regularize the moment inequalities by adding random noise to their coefficients under milder regularity conditions. Such operation restores a Gaussian limit of the perturbed estimated support function thus allowing for conventional bootstrap inference on the *enlarged* identified set. Since the inference is done on the enlarged identified set, one does not need to impose any constraint qualification conditions on the moment inequalities — they are satisfied automatically with probability 1 after addition of the random noise to the coefficients. Unfortunately, this approach results in confidence sets with zero power against all local alternatives corresponding to other nested enlargements of the identified set.

⁴In the asymptotic analysis, I only consider the case of fixed dimension of θ and a number fixed of inequalities for analytical simplicity. The goal of the present study is to focus on the computational difficulties resulting from the high dimensional moment inequalities. There are recent papers that are concerned with the impact of the growing number of inequalities on the statistical properties of the inference procedures, see for example Belloni et al. (2018).

⁵The first draft of Cho and Russell (2023) was circulated on 27 May 2019 on Arxiv depository two years after the initial distribution of the present paper as a dissertation chapter in Gafarov (2017).

In practical terms, it means that the confidence sets can be very conservative if too much noise was added or fail to control size in small samples if the added noise was insufficient (there is no theory that would determine the minimal level of required noise for a given sample size). In contrast, the theory provided in the present paper explicitly studies the impact of the tuning parameter choice on the size of the regularized identified set. Such analysis requires imposing constraint qualification conditions on the moment inequalities (Assumptions 4-5). Under these conditions, Theorem 1 shows that for sufficiently small value of the tuning parameters μ and κ both lower and upper bounds based on the regularized value function coincide exactly with the non-regularized support function in the regular case when the set of primal solutions is a singleton. More generally under the maintained assumptions, the regularized estimators remain consistent for the bounds on the sharp (non-enlarged) identified set thus providing non-trivial local power against the relevant alternatives.

Identified sets defined by affine inequalities appear in various economic applications—in particular, those dealing with discrete variables and shape restrictions. Linear models with interval outcome, which were originally studied in Manski and Tamer (2002) and Haile and Tamer (2003), are just one example of affine inequalities. Other examples include bounds on marginal effects in panel dynamic discrete-choice models (Honoré and Tamer (2006), Torgovitsky (2016, 2019)), bounds on average treatment effects (Kasy (2016), Laffers (2018), Russell (2017)), nonparametric IV models with shape restrictions (Manski and Pepper (2000), Freyberger and Horowitz (2015)), errors in variables (Molinari (2008)), intersection bounds (Honoré and Lleras-Muney (2006)), revealed-preference restrictions (Kline and Tartari (2016), Shi et al. (2018)), and game-theoretic models (Pakes et al. (2015), Syrgkanis et al. (2017)).

I also contribute to the growing literature on inference on nondifferentiable functions and regularized estimators. Bounds on components of a parameter characterized by linear moment conditions considered in this paper are an example of a nondifferentiable (*nonregular*) function of a parameter (the expectation of the data) that has an asymptotically normal estimator (the sample mean). Distributions of nondifferentiable functions of the normal estimator are hard to approximate using the standard methods. (See Section 2.3 for detailed discussion.) I propose differentiable lower and upper bounds that converge to the nondifferentiable function of interest as sample size grows. Since the bounds are regular parameters themselves, the standard delta method and bootstrap can be used to conduct one- or two-sided inference on the bounds. In the regular case, these bounds collapse and coincide with the original parameter of interest, which results in a \sqrt{n} -consistent and asymptotically normal estimator. In the nonregular case, the bounds converge at a slower rate and results in a locally biased estimator, which is acceptable for valid one-sided inference.

Another interesting statistical problem that appears in many applications is inference for extrema of finitely many means of random variables. It is known as the *intersection-bounds problem* (Hall and Miller (2010) and Chernozhukov et al. (2013)) and can be framed as a value of a linear program. The regularized support function estimator can also be used for uniform delta-method CSs in this setting (see Appendix Section B.3). The approach considered here is expected to have similar statistical properties to Chernozhukov et al. (2013) but has the advantage of closed-form standard errors and critical values, which correspond to the standard normal distribution.

The paper is structured as follows. Section 2 describes the setup, gives examples, and summarizes the available literature. Section 3 provides the main results, applies them to uniform inference on projections, and discusses extensions (overidentified inequality models, joint CSs, characterization of argmin sets). Section 4 provides the results of Monte Carlo experiments. Section 5 concludes.

I use the following notational conventions. \triangleq denotes definitions. $\mathbb{E}_P[\cdot]$ denotes expectation with

respect to a probability distribution P . Uppercase English letters denote random variables (scalar, vector, or matrix valued), and lowercase letters denote the corresponding realizations, for example, W_i and w_i . \mathbb{P}_n denotes the empirical distribution, and $f(0^+)$ denotes $\lim_{x \downarrow 0} f(x)$. The vector $e_j \triangleq (0, \dots, 1, \dots, 0)'$ is the j -th coordinate vector, where 1 occurs at position j . e_j is the projector on the j -th coordinate. The symbol \mathcal{J} denotes a finite set of indices $\mathcal{J} \triangleq \{i_1, \dots, i_\ell\} \subset \mathbb{N}$ and $\mathbb{J} \triangleq (e_{i_1}, \dots, e_{i_\ell})'$ as a coordinate projection matrix in the corresponding Euclidean space. $|\mathcal{J}|$ denotes the cardinality of the set \mathcal{J} . u.h.c. stands for upper hemicontinuous correspondence. The symbol $\text{sVar}(x)$ denotes the sample variance, where $\text{sVar}(x) = \frac{1}{n} \sum_{i=1}^n x_i^2 - \left(\frac{1}{n} \sum_{i=1}^n x_i\right)^2$.

2. SETUP, MOTIVATING EXAMPLES, AND RELATED LITERATURE

2.1. Support function and projections of identified sets

Consider a support function for a polygon $\Theta(P) \subset \mathbb{R}^d$ (a set defined by a system of linear equalities and inequalities) that depends on a data generating process parametrized by a measure P evaluated at a direction $e_1 \triangleq (1, 0, \dots)'$ in \mathbb{R}^d ,⁶

$$(2.1) \quad \underline{v}(P) \triangleq \min_{\theta \in \Theta(P)} e_1' \theta.$$

The set $\Theta(P)$, also referred to as an *identified set* for a parameter vector $\theta \in \mathbb{R}^d$, consists of all solutions to the following system of affine moment equalities/inequalities:

$$(2.2) \quad \begin{cases} \mathbb{E}_P g_j(W, \theta) = 0, & j \in \mathcal{J}^{eq}, \\ \mathbb{E}_P g_j(W, \theta) \leq 0, & j \in \mathcal{J}^{ineq}. \end{cases}$$

Here, $g_j(W, \theta) \triangleq \sum_{\ell=1}^d W_{j\ell} \theta_\ell - W_{j(d+1)}$. I consider the setup with finitely many unconditional moment functions (that is, $|\mathcal{J}^{eq} \cup \mathcal{J}^{ineq}| = k$, $|\mathcal{J}^{eq}| = p$, $0 \leq p \leq d$, $k < \infty$)⁷. The random matrix corresponding to an individual observation W has probability measure P with sample space $\mathbb{R}^{k \times (d+1)}$.

This setup reduces to a finite-dimensional parametric statistical model once the support function evaluated at e_1 is represented as a linear program,

$$(2.3) \quad \begin{aligned} \underline{v}(P) &= \min_{\theta \in \mathbb{R}^d} e_1' \theta \\ \text{s.t. } & e_j' A_P \theta = e_j' b_P, \quad j \in \mathcal{J}^{eq}, \\ & e_j' A_P \theta \leq e_j' b_P, \quad j \in \mathcal{J}^{ineq}. \end{aligned}$$

Here, the coefficients on the left-hand side, A_P , and the right-hand side, b_P , taken together constitute matrix $\mathbb{E}_P W \triangleq (A_P | b_P)$. Following optimization theory terminology, I occasionally refer to the support function at e_1 as the *value* of Program (2.3) and to the corresponding argmin set as the set of *optimal solutions*.

⁶Rockafellar (see, for example, 1970, chapter 13)

⁷There are, of course, interesting economic applications that involve overidentifying moment equalities which results in $p > d$. In particular, conditional moment equalities fall into this category since they imply potentially infinitely many unconditional moments corresponding to various transformations of the instrumental variables. In order to apply the approach proposed in this paper one needs to find a good just- or under-identified combination of the moment equalities.

The focus on the first coordinate of θ as an objective of (2.1) is without loss of generality. As discussed in Section 3.3.2, the support function evaluated at any unit direction vector $a \in \mathbb{R}^d$ can be represented in form (2.1) after some redefinition of the parameter space (matrices A_P and b_P can be functions of the unit vector a).

To ensure boundedness of the support function, I assume that System (2.2) includes inequalities that make the identified set compact,

$$(2.4) \quad -\infty < -\underline{c}_\ell \leq \theta_\ell \leq \bar{c}_\ell < \infty \quad \text{for } \ell = 1, \dots, d,$$

for some constants $\underline{c}, \bar{c} \in \mathbb{R}_+^d$. Moreover, I impose the following assumption:

ASSUMPTION 1 $\Theta(P)$ is non-empty for the probability measure P .

This assumption implies that $\underline{v}(P) < +\infty$. It is valid whenever the model corresponding to (2.2) is correctly specified.

Under Assumption 1, support functions evaluated at $\pm e_j$ characterize projections of $\Theta(P)$ on individual coordinates j of θ . In particular, the marginal (projected) identified set for the coordinate θ_1 can be represented as an interval $\mathcal{S}(P) = [\underline{v}(P), \bar{v}(P)]$, where the bounds are support functions evaluated at directions e_1 and $-e_1$. Indeed, the upper bound can also be written as a (minus) support function at e_1 ,

$$(2.5) \quad \bar{v}(P) \triangleq \max_{\theta \in \Theta(P)} \{e_1' \theta\} = - \min_{\theta \in \Theta(P)} \{-e_1' \theta\}.$$

The analysis for the upper bound is analogous to the one for the lower bound, so from here on I focus on the lower bound.

2.2. Motivating examples

The polygon-shaped identified sets considered in this paper appear in many econometric models that feature discrete data. I use the linear IV model with interval outcome (for example, Chernozhukov et al., 2007) to illustrate ideas throughout the paper.

EXAMPLE 1 (Linear IV model with interval-valued outcome) Consider a linear model for a random vector (Y, X, Z) satisfying

$$\mathbb{E}_P [Y - \theta' X | Z] = 0,$$

where $\theta \in \mathbb{R}^d$ is the vector of regression coefficients. The true outcome Y is unobserved; only its a.s. bounds $[\underline{Y}, \bar{Y}]$ are observed. Suppose that the vector of instrumental variables Z has a finite support $\{z_1, \dots, z_K\} \subset \mathbb{R}^d$. In this case, the model implies a polygon-shaped identified set $\Theta(P)$ for θ defined by a set of moment inequalities,

$$(2.6) \quad \begin{cases} \mathbb{E}_P [\underline{Y} 1\{Z = z_j\}] \leq \theta' \mathbb{E}_P [X 1\{Z = z_j\}], & j = 1, \dots, K, \\ \mathbb{E}_P [\bar{Y} 1\{Z = z_{j-K}\}] \geq \theta' \mathbb{E}_P [X 1\{Z = z_{j-K}\}], & j = K + 1, \dots, 2K. \end{cases}$$

These inequalities can be represented in the standard form (2.2) with $p = 0$, $k = 2K$, and the following observation matrix:

$$W_{j\ell} \triangleq \begin{cases} -X_\ell 1\{Z = z_j\}, & \text{for } j = 1, \dots, K, \\ X_\ell 1\{Z = z_{j-K}\}, & \text{for } j = K + 1, \dots, 2K, \end{cases}$$

$$W_{j(d+1)} \triangleq \begin{cases} \underline{Y} 1\{Z = z_j\}, & \text{for } j = 1, \dots, K, \\ -\bar{Y} 1\{Z = z_{j-K}\}, & \text{for } j = K + 1, \dots, 2K. \end{cases}$$

If it is known a priori that for some support points j

$$\mathbb{E}_P [\underline{Y} 1\{Z = z_j\}] = \mathbb{E}_P [\bar{Y} 1\{Z = z_{j-K}\}],$$

then one can replace the corresponding pair of inequalities with a single equality,

$$(2.7) \quad \mathbb{E}_P \left[\frac{1}{2} (\underline{Y} + \bar{Y}) 1\{Z = z_j\} \right] = \theta' \mathbb{E}_P X 1\{Z = z_j\}.$$

In this case, p (the number of equality restrictions) is equal to the number of such support points. One can further incorporate additional shape restrictions information such as signs of components of θ in a form of linear inequalities to narrow the identified set. \square

The following example provides a particular data set and an economic application of the interval-outcome IV regression with discrete regressors and a large-dimensional parameter space.

EXAMPLE 2 (The returns to schooling) [Trostel et al. \(2002\)](#) study economic returns to schooling for twenty-eight countries using International Social Survey Programme data from 1985 to 1995. They estimate the conventional [Mincer et al. \(1974\)](#) model of earnings (the human capital earnings function), which has Y , the log of hourly wages, satisfying

$$(2.8) \quad \mathbb{E}_P [Y - \theta' X | Z] = 0,$$

where the first regressor, X_1 , is the years of schooling; the other components of X play a role of additional controls. The component θ_1 is then interpreted as the return to schooling. It is equal to the percentage change in wages due to an additional year of schooling. To correct for the endogeneity bias in θ_1 resulting from omitting the latent-ability variable, one can use an instrument vector Z that correlates with X_1 (for example, an indicator of whether a good school is in proximity or an indicator of the quarter of birth).

Exact measurements of Y are not available for some countries (including the US); only hourly-income-bracket data $[\underline{Y}, \bar{Y}]$ are available.

Instrumental variables Z (coinciding with the control variable) considered in [Trostel et al. \(2002\)](#) take discrete values. The variables include annual fixed effects, union status, marital status, age and age squared, and country-year dummies (in the case of the aggregate equation). With inclusion of the country and time effects, d can be larger than 60. Because of the large number of support points for Z , the corresponding system (2.2) would also have a large number of linear moment inequalities k . \square

The conventional technique to estimate the IV regression with interval-outcome data is to replace the interval data with the corresponding midpoints and estimate the model using the OLS method.

This technique is valid only under the unreasonably strong condition

$$(2.9) \quad \mathbb{E}_P \left[\left(Y - \frac{1}{2}(\underline{Y} + \bar{Y}) \right) Z \right] = 0.$$

If Equation (2.9) is violated, then the OLS estimator with midpoints is inconsistent for the true parameter θ .⁸ Without assuming that (2.9) is satisfied, support function estimators (defined below) provide consistent bounds on marginal identified sets of the true parameter θ (see Beresteanu and Molinari, 2008; Bontemps et al., 2012).

When constructing valid CSs on projections of identified sets, strictly speaking, one needs only a lower bound on the support function at e_1 in (2.1). Other applications may instead require estimators of an upper bound on the (minimum) value of an optimization problem. The upper bound can be used, for example, to bound the set of optimal solutions or construct a consistent *inner* estimator of a convex identified set (the inner set estimator has important applications in inference; see, for example, Bugni et al., 2017).

2.3. Review of existing results about asymptotic distribution of support function estimators for polygon-shaped identified sets

Following Beresteanu and Molinari (2008) the parameter $\underline{v}(P)$ (the support function at e_1) can be estimated using a sample analog,

$$(2.10) \quad \hat{v}_n = \min_{\theta \in \mathbb{R}^d} e_1' \theta$$

$$(2.11) \quad \text{s.t.} \quad \begin{cases} \frac{1}{n} \sum_{i=1}^n g_j(W_i, \theta) = 0, & j \in \mathcal{J}^{eq}, \\ \frac{1}{n} \sum_{i=1}^n g_j(W_i, \theta) \leq 0, & j \in \mathcal{J}^{ineq}, \end{cases}$$

where the observations W_i are independent copies of the random matrix W . The asymptotic distribution of this estimator has been extensively studied in the case of the strictly convex identified set. The strict convexity implies that the support function at e_1 is differentiable in the data W (under additional regularity conditions discussed below). As a consequence, its sample analog has an asymptotic Gaussian distribution, admits of bootstrap inference, and attains semiparametric efficiency (Kaido and Santos, 2014). In contrast to the strictly convex case, the Gaussian limit is not guaranteed anymore in the general (non-deterministic) linear moment inequality model (Kaido and Santos (2014) only allow for deterministic linear inequalities under particular regularity conditions).

The stochastic programming approach (Shapiro, 1991) enables an asymptotic analysis of support function estimators for a fixed direction in the general convex case, which includes the linear moment inequality model. In this section, I briefly review the two main ideas in this approach, Lagrangian duality of convex programs and the delta method for directionally differentiable functions. I conclude with an overview of the alternative inference approaches and their relation to the stochastic programming methods.

The Lagrangian duality theory provides additional, often more convenient, formulations of convex programs. Namely, the (primal) program in (2.1) has the same value as its Lagrangian dual formulation,

$$(2.12) \quad \underline{v}(P) = \max_{\lambda \in \mathbb{R}^p \times \mathbb{R}_+^{k-p}} \{-\lambda' b_P\},$$

⁸The OLS is, however, consistent for the best linear predictor of the midpoint that may not have the desirable economic interpretation; see, for example, Shi (2020).

$$\text{s.t. } \lambda' A_P = -e_1',$$

given Assumption 1. If the dual program has a bounded set of solutions $\underline{\lambda}(P)$, the support function for a given direction has bounded directional derivatives in A_P, b_P (defined precisely below). Proposition 5.45 in [Bonnans and Shapiro \(2000\)](#) provides a necessary and sufficient condition for $\underline{\lambda}(P)$ to be bounded:

ASSUMPTION 2 (Slater's condition) *There exist $\theta \in \Theta(P)$ s.t. $\mathbb{E}_P g_j(W, \theta) < 0$ for all $j \in \mathcal{J}^{ineq}$.*

Assumption 2 enables another, Lagrangian min/max, representation of (2.1), which is particularly convenient for computing derivatives of the support function at a given direction for the delta-method-based inference procedures. Suppose that $\Lambda \subset \mathbb{R}^p \times \mathbb{R}_+^{k-p}$ is some compact set that contains $\underline{\lambda}(P)$. The support function for a given direction e_1 can then be represented as (see, for example, [Bonnans and Shapiro, 2000](#), p. 437)

$$(2.13) \quad \underline{v}(P) = \min_{\theta \in \Theta} \max_{\lambda \in \Lambda} \{e_1' \theta + \lambda'(A_P \theta - b_P)\}.$$

The directional derivative for the value of this min/max program is provided by a corresponding envelope theorem (for example, [Shapiro et al., 2014](#), Theorem 7.28). Namely, the envelope theorem gives a derivative of any perturbed version of the min/max program,

$$(2.14) \quad \underline{v}(P, t) \triangleq \min_{\theta \in \Theta} \max_{\lambda \in \Lambda} \{e_1' \theta + \lambda'((A_P + t h_{A,t}) \theta - (b_P + t h_{b,t}))\},$$

with respect to a scalar t for any uniformly converging sequence of directions $h_t \triangleq (h_{A,t}, h_{b,t}) \rightarrow (h_A, h_b)$. The derivative takes the form

$$(2.15) \quad \lim_{t \rightarrow 0} \frac{\underline{v}(P, t) - \underline{v}(P)}{t} = \min_{\theta \in \underline{\theta}(P)} \max_{\lambda \in \underline{\lambda}(P)} \{\lambda'(h_A \theta - h_b)\},$$

where $\underline{\theta}(P)$ and $\underline{\lambda}(P)$ are sets of primal and dual optima for (2.1). Either of the two sets can be nonsingletons as illustrated below.


EXAMPLE 3 (Bivariate linear IV with an interval outcome) Revisit Example 1. Suppose that $\theta \in \mathbb{R}^2$, $X = Z$, and regressors Z_1 and Z_2 take values in $\{0, 1\}$ with $\mathbb{E}Z_1 = \mathbb{E}Z_2 = \frac{1}{2}$. As in (2.6), the identified set for θ can be characterized by eight inequality constraints,

$$(2.16) \quad \mathbb{E}[Y \psi_z(Z)] \leq \mathbb{E}[Z_1 \psi_z(Z)] \theta_1 + \mathbb{E}[Z_2 \psi_z(Z)] \theta_2 \leq \mathbb{E}[\bar{Y} \psi_z(Z)],$$

where indicator functions $\psi_z(Z) = 1\{Z = z\}$ correspond to all combinations of $z \in \{0, 1\}^2$. Suppose, for illustrative purposes, we are interested only in the identified set $\Theta(P)$ defined by the following subsystem of four inequalities:

$$(2.17) \quad \mathbb{E}[Y Z_1] \leq \frac{1}{2} \theta_1 + \theta_2 \mathbb{E}[Z_1 Z_2] \leq \mathbb{E}[\bar{Y}_i Z_1], \mathbb{E}[Y(1 - Z_1)] \leq \theta_2 \mathbb{E}[(1 - Z_1) Z_2] \leq \mathbb{E}[\bar{Y}(1 - Z_1)].$$

In order to represent the identified set on a diagram, suppose further that the a.s. bounds on the outcome variable Y satisfy $\mathbb{E}_P[\bar{Y}|Z_1 = i] = -\mathbb{E}_P[Y|Z_1 = i] = \frac{1}{2} \Delta_i \geq 0$ for $i \in \{0, 1\}$ (that is, Δ_i is the average length of the outcome interval depending on Z_1).

FIGURE 1.— The identified sets in Example 3 for various values of ρ .

The shape of the full identified set $\Theta(P)$ depends on the value of $\rho \triangleq \mathbb{E}(Z_1 Z_2)$. The corresponding marginal identified set for θ_1 can be written in an explicit form, $\mathcal{S}(P) = [-\Delta_1 - 2|\rho|\Delta_0, \Delta_1 + 2|\rho|\Delta_0]$. If $\rho = 0$, $\underline{\theta}_2$ —the second coordinate of the set of primal optima of the program in (2.1)—is not uniquely defined.

The case of a nonsingleton set of dual optima occurs when the number of binding inequalities at the primal optimum is larger than the dimension of the parameter space. Suppose that we further restrict the parameter space by imposing $\theta_2 = 0$. Consider the following two moment inequalities from system (2.16):

$$(2.18) \quad -\frac{1}{2}\theta_1 \leq -\mathbb{E}[\underline{Y}Z_1], \quad -\frac{1}{2}\theta_1 \leq -\mathbb{E}[\underline{Y}].$$

The support function at e_1 corresponding to (2.18) takes the explicit form

$$(2.19) \quad \underline{v}(P) = \underline{\theta}_1 = 2 \min\{\mathbb{E}[\underline{Y}Z_1], \mathbb{E}[\underline{Y}]\}.$$

If $\mathbb{E}[\underline{Y}] = \mathbb{E}[\underline{Y}Z_1]$, then both inequalities in (2.18) are binding at the optimum, resulting in a nonsingleton dual-optimum set (compare with the parameter-on-the-boundary problem and the intersection-bounds problem considered in Andrews (2001) and Chernozhukov et al. (2013), respectively). Indeed, the dual formulation of the support function at e_1 takes form

$$(2.20) \quad \underline{v}(P) = \max_{\lambda \in \mathbb{R}_+^2} \{\mathbb{E}[\underline{Y}] \lambda_1 + \mathbb{E}[\underline{Y}Z_1] \lambda_2\} = \mathbb{E}[\underline{Y}] \max_{\lambda \in \mathbb{R}_+^2} \{\lambda_1 + \lambda_2\},$$

$$\text{s.t. } -\frac{1}{2}\lambda_1 - \frac{1}{2}\lambda_2 = -1.$$

This program has the constant value of the objective function on the entire optimization domain, which coincides with its argmin set,

$$(2.21) \quad \underline{\lambda} = \{\lambda \in \mathbb{R}_+^2 \mid \lambda_1 + \lambda_2 = 2\}. \square$$

The derivative (2.15) depends on particular optimal primal and dual solutions that are selected by a given perturbation (h_A, h_b) , unless both sets of solutions are singletons. I refer to parameter values A_P and b_P resulting in nonsingleton (primal or dual) solutions as *nonregular*.

Shapiro (1991) proposes a generalization of the delta method for the directionally differentiable functions of asymptotic normal estimators. Under Assumption 2, the sample analog of $\underline{v}(P)$ has a

solution (and thus is well defined) with probability approaching 1. Theorems 3.4 and 3.5 in [Shapiro \(1991\)](#) provide an asymptotic distribution of the sample support function at a given direction. It takes form

$$(2.22) \quad \frac{1}{\sqrt{n}}(\hat{v}_n - v(P)) \rightsquigarrow \min_{\theta \in \underline{\theta}(P)} \max_{\lambda \in \underline{\lambda}(P)} \{\lambda'(\mathbb{G}_A(P)\theta - \mathbb{G}_b(P))\},$$

where $\mathbb{G}(P) \triangleq (\mathbb{G}_A(P), \mathbb{G}_b(P))$ is the limiting zero-mean Gaussian process for $\mathbb{G}_n(P) \triangleq \frac{1}{\sqrt{n}} \sum_{i=1}^n (w_i - \mathbb{E}_P W)$ indexed by P . In general, the limit (2.22) is non-Gaussian since either of the two sets $\underline{\theta}$ and $\underline{\lambda}$ can be nonsingleton. (It does reduce to a Gaussian distribution when both sets are singletons.)

A large number of robust methods have been designed for inference on components of θ defined by moment inequalities (nonlinear in general). The dominant approach has been to test individual values of θ or their subvectors and then invert the tests (for example, [Andrews, 2001](#); [Andrews et al., 2019](#); [Chernozhukov et al., 2019, 2007](#); [Cox and Shi, 2019](#); [Kaido et al., 2019a](#)). The main advantage of the test-inversion approach is that fixing θ simplifies the asymptotic distribution of the test statistics and allows for valid inference under weak assumptions (in particular, one can allow for an empty true identified set as in [Andrews and Kwon \(2019\)](#)).

While having attractive statistical properties, the test-inversion methods can become computationally intractable in settings with high-dimensional parameters θ since they are based on grid-search and resampling methods.⁹ In contrast, delta-method inference using the support-function estimators for a fixed direction proposed in this paper remains a computationally tractable (frequentist) option in high-dimensional settings such as [Example 2](#) since it fully uses the linear-programming structure of the problem.

Inference using the sample support function for a fixed direction in the nonregular case faces three challenges: (i) nondifferentiability makes the standard bootstrap inconsistent (see [Fang and Santos, 2018](#)); (ii) discontinuity in the directional derivative (2.15) results in poor (uniform) approximation of (2.22) by the numerical bootstrap (see [Dümbgen, 1993](#); [Hong and Li, 2018](#)); (iii) the estimator is necessarily (asymptotically) biased (see [Hirano and Porter, 2012](#)). The regularized support-function estimator proposed in the next section addresses these challenges in a robust and computationally tractable way.

3. MAIN RESULTS

3.1. Bounds based on the regularized primal program

The main source of the inference complications, the lack of smoothness in the support function, resolves itself when the sets of primal, $\underline{\theta}$, and dual, $\underline{\lambda}$, optima are singletons (see equations (2.15) and (2.22)). Correspondingly, my proposal is to consider a regularized support function at e_1 that has a unique solution and that approximates the original support function from a known direction, either from above or from below.¹⁰ Since the direction of such a regularization bias is known,

⁹See further discussion of the computational issues in [Appendix C](#).

¹⁰If both primal and dual solutions of the regularized program are unique, all the directional derivatives of the regularized support function at e_1 given by (2.15) coincide and are given by formulas

$$(3.1) \quad \frac{\partial v(\mu_n, P)}{\partial (A_P)_{ij}} = \lambda_i(\mu_n, P) \underline{\theta}_j(\mu_n, P) \text{ and } \frac{\partial v(\mu_n, P)}{\partial (b_P)_i} = -\lambda_i(\mu_n, P),$$

where $\underline{\theta}(\mu_n, P)$ and $\underline{\lambda}(\mu_n, P)$ are, correspondingly, the primal and dual solutions of the regularized program with a regularization parameter μ_n .

the corresponding estimators of the regularized support function at e_1 admit standard one-sided delta-method and bootstrap inference based on the normal limiting distribution.

Without regularization both primal and dual solutions of (2.13) can be nonsingletons. To keep the problem analytically tractable, I focus on the case in which the dual set is a singleton by assumption, and I propose a regularization for the primal problem. This way I can explicitly consider only the bias from the primal regularization and develop the necessary bias correction. Appendix Section B.3 discusses the complementary case of dual regularization.

The following *regularized primal program* is strictly convex for any $\mu > 0$ and hence has a unique primal solution $\underline{\vartheta}(\mu, P)$ and approximates the optimal value of Program (2.1):

$$(3.2) \quad \underline{\vartheta}(\mu, P) = \min_{\theta \in \Theta(P)} \left\{ e_1' \theta + \mu \|\theta\|^2 \right\}.$$

The corresponding dual program has a unique solution $\underline{\lambda}(\mu, P)$ iff the set of constraints satisfies the linear independence constraint qualification (LICQ; see Shapiro (1991, p.178) and Wachsmuth (2013)),

ASSUMPTION 3 (LICQ) *The matrix of gradients of binding constraints has a full rank for any $\theta \in \Theta(P)$.*

REMARK 1 *There is a purely computational reason to ensure LICQ holds. If it is violated, Newton-type algorithms, which typically guarantee a (fast) quadratic rate of convergence to a stationary point, have a linear rate of convergence or do not converge at all (see, for example, Golishnikov and Izmailov (2006)).*

The set of DGPs that satisfy LICQ is not closed since a limit of a sequence of linearly independent matrices can be a reduced rank matrix. It means that depending on the size of the smallest singular value of the gradients of the set of binding constraints at any point may be arbitrary small while still satisfying LICQ. As a result, the dual solutions $\underline{\lambda}(\mu, P)$ may be arbitrary large implying large derivatives of $\underline{\vartheta}(\mu_n, P)$ and may require very high sample sizes for reasonable precision of the delta-method inference (see equation (A.32) in Lemma 7 in Appendix for details). In the next section I provide sufficient conditions for LICQ that explicitly ensure that the class of DGP under consideration is closed, so that we can uniformly control the quality of the delta-method inference within this class. Then I illustrate these conditions in the context of Example 3.

3.1.1. Testable sufficient conditions for uniqueness of dual solutions

LICQ is often considered hard to verify (see Kaido et al., 2019b). Indeed, a direct test of this assumption would face two problems: (i) the set $\Theta(P)$ is unknown but can only be estimated with an error; (ii) the set of binding inequalities at each θ is also unknown. Because of their multiple-testing nature, both problems would complicate the inference beyond the practical level. Moreover, the LICQ does not provide an explicit bound on the dual variables required for uniform-validity analysis.¹¹

¹¹Some authors including Hsieh et al. (2021) make a high level assumption about boundedness of the dual variables. Others (KMS and Andrews et al. (2019)) focus on test inversion and thus only require restriction on non-degenerate covariance matrix of the moment functions. Both KMS and Andrews et al. (2019) compute critical values of a test at a particular point θ which they then invert. (Kaido et al. (2019a)) additionally propose a method for computationally

I propose two sufficient conditions for LICQ in a form of bounds on the values of two auxiliary optimization programs. Values of these programs, in turn, explicitly characterize an upper bound on the dual variables, thus allowing for uniform asymptotic analysis. Specifically, within this class of DGP satisfying this assumptions, we can uniformly control the quality of the delta-method inference by establishing an explicit uniform bound on the relevant second order directional derivatives.

To introduce the sufficient conditions for LICQ, I use the following notation. For any $\mathcal{J} \subset \mathcal{J}^{ineq}$, let the matrix $\mathbb{J}^a = (e_{i_1}, \dots, e_{i_\ell})'$ correspond to $\mathcal{J}^a \triangleq \mathcal{J}^{eq} \cup \mathcal{J} = \{i_1, \dots, i_\ell\}$, a set of active constraints. Let $\eta_1(\cdot)$ be the smallest left-singular-value function of a matrix; that is, $\eta_1(A) \triangleq \sqrt{\min_u (u' A A' u / u' u)}$. The sufficient conditions can now be formulated as the following two assumptions on every submatrix $\mathbb{J}^a(A_P | b_P)$ that include all p equality constraints and either $d - p$ or $1 + (d - p)$ inequality constraints.

ASSUMPTION 4 *Measure P is such that for any combination \mathcal{J} consisting of all p equality constraints and $d - p$ inequality constraints, the corresponding submatrices of coefficients $\mathbb{J}^a(A_P | b_P)$ of the full set of constraints $(A_P | b_P)$ have full rank and singular values that are uniformly bounded from below by a positive number $\eta(P)$.*

Without loss of generality we can chose

$$(3.3) \quad \eta(P) \triangleq \min_{\substack{\mathcal{J} \subset \mathcal{J}^{ineq} \\ \text{s.t. } |\mathcal{J}| = d - p}} \eta_1(\mathbb{J}^a(A_P | b_P)) > 0.$$

ASSUMPTION 5 *Measure P is such that for any combination \mathcal{J} consisting of all p equality constraints and $d - p + 1$ inequality constraints cannot be simultaneously satisfied as equality at any point $\theta \in \Theta(P)$,*

$$(3.4) \quad s(P) \triangleq \min_{\substack{\mathcal{J} \subset \mathcal{J}^{ineq} \\ \text{s.t. } |\mathcal{J}| = d - p + 1 \\ \theta \in \Theta(P)}} \|\mathbb{J}^a(A_P \theta - b_P)\| > 0.$$

It is instructive to see what Assumptions 4–5 imply for our running example.

EXAMPLE 4 (Example 3 continued) In this setup, there are four moment inequality conditions defined in (2.17). They correspond to the following matrix of coefficients:

$$(3.5) \quad (A_P | b_P) = \left(\begin{array}{cc|cc} \frac{1}{2} & \mathbb{E}[Z_1 Z_2] & \mathbb{E}[\bar{Y} Z_1] & \\ -\frac{1}{2} & -\mathbb{E}[Z_1 Z_2] & -\mathbb{E}[\underline{Y} Z_1] & \\ 0 & \mathbb{E}[(1 - Z_1) Z_2] & \mathbb{E}[\bar{Y} (1 - Z_1)] & \\ 0 & -\mathbb{E}[(1 - Z_1) Z_2] & -\mathbb{E}[\underline{Y} (1 - Z_1)] & \end{array} \right)$$

efficient interpolation of the confidence sets based on test inversion.) KMS, for example, only need bounded dual variable for each of the bootstrap draws for their purposes. The local linear program that is used in KMS bootstrap has bounded dual almost surely under the aforementioned covariance constraints since its inequality constraints have random bootstrapped coefficients (see Remark 2 below). In contrast, I need stronger assumptions to be able to estimate the nuisance parameter, the argmin set $\underline{\theta}((\mu_n, P))$ and the dual solutions $\underline{\lambda}((\mu_n, P))$, that allows me to avoid the test inversion stage and achieve computational gains.

Here $p = 0$, so to check Assumption 4 one need to consider all subsets with two rows out of four, six combinations in total. For example, the submatrix with rows $\mathcal{J} = \{1, 2\}$ takes form

$$(3.6) \quad \mathbb{J}^a(A_P|b_P) = \left(\begin{array}{cc|c} \frac{1}{2} & \mathbb{E}[Z_1 Z_2] & \mathbb{E}[\bar{Y} Z_1] \\ -\frac{1}{2} & -\mathbb{E}[Z_1 Z_2] & -\mathbb{E}[\underline{Y} Z_1] \end{array} \right).$$

This matrix has a full row rank iff $\mathbb{E}[\bar{Y} Z_1] \neq \mathbb{E}[\underline{Y} Z_1]$ or $\Delta_1 > 0$. As result, we just verified that for $\mathcal{J} = \{1, 2\}$ we have $\eta_1(\mathbb{J}^a(A_P|b_P)) > 0$. Similarly, the matrix corresponding to rows $\mathcal{J} = \{3, 4\}$ has full rank iff both $\Delta_0 > 0$ and $\mathbb{E}[(1 - Z_1) Z_2] \neq 0$. Under those conditions submatrices with pairs of rows $\mathcal{J} \in \{\{1, 3\}, \{1, 4\}, \{2, 3\}, \{2, 4\}\}$ are also all full rank and thus their left singular values are all positive. To summarize, Assumption 4 is satisfied iff

$$(3.7) \quad \Delta_0 > 0, \Delta_1 > 0,$$

$$(3.8) \quad \mathbb{E}[(1 - Z_1) Z_2] \neq 0.$$

What do these conditions mean in practical terms? Inequalities (3.7) imply that the upper and lower bounds on Y are different from each other on average, conditional on Z_1 , while inequality (3.8) implies that the instrument Z_2 (with values in $\{0, 1\}$) is not perfectly correlated with Z_1 . Indeed, in the case in which the bounds \underline{Y} and \bar{Y} coincide with probability 1, one can replace the corresponding pair of moment inequalities with a single equality, as seen in Example 1. The degenerate case $\mathbb{E}[(1 - Z_1) Z_2] = 0$ is analogous to the multicollinearity problem in the usual linear regression setup.

Inequalities (3.7) imply that Assumption 5 is satisfied. To illustrate a violation of Assumption 5, suppose that we add one more moment condition corresponding to instrument Z_2 ,

$$\mathbb{E}[\underline{Y} Z_2] \leq \theta_1 \mathbb{E}[Z_1 Z_2] + \frac{1}{2} \theta_2.$$

Assumption 5 would be violated if, for example, this additional inequality was binding at the corner points of the original identified set, $\theta = (-\Delta_1 \mp 2\rho\Delta_0, \pm\Delta_0)$.

It is worth contrasting Assumptions 4-5 with Assumption 3: the latter only restricts the gradients of the submatrices of A_P at any point where the corresponding constraints are active. It turns out that instead of checking the active constraints at any given point, one can restrict the singular values of submatrices $\mathbb{J}^a(A_P|b_P)$. In this simple example, we can manually verify that LICQ holds under Assumptions 4-5 (the general proof is given in Lemma 2 in the Appendix). First, lets consider pairs of constraints with collinear gradients that potentially could violate LICQ. For example, the submatrix of A_P corresponding to rows $\mathcal{J} = \{1, 2\}$,

$$(3.9) \quad \mathbb{J}^a A_P = \left(\begin{array}{cc} \frac{1}{2} & \mathbb{E}[Z_1 Z_2] \\ -\frac{1}{2} & -\mathbb{E}[Z_1 Z_2] \end{array} \right),$$

is always a reduced rank matrix. Nevertheless, the corresponding two inequality constraints cannot be binding simultaneously by the Rouché-Capelli theorem (also known as Kronecker–Capelli theorem) since $\mathbb{J}^a(A_P|b_P)$ for them has rank equal to 2 which is larger than 1, the rank of $\mathbb{J}^a A_P$. Correspondingly, these two constraints do not violate LICQ at any point (since they do not intersect) despite having collinear gradients. Second, for constraint pairs like $\mathcal{J} = \{1, 3\}$ we have $\text{rk}(\mathbb{J}^a A_P) = \text{rk}(\mathbb{J}^a(A_P|b_P)) = 2$, and hence the corresponding corner point (intersection of these constraints) exist as a unique solution to $(\mathbb{J}^a A_P)\theta = \mathbb{J}^a b_P$. That corner point does not violate LICQ since

$\text{rk}(\mathbb{J}^a A_P) = d = 2$. By this logic, we can prove that all the points in this example have at most 2 binding constraints, all of which have linearly independent gradients. Third, faces of the polygon are always defined by a subset of the inequalities defining their corners. Hence the inequalities defining the faces are also linearly independent (a matrix of full rank has all submatrices of full rank) and any point on a face does not violate LICQ. Finally, we do not need to check the LICQ in the internal points of $\Theta(P)$ since there are no binding constraints at those points by definition. Thus we just verified in this example that Assumption 3 indeed holds under Assumptions 4-5. \square

REMARK 2 *Assumptions 4 and 5 are generically satisfied as long as the identified set is not a singleton. Indeed, if the matrix A_P and vector b_P were randomly drawn from some continuous distribution, a probability that any submatrix with exactly d rows of the joint matrix $(A_P|b_P)$ would be linearly dependent is zero. Correspondingly, it will have a positive smallest singular value. Since there are finitely many of such submatrices, there is always some positive uniform lower bound on the smallest singular value $\eta(P) > 0$. For the same reason, for a randomly drawn coefficients $(A_P|b_P)$, the probability of getting overidentifying conditions at any point would be zero, thus the inequality $s(P) > 0$ is also true generically.¹² As a result, in practice, unless one adds an overidentifying moment equality (case $p \geq d$, see also footnote 7 in Subsection 2.1), intentionally adds copies of the same inequality, or has a singleton identified set $\Theta(P)$, the delta-method inference proposed in this paper should work reasonably well.*

In contrast, for any given identified set $\Theta(P)$ defined by linear moment conditions there are multiple directions that are orthogonal to faces of $\Theta(P)$; the support function is not differentiable at those directions. As a result, the primal regularization is almost surely necessary at least for some directions to restore the Gaussian limiting distribution of the support function for those directions.

REMARK 3 (Sufficient but not necessary for LICQ) *Assumption 3 directly implies Assumption 5 (see Remark 7 in Appendix A.1). But Assumption 4 also restricts gradients of the constraints that can intersect outside of the identified set, in contrast to Assumption 3. So, for example, if there are two copies of a redundant inequality constraint (i.e. the one that is not active at any point in $\Theta(P)$), Assumption 4 is violated and Assumption 3 is not. In the light of Remark 2 such counterexamples, however, are rather artificial. Moreover, the truly redundant constraints in sufficiently large samples will not affect the estimator with probability approaching 1. So the gap between Assumptions 4-5 and Assumption 3 is not likely to be relevant in practice.*

REMARK 4 (Interpretation as LICQ imposed over the entire space) *Assumption 5 is formulated in form of restriction on a norm of a slack vector, $\|\mathbb{J}^a A_P \theta - \mathbb{J}^a b_P\|$ of any $d+1$ constraints at every point $\theta \in \Theta(P)$. If we require instead no overidentifying inequalities for all $\theta \in \mathbb{R}^d$, one can write this requirement in a form similar to Assumption 4,*

$$(3.10) \quad \min_{\mathcal{J} \subset \mathcal{J}^{ineq}} \eta_1(\mathbb{J}^a(A_P|b_P)) > 0.$$

s.t. $|\mathcal{J}| = d - p + 1$

The bound on the singular value (3.10) implies that $\text{rk}(\mathbb{J}^a(A_P|b_P)) = d+1$. Since for any combination of $d+1$ constraints $\text{rk}(\mathbb{J}^a A_P) \leq d$, by the aforementioned Rouché-Capelli theorem any combination

¹²This observation was used in a recent paper [Cho and Russell \(2023\)](#), who proposed adding a continuously distributed noise to b_p to guarantee LICQ at the expense of enlarging the identifying set. In contrast, in the present paper the goal is to study sufficient conditions under which such enlargement is not necessary, so the inference on the sharp identified set is possible.

of $d + 1$ linear inequalities/equalities cannot all be satisfied at once, i.e. $\|\mathbb{J}^a A_P \theta - \mathbb{J}^a b_P\| \geq 0$. The converse is also true by the same theorem.

One can notice that equation (3.10) (bound on singular values and rank of $(d+1) \times (d+1)$ matrices) implies equation (3.3) (bound on singular values and rank of its $d \times (d+1)$ submatrices) for some particular choice of bounds on the minimal singular value $\eta(P)$. As a result, (3.10) implies both Assumptions 4 and Assumption 5 as well as the LICQ for all $\theta \in \mathbb{R}^n$ (see Lemma 2 in Appendix). By the argument in Remark 7 in Appendix, LICQ for all $\theta \in \mathbb{R}^n$ would in turn imply (3.10). So (3.10) is necessary and sufficient condition for LICQ to hold for all $\theta \in \mathbb{R}^n$ (not just on the identified set). In practical terms then Assumptions 4-5 are satisfied if there are no overidentifying inequality and equality constraints for all $\theta \in \mathbb{R}^n$.

REMARK 5 (Implications for bounds on dual solutions) *In addition to implying uniqueness of the dual solutions (and Assumption 3), Assumptions 4 and 5 imply an explicit bound (uniform in all directions for the support function) on the dual solutions that depends only on $\eta(P)$ and C_Θ , a maximal norm of θ in Θ (see Lemma 4 in Appendix A.1). This explicit bound is particularly useful when establishing uniform bounds on the error of the asymptotic approximation of the proposed inference procedure (see Lemma 7 in Appendix).*

Although the minimal slack $s(P)$ parameter does not appear in the bound on the dual variables λ , it determines how large the sample size is needed for the estimated identified set to satisfy the Slater's condition (that is, Assumption 2) with a sufficiently high probability. This property is necessary to make sure that the sample support function at e_1 is well defined with probability approaching 1 in sufficiently large samples (see Lemma 10 in Appendix).

REMARK 6 (Testability of the assumptions) *An attractive feature of Assumption 4 is that it does not require a pretest of active constraints for every θ as Assumption 3 would. In fact, Assumption 4 can be rewritten as a minimal eigenvalue restriction on a single-large-block diagonal matrix with all possible matrices $(\mathbb{J}^a(A_P|b_P))(\mathbb{J}^a(A_P|b_P))'$ on the diagonal. Accordingly, the standard eigenvalue tests (analogous to the multicollinearity pretests) can be applied to it once appropriate adjustments for the high dimension of the parameter matrix are made. (I leave an explicit testing of this assumption for future research.)*

Assumption 5 rules out more than d binding inequality constraints at any point $\theta \in \Theta(P)$. There are some special empirical applications resulting in a singleton identified set $\Theta(P)$ in which such *overidentifying* inequality constraints can appear (for example, Gafarov (2014) and Shi et al. (2018) who considers a case of infinitely inequality conditions), which can result in multiplicity of the dual solutions. This multiplicity can be eliminated using a regularization of the dual program (2.12); that is, Assumption 5 can be relaxed within the framework proposed in this paper, but such an extension is left for future research. I briefly discuss this proposal in Appendix Section B.3.

3.1.2. Tighter bounds on the support function at a given direction

Clearly, for any positive μ , the value of the regularized program (3.2) $\underline{v}(\mu P)$ is larger than $\underline{v}(P)$ by at least $\mu \|\underline{\theta}(\mu, P)\|^2$. A tempting approach would be to correct this regularization bias by subtracting $\mu \|\underline{\theta}(\mu, P)\|^2$. This correction, unfortunately, makes the corrected value function $\underline{v}(\mu, P) - \mu \|\underline{\theta}(\mu, P)\|^2$ non-differentiable in the parameters A_P, b_P if the non-regularized value function $\underline{v}(P)$ itself is non-differentiable. So to preserve differentiability of the bias-corrected value

function, which is crucial for delta-method inference, one cannot use argmin of the actual regularized program for bias corrections. Instead, I suggest to tighten the bounds on $\underline{v}(P)$ using

$$(3.11) \quad \underline{v}^{in}(\mu, \kappa, P) \triangleq \underline{v}(\mu, P) - \mu \|\underline{\theta}(\kappa, P)\|^2,$$

$$(3.12) \quad \underline{v}^{out}(\mu, P) \triangleq \underline{v}(\mu, P) - \mu \|\theta^*\|^2,$$

where $\underline{\theta}(\kappa, P)$ is an argmin of the regularized program (3.2) with a larger tuning parameter κ instead of μ , θ^* is any point in $\underline{\theta}(P)$. As μ and κ shrink to zero, these bounds continuously shrink toward $\underline{v}(P)$ from above and below, correspondingly.

THEOREM 1 *For any DGP parameterized by P satisfying Assumption 1 and any $\kappa \geq \mu \geq 0$, the following bounds hold:*

$$(3.13) \quad \underline{v}^{out}(\mu, P) \leq \underline{v}(P) \leq \underline{v}^{in}(\mu, \kappa, P).$$

If, in addition, P satisfies Assumptions 4–5, then there exist $\bar{\mu}(P) > 0$ such that $\underline{v}^{in}(\mu, \kappa, P) = \underline{v}(P)$ for any $\mu < \kappa < \bar{\mu}(P)$. Further, if $\underline{\theta}(P)$ is a singleton, then $\underline{v}^{out}(\mu, P) = \underline{v}(P)$ for any $\mu < \bar{\mu}(P)$.

PROOF: See Appendix Section A.3.

Q.E.D.

The bounds (3.13) are valid even in the general nonlinear case, which is beyond the scope of this paper. For linear programs, the stronger result holds: there exists a finite cutoff value $\bar{\mu}(P) > 0$ at which the upper bound $\underline{v}^{in}(\mu, \kappa, P)$ coincides with $\underline{v}(P)$. As a result, the sample analog of $\underline{v}^{in}(\mu, \kappa, P)$ with appropriately chosen shrinking sequences μ_n, κ_n will be asymptotically unbiased for $\underline{v}(P)$ given any (fixed) DGP P . Moreover, in the regular case, a sample analog of the outer bound $\underline{v}^{out}(\mu, P)$ also has this property. This is important since it shows that in the regular case, sample analogs of either of the two bounds (asymptotically) coincide with the usual support function estimator that attains the semiparametric efficiency bound for convex moment inequality models (see [Kaido and Santos, 2014](#)).

It is tempting to find ways to pick the ideal tuning parameters $\mu_n = \kappa_n = \bar{\mu}(P)$ that would result in exact regularization. Even though for a given DGP P , characterized by (A_P, b_P) , this cut-off $\bar{\mu}(P)$ is well defined, it can discontinuously change as a result of a small change in the DGP. Moreover, $\bar{\mu}$ can take values that are arbitrarily close to zero. So it cannot be uniformly consistently estimated or non-trivially bounded from below over the relevant class of DGP \mathcal{P} (the trivial zero bound corresponds to absence of regularization). As a practical solution in the subsequent sections, I suggest using two positive shrinking sequences, μ_n and κ_n , such that in sufficiently large samples for any fixed DGP they become smaller than any fixed $\bar{\mu}(P)$.¹³ As a result, for a fixed DGP P the value $\underline{v}^{in}(\mu_n, \kappa_n, P)$ will coincide with $\underline{v}(P)$ with probability 1 as sample size grows. In the regular case, the same is true for $\underline{v}^{out}(\mu_n, P)$.

In small samples, the expressions in (3.13) provide valid conservative bounds from above and below for $\underline{v}(P)$ that are useful for uniform one-sided inference (coverage probability in this case can be higher than the nominal level). The gap in small samples between $\underline{v}(P)$ and $\underline{v}^{out}(\mu_n, P)$ is at least $\mu_n(\|\theta^*\|^2 - \|\underline{\theta}(\mu_n, P)\|^2) \geq 0$; for $\underline{v}^{in}(\mu_n, \kappa_n, P)$ the gap is at most $\mu_n(\|\underline{\theta}(\mu_n, P)\|^2 - \|\underline{\theta}(\kappa_n, P)\|^2) \leq 0$.

¹³There is an analogy here between replacing μ_n with larger κ_n (which equivalent to using smaller sample size) in the present context and m out of n subsampling procedures that are sometimes required in non-regular estimation problems for point-wise consistency of inference. This observation was pointed out in a private conversation with professor Jinyong Hahn.

In the non-regular case, the gap for $\underline{v}^{out}(\mu_n, P)$ is shrinking to 0 at rate μ_n which still results in a consistent estimators of $\underline{v}(P)$ at rate μ_n that can be slightly slower than the regular parametric rate $1/\sqrt{n}$ (for example, $\sqrt{\ln n}/\sqrt{n}$). As a result, the corresponding confidence sets cover the projections of the sharp identified set, not an enlargement of it (in contrast to other studies, including [Cho and Russell \(2023\)](#), who do not use consistent estimators of the sharp bounds on the identified set and study an enlarged identified set instead). The gap for $\underline{v}^{in}(\mu_n, \kappa_n, P)$ becomes exactly equal to zero for some sufficiently large n in both regular and non-regular cases, resulting in exact corresponding point-wise one-side confidence sets.

Theorem 1 is crucial for valid one-sided inference on $\underline{v}(P)$. Under the assumptions of the theorem, the value of Program (3.2) for $\mu > 0$ is differentiable in $(A_P|b_P)$ (see Lemma 7 in Appendix) so that the sample analog (defined in Subsection 3.2) has an asymptotic Gaussian distribution. As a result, the analog estimators of \underline{v}^{out} and \underline{v}^{in} are half-median unbiased estimators for \underline{v} . In other words, they can be biased, but the direction of the bias is known, which allows for valid one-sided CSs (compare [Chernozhukov et al., 2013](#); [Hirano and Porter, 2012](#)).

3.2. Large-sample results for the estimated regularized support function

Consider the analog estimator of $\underline{v}(\mu_n, P)$ for some sequence μ_n ,

$$(3.14) \quad \underline{v}(\mu_n, \mathbb{P}_n) \triangleq \min_{\theta \in \Theta(\mathbb{P}_n)} \left\{ e_1' \theta + \mu_n \|\theta\|^2 \right\}.$$

As we will see in this subsection, this estimator admits an asymptotic linear (Bahadur-Kiefer) expansion with an explicit approximation error, which determines the precision of asymptotic normal and bootstrap inference. This allows me to establish validity of the corresponding inference methods *uniformly* over a reasonable class of DGP. The uniform asymptotic validity is crucial for small-sample performance of inference methods in the presence of possible discontinuous changes of the asymptotic distribution of estimators (for example, $\underline{v}(\mu_n, \mathbb{P}_n)$) with respect to the DGP parameters (for example, P).

Throughout the paper, I focus on the i.i.d. setup.

ASSUMPTION 6 $\{w_i \in \mathbb{R}^{k \times (d+1)} \mid i = 1, \dots, n\}$ is an i.i.d. sample of W .

There is a straightforward way to extend the analysis to the case of dependent data as long as a CLT remains valid.

3.2.1. A class of DGPs under consideration

I consider the class of all measures $\mathcal{P} = \mathcal{P}(\eta, \underline{s}, \varepsilon, \bar{M})$ that satisfy Assumptions 1, 4, and 5 together with Assumption 7 with some positive constants η , \underline{s} , ε , and \bar{M} .

ASSUMPTION 7 There exist $\varepsilon > 0$ and $\bar{M} < \infty$ such that

$$(3.15) \quad \mathbb{E}_P \|W\|^{2+\varepsilon} < \bar{M}.$$

To summarize, every $P \in \mathcal{P}$ satisfies $\Theta(P) \neq \emptyset$, $\eta(P) > \eta$, $s(P) > \underline{s}$, and $\mathbb{E}_P \|W\|^{2+\varepsilon} < \bar{M}^{2+\varepsilon}$.

The class \mathcal{P} includes DGPs (parameterized by measure \bar{P}) with multiplicity of primal solutions of (2.1) but assumes uniqueness of the dual solutions. This property is necessary to show that

Program (3.14) has a non-empty sample argmin and a unique vector of Lagrange multipliers in large enough samples with probability approaching 1 uniformly in $P \in \mathcal{P}$ as $n \rightarrow \infty$ (see Lemma 10 in Appendix). As discussed in Section 3.1.1, one can study a case with multiple dual solutions and a unique primal solution in a similar way (see Appendix Section B.3).

3.2.2. A higher-order envelope theorem and a strong approximation of the regularized support function

A Bahadur-Kiefer expansion of estimator defined in Program (3.14) can be established using the envelope theorem (2.15). As discussed earlier in Subsection 2.3, this theorem typically holds for the sample support function at e_1 . However, this theorem does not provide bounds on the higher-order directional derivatives. So without additional assumptions, the directional delta method of Shapiro (1991) does not provide means to evaluate the error of the asymptotic approximation of the sample support function at e_1 by (2.22). Since the limiting distribution changes discontinuously with the DGP P , one should expect poor performance of the asymptotic methods based on the limit (2.22) when parameter P is close to a discontinuity point. In other words, inference procedures that estimate (2.22) directly (for example, subsampling or directional bootstrap) will inevitably be only point-wise valid and can have poor finite-sample performance in such cases. (It is an implication of the impossibility theorem of Hirano and Porter (2012) for a directionally differentiable functional.)

In contrast, the regularized support function at e_1 admits a stronger version of the envelope theorem that provides a bound on the error of the linear approximation uniformly over $P \in \mathcal{P}$. I developed this novel bound using a *second-order directional Taylor expansion* of a system of generalized inequalities that define the optimal solutions to the regularized program (see Lemmas 6 and 7 in Appendix). Using this result, the asymptotic linear (Bahadur-Kiefer) representation of the value function in Program (3.14) follows almost immediately (see Lemma 11 in Appendix):

$$(3.16) \quad \sqrt{n}(\underline{v}(\mu_n, \mathbb{P}_n) - \underline{v}(\mu_n, P)) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \underline{\lambda}(\mu_n, P)' g(w_i, \underline{\theta}(\mu_n, P)) + O_{\mathcal{P}}\left(\frac{1}{\mu_n \sqrt{n}}\right).$$

The first term on the right-hand side of this representation is a scaled sample average of a zero-mean random variable, which admits a uniform Gaussian approximation. (Indeed, the binding constraints have zero mean at $\underline{\theta}$, while the nonbinding constraints are multiplied by zero dual variables $\underline{\lambda}$.) The residual term $O_{\mathcal{P}}(1)$ denotes a uniformly tight sequence of a random process indexed by $P \in \mathcal{P}$. Analogously, I denote any random sequence as $o_{\mathcal{P}}(1)$ if $\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P(\|\zeta_n(P)\| \geq \epsilon) = 0$.

The Bahadur-Kiefer representation (3.16) is important for three reasons. First, it suggests a coupling of $\underline{v}(\mu_n, \mathbb{P}_n)$ with a Gaussian process (through the Yurinsky theorem). Second, it implies uniform validity of the optimization-free score bootstrap, which is particularly computationally convenient. Third, it suggests an analog estimator of the asymptotic variance of the regularized support-function estimator,

$$(3.17) \quad \underline{\sigma}^2(\mu_n, \mathbb{P}_n) = \frac{1}{n} \sum_{i=1}^n (\underline{\lambda}'(\mu_n, \mathbb{P}_n) g(w_i, \underline{\theta}(\mu_n, \mathbb{P}_n)))^2,$$

where $\underline{\theta}(\mu_n, \mathbb{P}_n)$ and $\underline{\lambda}(\mu_n, \mathbb{P}_n)$ are, respectively, the optimum and the vector of Lagrange multipliers of (3.14), which are provided by common constraint-optimization software packages.

These implications are summarized in the following theorem.

THEOREM 2 Consider any sequence μ_n such that $\mu_n \rightarrow 0$ and $\mu_n \sqrt{n} \rightarrow \infty$. Then with probability approaching 1 uniformly in $P \in \mathcal{P}$,

$$(3.18) \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \pi(\sqrt{n}(\underline{v}(\mu_n, \mathbb{P}_n) - \underline{v}(\mu_n, P)), N(0, \underline{\sigma}^2(\mu_n, P))) = 0,$$

$$(3.19) \quad \underline{\theta}(\mu_n, \mathbb{P}_n) = \underline{\theta}(\mu_n, P) + O_{\mathcal{P}}\left(\frac{1}{\mu_n \sqrt{n}}\right),$$

$$(3.20) \quad \underline{\lambda}(\mu_n, \mathbb{P}_n) = \underline{\lambda}(\mu_n, P) + O_{\mathcal{P}}\left(\frac{1}{\sqrt{n}}\right),$$

$$(3.21) \quad \underline{\sigma}(\mu_n, \mathbb{P}_n) = \underline{\sigma}(\mu_n, P) + o_{\mathcal{P}}(1).$$

The function $\pi(\cdot, \cdot)$ is the Levy-Prohorov metric, which metricizes the weak topology of probability measures (see *van der Vaart and Wellner (1996)*).

PROOF: The strong-approximation result (3.18) is based on the uniform bound on the higher-order directional derivatives (implied by Assumptions 1, 4, and 5) and the generalization of the *Yurinskii (1978)* coupling proposed in *van der Vaart and Wellner (1996, Proposition A.5.2 on p. 457)* (it is implied by Assumptions 6 and 7). See Appendix A.4 for details. Q.E.D.

The coupling with a Gaussian random process (3.18) is a *strong-approximation* result, which can be understood using a geometric interpretation. The distance between the difference $\sqrt{n}(\underline{v}(\mu_n, \mathbb{P}_n) - \underline{v}(\mu_n, P))$ and some sequence of zero-mean Gaussian r.v.s $N(0, \underline{\sigma}^2(\mu_n, P))$ converges to zero with the same *uniform* rate for all DGP $P \in \mathcal{P}$. This property is stronger than the conventional CLT-type results since it does not require the existence of a limiting distribution for the approximating Gaussian variables. (For example, the limit can vary depending on the choice of the tuning sequence μ_n or the DGP P , which itself can vary in the asymptotic experiment.) In practical terms, for the regularized support function at e_1 , it means the finite-sample performance of such a *uniform* asymptotic Gaussian approximation will work well even if the DGP parameter P results in a nondifferentiable support function in the absence of regularization (the case in which the usual sample support function at e_1 is no longer asymptotically Gaussian). The implementation of the Gaussian approximation can be based on the delta method or the score bootstrap using the sample analogs of the optimal solutions. In this paper, however, I focus on the delta-method implementation.

3.3. Uniformly valid inference

One can use Theorems 1 and 2 to construct uniformly valid confidence bands (one-sided CSs) for $\underline{v}(P)$. The corresponding generic algorithm takes the following form:

- Step 1. Compute the regularized sample support function $\underline{v}(\mu_n, \mathbb{P}_n)$ at e_1 defined in (3.14).
- Step 2. Compute the standard error using (3.17).
- Step 3. Compute a bias adjustment using sample analogs of either $\mu_n \|\underline{\theta}(\kappa_n, P)\|^2$ (for an upper confidence band) or the norm of some point in the argmin set $\mu_n \|\theta^*\|^2$ (for a lower confidence band).

The following subsections explain specific implementations of this algorithm for uniformly valid CSs for a projection of the identified set on a single coordinate or multiple coordinates and for the argmin set of a linear program with estimated coefficients.

3.3.1. Application to confidence sets for a scalar projection of the identified set

Let's revisit one of the primary objects of interest in the moment-inequality models: CSs on projections of the identified set, $\mathcal{S}(P) \triangleq [\underline{v}(P), \bar{v}(P)]$ (see Examples 1 and 2).

Theorems 1 and 2 suggest that the outer-bound estimator for the minimal value,

$$(3.22) \quad \underline{v}^{out}(\mu_n, \mathbb{P}_n) \triangleq \underline{v}(\mu_n, \mathbb{P}_n) - \mu_n \|\theta^*(\mathbb{P}_n)\|^2,$$

is asymptotically unbiased (in the regular case) or biased downward (in the nonregular case). (Some downward bias is acceptable since the purpose of CSs is to cover the lowest point $\underline{v}(P)$ from below.) To achieve this property, the estimator $\theta^*(\mathbb{P}_n)$ should have a norm that is not smaller than the minimal norm in $\underline{\theta}(P)$ with probability approaching 1. We can use an estimator $\theta^*(\mathbb{P}_n)$ that converges to the point with the coordinates

$$(3.23) \quad \theta_i^*(P) \triangleq \max\{\theta_i^+(P), \theta_i^-(P)\},$$

where

$$(3.24) \quad \theta_i^\pm(P) \triangleq \left| \min_{\theta \in \Theta(P), \theta_1 \leq \underline{v}(P) + \mu_n} \{\pm \theta_i\} \right|.$$

By definition, $\|\theta^*\| \geq \|\theta\|$ for any $\theta \in \underline{\theta}(P)$. (See the proof of consistency in Lemma 12 in the Appendix.)

This bound on $\|\theta^*\|$ has two attractive properties. First, it can be (uniformly) consistently estimated using only $2k$ linear programs, which can be easily computed even in models with a very large-dimension k and a large number of inequalities using interior-point numerical optimization methods. Second, in the regular case, in which $\underline{\theta}(P)$ is a singleton, this estimator is asymptotically unbiased since it converges to $\|\theta^*\| = \|\underline{\theta}(P)\|$. So for any such fixed P , by Theorem 1 we have $\underline{v}^{out}(\mu_n, P) = \underline{v}(P)$ for sufficiently small μ_n ; that is, it is possible to construct CIs with correct (nonconservative) coverage in the regular case.

Using the bias-corrected estimator $\underline{v}^{out}(\mu_n, \mathbb{P}_n)$ and its analog for the upper bound, $\bar{v}^{out}(\mu_n, \mathbb{P}_n)$, I construct the following delta-method CSs:

$$(3.25) \quad \begin{cases} \text{CB}_{\alpha, n, \mathcal{P}} &= [\underline{v}^{out}(\mu_n, \mathbb{P}_n) - z_{1-\alpha} n^{-1/2} \hat{\underline{\sigma}}_n^{reg}, \infty), \\ \text{CI}_{\alpha, n, \mathcal{P}}^{\theta_1} &= [\underline{v}^{out}(\mu_n, \kappa_n, \mathbb{P}_n) - z_{1-\alpha} n^{-1/2} \hat{\underline{\sigma}}_n^{reg}; \bar{v}^{out}(\mu_n, \kappa_n, \mathbb{P}_n) + z_{1-\alpha} n^{-1/2} \hat{\bar{\sigma}}_n^{reg}], \\ \text{CI}_{\alpha, n, \mathcal{P}}^S &= \text{CI}_{\alpha/2, n, \mathcal{P}}^{\theta_1}. \end{cases}$$

Here, $z_{1-\alpha}$ is $1 - \alpha$ quantiles of the standard Gaussian distribution and

$$(3.26) \quad \hat{\underline{\sigma}}_n^{reg} \triangleq \max\{\underline{\sigma}(\mu_n, \mathbb{P}_n), \sigma_0\},$$

$$(3.27) \quad \hat{\bar{\sigma}}_n^{reg} \triangleq \max\{\bar{\sigma}(\mu_n; \mathbb{P}_n), \sigma_0\},$$

for some small positive number σ_0 . $\text{CB}_{\alpha, n, \mathcal{P}}$ is a one-sided CB for $\underline{v}(P)$ (and for θ_1 as well), $\text{CI}_{\alpha, n, \mathcal{P}}^{\theta_1}$ is a two-sided CI that covers any θ_1 in the identified set, and $\text{CI}_{\alpha, n, \mathcal{P}}^S$ is a two-sided CI (based on the Bonferroni inequality) that covers the entire marginal identified set $\mathcal{S}(P)$. The tuning parameter σ_0 is introduced to guarantee the nominal coverage in the cases in which only nonstochastic constraints are binding at the optimal solutions corresponding to the support functions; otherwise, this degeneracy can lead to superconsistent support-function estimators with a non-Gaussian limiting distribution. One can set $\sigma_0 = 0$ if this concern is not appropriate in a particular application.

As before, let \mathcal{P} contain all measures P that satisfy Assumptions 1 and 4–7 with some uniform positive constants $\underline{\eta}$, \underline{s} , ε , and \bar{M} .

THEOREM 3 *Suppose that Assumption 6 holds. In addition, suppose that $0 < \alpha < 1/2$, $\mu_n \rightarrow 0$, and $\mu_n \sqrt{n} \rightarrow \infty$. Then the following results hold:*

$$\begin{aligned} \liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{P}} P(\mathcal{S}(P) \subset CB_{\alpha, n, \mathcal{P}}) &= \liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{P}} \inf_{\theta \in \Theta(P)} P(\theta_1 \in CB_{\alpha, n, \mathcal{P}}) \geq 1 - \alpha, \\ \liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{P}} P(\mathcal{S}(P) \subset CI_{\alpha, n, \mathcal{P}}^S) &\geq 1 - \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{P}} \inf_{\theta \in \Theta(P)} P(\theta_1 \in CI_{\alpha, n, \mathcal{P}}^S) \geq 1 - \alpha. \end{aligned}$$

PROOF: See Appendix A.4.

Q.E.D.

Note that the worst-case (that is, the smallest) asymptotic coverage probability of $CB_{\alpha, n, \mathcal{P}}$ is exactly equal to $1 - \alpha$ for a fixed regular DGP such that $\underline{\theta}(P)$ is a singleton and $\lim_{n \rightarrow \infty} \underline{\sigma}^2(\mu_n, P) > \sigma_0 \geq 0$. For a nonregular DGP (or sequence of DGPs), the asymptotic coverage can only be larger than $1 - \alpha$. How conservative is $CB_{\alpha, n, \mathcal{P}}$ in the nonregular case? We can evaluate it by comparing its average length with that of a confidence bound that has exact point-wise asymptotic coverage probability in a Monte Carlo study (see Section 4). Theorem 4 in Appendix Section B.1 provides an alternative bias correction that remains non-conservative in the non-regular case and results in shorter point-wise confidence bounds. This approach is based on using an estimator of the inner bound $\underline{v}^{in}(\mu, \kappa, P)$.

I conclude this section with a brief discussion of the tuning parameters. The theory of optimal choice of tuning parameter is beyond the scope of this paper. The following considerations, however, can provide some guidance for the optimal choice. Theorem 1 suggests that the tuning parameters should be smaller than $\bar{\mu}(P)$ to avoid the bias in the first-order asymptotic distribution. This choice is infeasible since $\bar{\mu}(P)$ is unknown, so one has to let κ_n and μ_n go to zero. The optimal rates of κ_n and μ_n should balance the higher-order variance and the worst-case bias. A specific choice of the tuning parameters is discussed in Section 4.

3.3.2. Joint confidence sets for general subvectors

It is trivial to extend the analysis to

$$\underline{v}(P; a) = \min_{\theta \in \Theta(P)} a' \theta$$

for any $a \in R^d$ with $\|a\| = 1$. Indeed, Assumptions 1–7 are invariant with respect to orthogonal transformations of the coordinates; that is, they are satisfied for the following program (with $\tilde{\theta} = U' \theta$, $\tilde{A}_P = A_P U$, and $a' = e_1' U$ for any orthogonal matrix U):

$$(3.28) \quad \underline{v}(P; a) = \min_{\tilde{\theta} \in \mathbb{R}^d} e_1' \tilde{\theta}$$

$$(3.29) \quad \text{s.t.} \quad \begin{cases} e_j' \tilde{A}_P \tilde{\theta} = e_j' b_P, & j \in \mathcal{J}^{eq}, \\ e_j' \tilde{A}_P \tilde{\theta} \leq e_j' b_P, & j \in \mathcal{J}^{ineq}. \end{cases}$$

One can think about \tilde{A}_P as a coefficient matrix under a different measure \tilde{P} , $A_{\tilde{P}}$. The set of measures \mathcal{P} from Section 3.3.1 includes \tilde{P} corresponding to all orthogonal transformations of A_P .

The identified set $\Theta(P)$ is convex, so any projection of it can be characterized using support functions for a corresponding direction. One can construct a joint CS for $\Theta(P)$ as follows. For any set of directions $\mathcal{A} \subset R^d$, take

$$CS_{\alpha, n}^{\mathcal{A}} = \{\theta | a \in \mathcal{A}, a' \theta \leq -\underline{v}^{out}(\mu_n, \mathbb{P}_n; -a) + c_{1-\alpha} n^{-1/2} \max\{\underline{\sigma}(\mu_n, \mathbb{P}_n; -a), \sigma_0\}\},$$

where $c_{1-\alpha}$ is $1 - \alpha$ quantiles of the maximum of the corresponding asymptotic Gaussian variables (also known as *supt* statistics) that can be estimated using multiplier bootstrap enabled by the asymptotic linear representation, (3.16).¹⁴ One can also use the Bonferroni inequality-based standard Gaussian critical value $c_{1-\alpha} = z_{1-\alpha/|\mathcal{A}|}$.

By appropriately choosing the set of directions \mathcal{A} , we can construct joint CSs for projections of $\Theta(P)$ on any subvectors θ . If \mathcal{A} has finitely many elements, $\text{CS}_{\alpha,n}^{\mathcal{A}}$ is a polygon. So we can plot it directly without performing test inversion as in the one-dimensional case. The confidence set $\text{CI}_{\alpha,n,\mathcal{P}}^{\mathcal{S}}$ is a particular case of $\text{CS}_{\alpha,n}^{\mathcal{A}}$ corresponding to $\mathcal{A} = \{e_1, -e_1\}$ and the Bonferroni estimate of $c_{1-\alpha}$. The following example provides another interesting set of directions.

EXAMPLE 5 (A polygonal joint confidence set) It seems natural to construct a joint CS for θ based on directions corresponding to the normal vectors of the moment conditions. For simplicity, assume that $p = 0$. The original system (2.2) may have some inequalities that are slack for any point $\theta \in \Theta(P)$. We can characterize the identified set $\Theta(P)$ as the solution to a tight system of inequalities

$$(3.30) \quad e'_j A_P \theta \leq \underline{b}_j, j \in \mathcal{J}^{ineq},$$

where

$$(3.31) \quad \underline{b}_j = \max_{(\vartheta, \theta) \in \mathbb{R}^{d+1}} \vartheta$$

$$(3.32) \quad \text{s.t.} \quad \begin{cases} \vartheta & = e'_j A_P \theta, \\ e'_\ell A_P \theta & \leq e'_\ell b_P, \ell \in \mathcal{J}^{ineq}. \end{cases}$$

Every inequality in System (3.30) is active at least at one point in $\Theta(P)$ (any point in the argmax of (3.31)). Programs (3.31) meet Assumptions 1 and 4–7, and so the outer estimators \hat{b}_j^{out} are half-median unbiased with corresponding standard error estimators $\hat{\sigma}_j$. Then the following polyhedron CS will cover any point $\theta \in \Theta(P)$ with asymptotic probability of at least $1 - \alpha$ uniformly over $P \in \mathcal{P}$,

$$\text{CS}_{\alpha,n}^{\mathcal{N}} = \{\theta | e'_j \hat{A}_P \theta \leq \hat{b}_j^{\text{out}} + z_{1-\frac{\alpha}{k}} n^{-1/2} \max\{\hat{\sigma}_j, \sigma_0\}, j \in \mathcal{J}^{ineq}\},$$

where $z_{1-\frac{\alpha}{k}}$ is the critical value of the standard normal r.v. and k is the number of (in)equality restrictions. The generalization to the case $p \neq 0$ is straightforward. \square

4. MONTE CARLO EXPERIMENTS

4.1. Overview

In this Monte Carlo study our goal is to evaluate how the confidence bound's length and the corresponding coverage probability depends on a number of key factors: (i) how close gradients of the relevant moment inequality to the non-regular case (i.e. being collinear with the vector e_1); (ii) how tuning parameter choice affects the conservativeness of confidence bounds; (iii) how the dimension of the problem and the number of inequalities affect length, coverage and computational time for the proposed methods. The last exercise in the list also involves comparison with the AS projection implemented using KMS EM computational algorithm.

¹⁴I leave full analysis of multiplier bootstrap procedure in this setup for future work; see additional discussion in Appendix Section C.3.

4.2. Proximity to non-regular case

To illustrate advantages of uniformly coverage when compared to the point-wise coverage, we can study a simple design with $k = 4$ moment inequalities where angle ω between the gradient of one of the faces of the identified set and vector e_1 varies continuously. The expectation of the moment inequalities can be parametrized as follows

$$\mathbb{E}_P W = \mathbb{E}_P(A_P | b_P) = \left(\begin{array}{cc|c} -\cos(\frac{\omega\pi}{180}) & -\sin(\frac{\omega\pi}{180}) & \cos(\frac{\omega\pi}{180}) + \sin(\frac{\omega\pi}{180}) \\ \cos(\frac{\omega\pi}{180}) & \sin(\frac{\omega\pi}{180}) & \cos(\frac{\omega\pi}{180}) + \sin(\frac{\omega\pi}{180}) \\ 0 & -1 & 1 \\ 0 & 1 & 1 \end{array} \right).$$

The shape of this set is a parallelogram analogous to the one in Figure 1 in with $\rho \geq 0$. The parameter $\omega \in [0^\circ, 36^\circ]$ defines the angle between the normal vectors of the rear sides of the parallelogram and the horizontal axis. The value $\omega = 0^\circ$ corresponds to a square-shaped identified set, also referred to as a *non-regular case* since the sides are orthogonal to the objective function gradient. All other values are referred to as *regular*. In the vicinity of $\omega = 0^\circ$ point-wise valid $\text{CB}_{\alpha,n}$ may have coverage probability below the nominal level because it lacks uniform validity. The expectation $\mathbb{E}_P W$ is parameterized to guarantee $\underline{\theta}_1 = \underline{\theta}_2 = -1$ and $\bar{\theta}_1 = \bar{\theta}_2 = 1$ for all values of ω . The components of W_i are independent Gaussian random variables with variance $s_2^2 = 0.01$. For each value of ω , I compute the frequency of coverage and the excess average length over the identified set for $\text{CB}_{\alpha,n}$ and $\text{CB}_{\alpha,n,\mathcal{P}}$ based on sample sizes $n \in \{100, 10000\}$. The number of MC simulations is 1,000 for every combination of n and ω . The focus is on the nominal coverage probability $\alpha = 0.95$.

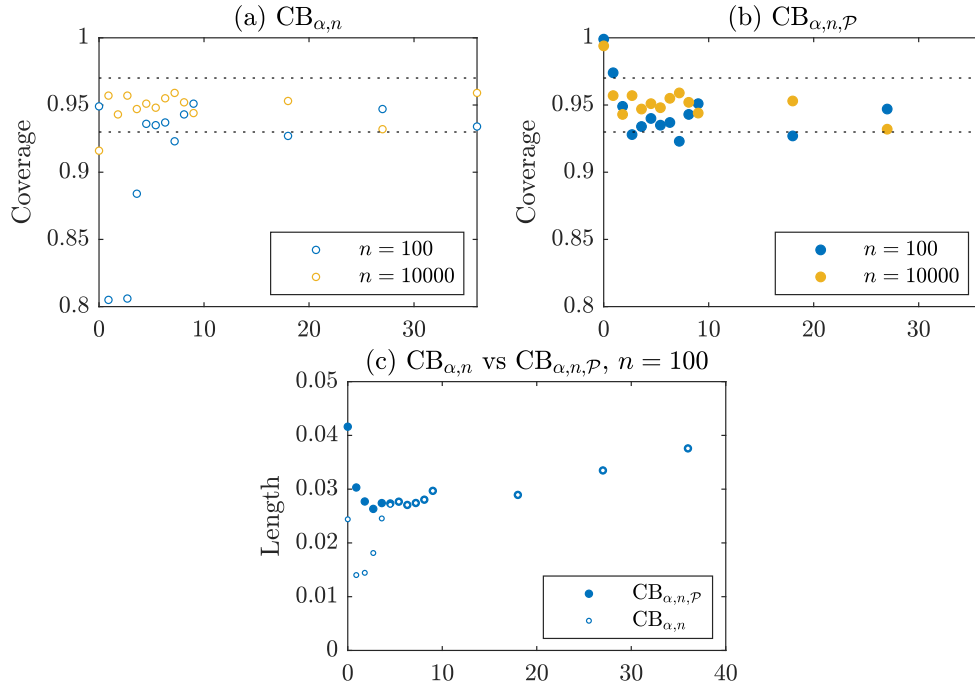
As the main tuning parameter choice, I use $\mu_n = \hat{\mu}_0 \sqrt{n^{-1} \ln \ln n}$ and $\kappa_n = \hat{\mu}_0 \sqrt{n^{-1} \ln n}$, where $\hat{\mu}_0 = \sqrt{n^{-1} \sum_{i=1}^n \lambda'(0, \mathbb{P}_n) w_i e_1}$. While a theory of optimal choice of $\hat{\mu}_0$ is beyond the scope of this paper, this particular choice has some advantages: (i) $\hat{\mu}_0$ depends only on the behavior of the relevant moment inequalities as selected by non-zero components of $\lambda'(0, \mathbb{P}_n)$; (ii) it does not depend on the dimension of the parameters θ , since only the first column of w_i is involved; (iii) it has the same scale as the standard deviation of the relevant components of w_i which justifies our perturbation analysis (the impact of the regularization term $\mu_n \|\theta\|^2$ is larger than the sample variation, $\sim \hat{\mu}_0 n^{-1/2}$). As a result, this choice resulted in a good alignment of the theoretical predictions with the simulations as could be seen below.

To appreciate the importance of the uniformly valid confidence bounds, we start our study with the point-wise valid confidence bounds $\text{CB}_{\alpha,n}$. Figure 2 panel (a) shows the corresponding coverage frequency for a small sample size $n = 100$ and a large sample with $n = 10,000$. For values of ω that are far from 0, both sample sizes have observed frequency of covering the correct bound of the identified set for θ_1 close to the nominal coverage probability of $\alpha = 0.95$. The same is true for the non-regular case with $\omega = 0$. However in the vicinity of the non-regular case, $\omega \in (0, 4^\circ]$, large sample sizes are necessary to achieve satisfactory coverage frequency. In contrast, the coverage frequency for the corresponding sample sizes for the uniformly-valid confidence bounds $\text{CB}_{\alpha,n,\mathcal{P}}$ given on Figure 2 panel (b) is uniformly at least as large as α for all values of ω . One can see that the only point on Figure 2 panel (b) with coverage is significantly higher than $\alpha = 0.95$ for $n = 10,000$ is $\omega = 0$, that is for most designs (or equivalently, directions of the support function) the uniformly valid confidence bounds $\text{CB}_{\alpha,n}$ have exact coverage.

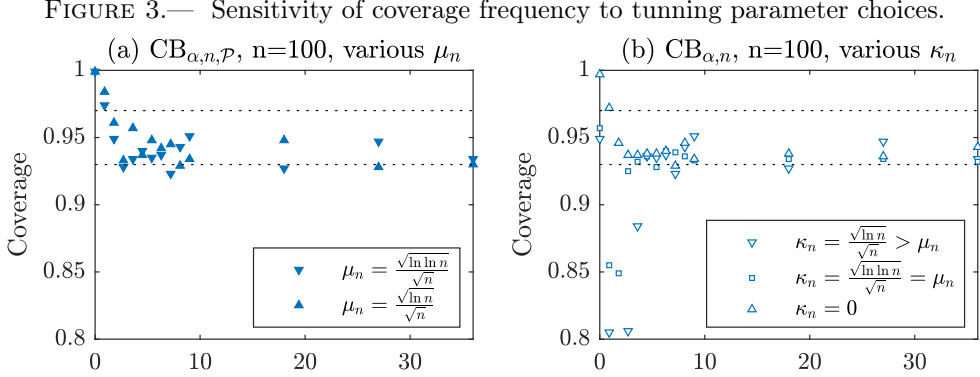
The difference in coverage frequencies of $\text{CB}_{\alpha,n}$ and $\text{CB}_{\alpha,n,\mathcal{P}}$ can be understood by considering the behaviour of the corresponding average excess length of the bounds, i.e. the Monte Carlo average of the difference between the corresponding confidence bounds and the true bound of the identified

set $\underline{\theta}_1 = -1$. Figure 2 panel (c) compares the lengths of the two bounds for the small sample case $n = 100$ where the point-wise inference becomes unreliable. One can see that the confidence bounds should have excess length approximately equal to 0.025 in order to match the coverage frequency with the nominal probability α in the proximity of $\omega = 0$. The length for $\text{CB}_{\alpha,n,\mathcal{P}}$ is larger than necessary for $\omega = 0$, while $\text{CB}_{\alpha,n}$ has the correct length for $\omega = 0$, but is too short in the neighborhood $(0, 4^\circ]$. For values $\omega \geq 5$ both confidence bounds have approximately the same length which also corresponds to the correct coverage probability as predicted by Theorems 1 and 4.

FIGURE 2.— Coverage frequency (a,b) and average excess length (c) in the 2-dimensional design for $\text{CB}_{\alpha,n}$ and $\text{CB}_{\alpha,n,\mathcal{P}}$ as function of ω in the 2-dimensional design.



Note: The dotted lines correspond to the asymptotic uniform 95% confidence interval for the parameter $p = 0.95$ of Bernoulli random variable based on a random sample of 1000 simulations based on Bonferroni correction for 14 hypothesis tests. Values of ω close to zero on Panel (a) result in non-negligible under-coverage. As sample size grows, the problematic area shrinks. Values of ω close to zero on Panel (b) result in non-negligible conservative coverage.



Note: The dotted lines correspond to the asymptotic uniform 95% confidence interval for the parameter $p = 0.95$ of Bernoulli random variable based on a random sample of 1000 simulations based on Bonferroni correction for 14 hypothesis tests.

4.3. Sensitivity to choice of μ_n and κ_n

First, consider the effect of the tuning parameter μ_n on the coverage probability for the uniformly valid confidence bounds $CB_{\alpha,n,\mathcal{P}}$. Figure 3 panel (a) compares the coverage frequency of $CB_{\alpha,n,\mathcal{P}}$ for sample size $n = 100$ as a function of ω for two choices: (i) $\hat{\mu}_0 \sqrt{n^{-1} \ln \ln n}$ (baseline), and (ii) $\hat{\mu}_0 \sqrt{n^{-1} \ln n}$ (large). The baseline option has slightly less conservative coverage in the neighborhood of the non-regular design ($\omega = 0$) and similar performance for other values of ω . As a result, the baseline option is used for all the other simulations.

Unlike its uniform counterpart, the point-wise valid confidence interval $CB_{\alpha,n}$ depends on additional tuning parameter sequence κ_n . Theorem 4 claims that for one-sided bounds $CB_{\alpha,n}$, the asymptotic coverage is exactly equal to α as long as κ_n shrinks to zero slower than μ_n . It allows values of κ_n to be smaller than μ_n and still result in (point-wise) valid inference. Figure 3 panel (b) compares three choices: (i) $\kappa_n = \hat{\mu}_0 \sqrt{n^{-1} \ln n} > \mu_n$; (ii) $\kappa_n = \mu_n$; (iii) $\kappa_n = 0 < \mu_n$. From the practitioner's perspective, options (ii) and (iii) are attractive since they only require specification of one tuning parameter μ_n . Both additional choices (ii) and (iii) result in valid point-wise coverage in large samples (see Appendix Figures 8 and 10).

For $n = 100$, $\omega > 4^\circ$ all three options have nearly indistinguishable coverage frequency. The behaviour is notably different between the three for $\omega \in [0, 4^\circ]$. The choice (ii) still results in lower than the required probability of coverage, but the problematic neighborhood $[0, 2^\circ]$ is smaller than for the choice (i). Choice $\mu_n = \kappa_n$ for $n = 10,000$ has particularly good alignment with the nominal coverage $\alpha = 0.95$. The choice (iii) essentially results in $CB_{\alpha,n}$ being the same length as $CB_{\alpha,n,\mathcal{P}}$ (see Appendix Figure 11). The similar performance in simulation suggests that $CB_{\alpha,n}$ with $\kappa_n = 0$ may have indeed uniformly valid coverage like $CB_{\alpha,n,\mathcal{P}}$. However, a formal study of uniform validity for $CB_{\alpha,n}$ with $\kappa_n = 0$ is more difficult to conduct than for $CB_{\alpha,n,\mathcal{P}}$.

4.4. *Effect of high dimensions*

In this section we study the effect of dimension of θ on the coverage probability and the length of the confidence bounds. The design of the constraints matrix $\mathbb{E}_P W = \mathbb{E}_P(A_P|b_P)$ is given by

$$A_P = \begin{pmatrix} 1 & -a \\ 0 & -I_{d-1} \\ 0 & I_{d-1} \end{pmatrix} \text{ and } b_P = \begin{pmatrix} 0 \\ \iota_2 \\ \iota_{d-3}/\sqrt{d-3} \\ \iota_2 \\ \iota_{d-3}/\sqrt{d-3} \end{pmatrix}.$$

The first line corresponds to an equality constraint that depends on a unit vector with direction $a \in \mathbb{R}^{d-1}$. This equality constraint is deterministic, i.e. $\text{Var } W_{1j} = 0$ for all $j = 1, \dots, d+1$. Other constraints are inequality constraints with coefficients that are i.i.d Gaussian r.v. with $\text{Var } W_{ij} = 0.01$ for all $j = 1, \dots, d+1$ and $i = 2, \dots, (1+2d)$. The identified set for the first coordinate θ_1 in this design corresponds to values of a support function of a $d-1$ dimensional rectangular box $[-1, 1]^2 \times [-1/\sqrt{d-3}, 1/\sqrt{d-3}]^{d-3}$ in direction a (and $-a$). As before, the focus on the lower bound, since by design the upper bound is symmetric. When $d = 3$ and $a = (1, 0)'$, this design reduces to the two-dimensional design considered in the previous subsection with $\omega = 0$.

We consider two possible values for a : (a) regular case, $a = \iota_{d-1}/\sqrt{d-1}$; (b) non-regular case, $a = (1, 0, \dots, 0)' \in \mathbb{R}^{d-1}$. In the regular case $\underline{\theta} = (-\iota_2, -\iota_{d-3}/\sqrt{d-3})$. In the non-regular case, the argmin set $\underline{\theta}$ is a convex hull of 2^{d-1} corner points with coordinates $(-1, \pm 1, \pm 1/\sqrt{d-3})$. Every such a corner point has the same distance to 0, equal to $\sqrt{3}$, for $d > 2$ regardless of the dimension d . This normalization was chosen to make performance comparisons as dimension grows while keeping the diameter of the identified set fixed. This way we isolate the effect of the number of dimensions and constraint from the effect of the diameter of the identified on the the dual variables. By design, the asymptotic variance of the regularized estimators will not change with the dimension. We use $n = 1000$ data points and $N = 1000$ Monte Carlo simulations. We use the first option for the tuning sequences, namely $\mu_n = \hat{\mu}_0 \sqrt{n}^{-1} \ln \ln n$ and $\kappa_n = \hat{\mu}_0 \sqrt{n}^{-1} \ln n$.

First, consider the coverage probability of the confidence bounds $\text{CB}_{\alpha, n}$ and $\text{CB}_{\alpha, n, \mathcal{P}}$. Panel (a) on Figure 4 shows that in the regular case both uniform and point-wise valid confidence set have coverage frequency that is within simulation error bound from the nominal coverage probability of $\alpha = 0.95$. At the same time, the coverage of the two procedures becomes noticeably different from each other for $d > 40$. It suggests that the size of the neighborhood where the uniform inference become important gets wider with d . The non-regular design provided on Panel (b) on Figure 4 shows that uniformly valid confidence bounds have coverage frequency nearly equal to 100%. The point-wise bounds $\text{CB}_{\alpha, n}$ are less conservative for all d , but still reach 100% coverage frequency for $d \geq 21$. It suggests that sample size required for exact asymptotic coverage for $\text{CB}_{\alpha, n}$ gets larger as dimension grows.

The coverage probability above the nominal level $\alpha = 0.95$ in the non-regular case in with high dimension d reflects the fact that as number of the corner points grows exponentially with d . As a result the non-regularized support function gets increasingly biased as it has asymptotic distribution of a minimum of 2^{d-1} Gaussian r.v. as evident from (2.22). Fortunately, this extreme conservatism only appears in the worst possible case. For a randomly chosen direction of a support function $a \in \mathbb{R}^{d-1}$, the performance is expected to be closer to the one on Panel (a) on Figure 4, at least for sufficiently large sample sizes n .

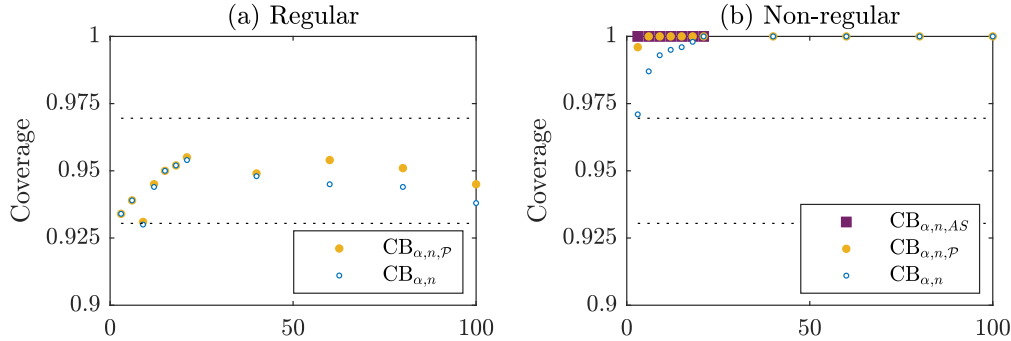
As a benchmark, I use $\text{CB}_{\alpha, n, AS}$ —one-sided confidence bounds for θ_1 based on [Andrews and Soares \(2010\)](#) with Bonferroni critical values implemented using the fast E-A-M algorithm of [Kaido](#)

et al. (2015). This choice of benchmark is one of the fastest available uniformly valid procedure in the literature. The two alternative approaches, Bugni et al. (2016) and Kaido et al. (2015), can provide uniformly valid CSs with potentially shorter average length than $CB_{\alpha,n,AS}$. Both of them, however, are expected to take considerably more time to compute because they add time-consuming profiling or calibration steps. AS results are only available in the non-regular case because of the software restrictions.¹⁵ Nevertheless, AS is not adaptive and is expected to have the similar length regardless whether we are at a regular or non-regular directions a . Coverage probability equal of confidence sets to 100% also applies to $CB_{\alpha,n,AS}$. Figure 5 shows that for $d \geq 9$ for $n = 1000$, the uniformly valid $CB_{\alpha,n,\mathcal{P}}$ is less conservative than $CB_{\alpha,n,AS}$.

4.5. Computational time

The main advantage of $CB_{\alpha,n,\mathcal{P}}$ when compared to $CB_{\alpha,n,AS}$ is computational gain that is achieved because of two factors: (i) using only linear and convex quadratic programs; (ii) no need to use multi-start procedures for global optimization of non-convex programs. Other alternative procedures like KMS and BCS would have additional burden of computing simulation based critical values and are expected to be much slower. Note, that the computational cost does not depend on whether design is regular or not, so Figure 6 compares the average computation time as a function of dimension d on a modern multi-core laptop.¹⁶ The computational time for the delta-method for regularized support functions grows very slowly with dimension d ; the average time for $d = 100$ is about 2 seconds. In contrast, $CB_{\alpha,n,AS}$ is nearly 1000 times slower already for $d = 21$.

FIGURE 4.— Coverage frequency in the d -dimensional design as function of d in (a) regular and (b) non-regular cases.

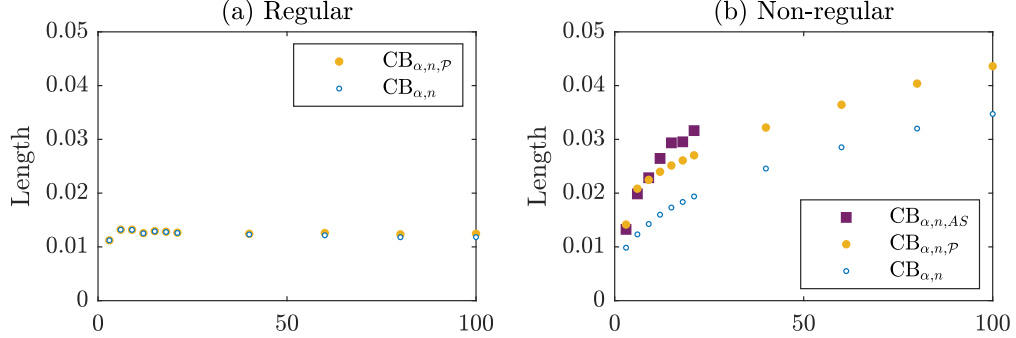


Note: The dotted lines correspond to the asymptotic uniform 95% confidence interval for the parameter $p = 0.95$ of Bernoulli random variable based on a random sample of 1000 simulations based on Bonferroni correction for 11 hypothesis tests. AS coverage frequency is computed using 100 simulations and for $d \leq 21$, while the delta-method confidence bounds are based on 1000 simulations. In all cases each simulation is based on $n = 1000$ observations.

¹⁵The code is available at <https://molinari.economics.cornell.edu/programs.html>. This code is highly optimized for the case of inference on bounds on individual coordinates of the identified set. The code does not allow making confidence sets for arbitrary directions a .


¹⁶All simulations are using 2023 Macbook Pro with M2 Max CPU (12 cores) and 96 GB of RAM. The EAM algorithm for $CB_{\alpha,n,AS}$ takes advantage of all 12 cores.

FIGURE 5.— Length in the d -dimensional design as function of d in (a) the regular and (b) non-regular cases.



Note: AS length is computed using 100 simulations and for $d \leq 21$, while the delta-method average length are based on 1000 simulations. In all cases each simulation is based on $n = 1000$ observations.

FIGURE 6.— Computational time in the d -dimensional design as function of d .



Note: The scale is logarithmic. AS average time is computed using 100 simulations and for $d \leq 21$ only, while the delta-method average times are based on 1000 simulations. In all cases each simulation is based on $n = 1000$ observations.

5. CONCLUSION

This paper demonstrated that the regularization approach provides a fast way to construct both point-wise and uniform CSs for a θ_1 that has comparable or shorter length to those in the existing literature. Monte Carlo simulations showed that the proposed CSs have good finite-sample-coverage properties. The computational benefits of the new approach are particularly prominent if the dimension of θ is large. The regularization framework can be extended in a number of ways to allow for overidentification and joint inference. The proposed approach is attractive in applications such as a linear model with an interval-valued outcome variable and a large number of regressors and in problems with parameters represented as intersection bounds.

Focus on the affine inequalities simplifies the large sample analysis. However, there are many interesting applications that characterize the identified set for the structural parameters using non-linear moment conditions, in particular, among the structural models of the industrial organization

(Pakes et al., 2015). Analysis of such non-linear moment inequalities goes beyond the scope of the current study, but would be a promising direction of future research.

REFERENCES

- ANDREWS, D. W. (2001): “Testing when a parameter is on the boundary of the maintained hypothesis,” *Econometrica*, 69, 683–734.
- ANDREWS, D. W. AND S. KWON (2019): “Inference in moment inequality models that is robust to spurious precision under model misspecification,” .
- ANDREWS, D. W. AND G. SOARES (2010): “Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection,” *Econometrica*, 78, 119–157.
- ANDREWS, I., J. ROTH, AND A. PAKES (2019): “Inference for linear conditional moment inequalities,” Tech. rep., National Bureau of Economic Research.
- BELLONI, A., F. BUGNI, AND V. CHERNOZHUKOV (2018): “Subvector inference in partially identified models with many moment inequalities,” *arXiv preprint arXiv:1806.11466*.
- BERESTEANU, A. AND F. MOLINARI (2008): “Asymptotic properties for a class of partially identified models,” *Econometrica*, 763–814.
- BERGE, C. (1963): *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity*, Courier Corporation.
- BONNANS, J. F. AND A. SHAPIRO (2000): *Perturbation Analysis of Optimization Problems*, Springer Science & Business Media.
- BONTEMPS, C., T. MAGNAC, AND E. MAURIN (2012): “Set identified linear models,” *Econometrica*, 80, 1129–1155.
- BOYD, S. AND L. VANDENBERGHE (2004): *Convex optimization*, Cambridge university press.
- BUGNI, F., I. CANAY, AND X. SHI (2016): “Inference for functions of partially identified parameters in moment inequality models,” Tech. rep., cemmap working paper, Centre for Microdata Methods and Practice.
- BUGNI, F. A., I. A. CANAY, AND X. SHI (2017): “Inference for subvectors and other functions of partially identified parameters in moment inequality models,” *Quantitative Economics*, 8, 1–38.
- CHERNOZHUKOV, V., D. CHETVERIKOV, AND K. KATO (2019): “Inference on causal and structural parameters using many moment inequalities,” *The Review of Economic Studies*, 86, 1867–1900.
- CHERNOZHUKOV, V., H. HONG, AND E. TAMER (2007): “Estimation and confidence regions for parameter sets in econometric models,” *Econometrica*, 75, 1243–1284.
- CHERNOZHUKOV, V., S. LEE, AND A. M. ROSEN (2013): “Intersection Bounds: Estimation and Inference,” *Econometrica*, 81, 667–737.
- CHERNOZHUKOV, V., W. K. NEWEY, AND A. SANTOS (2023): “Constrained conditional moment restriction models,” *Econometrica*, 91, 709–736.
- CHO, J. H. AND T. M. RUSSELL (2023): “Simple inference on functionals of set-identified parameters defined by linear moments,” *Journal of Business & Economic Statistics*, 1–16.
- COX, G. AND X. SHI (2019): “A simple uniformly valid test for inequalities,” *arXiv preprint arXiv:1907.06317*.
- DÜMBGEN, L. (1993): “On nondifferentiable functions and the bootstrap,” *Probability Theory and Related Fields*, 95, 125–140.
- FANG, Z. AND A. SANTOS (2018): “Inference on Directionally Differentiable Functions,” *The Review of Economic Studies*, rdy049.
- FISCHER, A. (1992): “A special Newton-type optimization method,” *Optimization*, 24, 269–284.
- FREYBERGER, J. AND J. L. HOROWITZ (2015): “Identification and shape restrictions in nonparametric instrumental variables estimation,” *Journal of Econometrics*, 189, 41 – 53.
- GAFAROV, B. (2014): “Identification in Dynamic Models Using Sign Restrictions,” *Available at SSRN 2384811*.
- (2017): “Essays on Partially Identified Models,” Ph.D. thesis, Pennsylvania State University.
- GAFAROV, B., M. MEIER, AND J. L. M. OLEA (2018): “Delta-Method inference for a class of set-identified SVARs,” *Journal of Econometrics*, 203, 316–327.
- GOLISHNIKOV, M. AND A. F. IZMAILOV (2006): “Newton-type methods for constrained optimization with nonregular constraints,” *Computational Mathematics and Mathematical Physics*, 46, 1299–1319.
- HAILE, P. A. AND E. TAMER (2003): “Inference with an incomplete model of English auctions,” *Journal of Political Economy*, 111, 1–51.
- HALL, P. AND H. MILLER (2010): “Bootstrap confidence intervals and hypothesis tests for extrema of parameters,” *Biometrika*, 97, 881–892.
- HIRANO, K. AND J. R. PORTER (2012): “Impossibility results for nondifferentiable functionals,” *Econometrica*, 1769–1790.

- HONG, H. AND J. LI (2018): “The numerical delta method,” *Journal of Econometrics*, 206, 379–394.
- HONORÉ, B. E. AND A. LLERAS-MUNEY (2006): “Bounds in competing risks models and the war on cancer,” *Econometrica*, 74, 1675–1698.
- HONORÉ, B. E. AND E. TAMER (2006): “Bounds on parameters in panel dynamic discrete choice models,” *Econometrica*, 74, 611–629.
- HSIEH, Y.-W., X. SHI, AND M. SHUM (2021): “Inference on estimators defined by mathematical programming,” *Journal of Econometrics*.
- IMBENS, G. W. AND C. F. MANSKI (2004): “Confidence Intervals for Partially Identified Parameters,” *Econometrica*, 72, 1845–1857.
- KAIDO, H., F. MOLINARI, AND J. STOYE (2015): “Inference for projections of identified sets,” *manuscript*.
- (2019a): “Confidence intervals for projections of partially identified parameters,” *Econometrica*, 87, 1397–1432.
- (2019b): “Constraint qualifications in partial identification,” *Econometric Theory*, 1–24.
- KAIDO, H. AND A. SANTOS (2014): “Asymptotically Efficient Estimation of Models Defined by Convex Moment Inequalities,” *Econometrica*, 82, 387–413.
- KASY, M. (2016): “Partial identification, distributional preferences, and the welfare ranking of policies,” *Review of Economics and Statistics*, 98, 111–131.
- KLINE, P. AND M. TARTARI (2016): “Bounding the labor supply responses to a randomized welfare experiment: A revealed preference approach,” *American Economic Review*, 106, 972–1014.
- LAFFÈRS, L. (2018): “Bounding average treatment effects using linear programming,” *Empirical Economics*.
- MANSKI, C. F. (2003): *Partial identification of probability distributions*, Springer Science & Business Media.
- MANSKI, C. F. AND J. V. PEPPER (2000): “Monotone instrumental variables: With an application to the returns to schooling,” *Econometrica*, 68, 997–1010.
- MANSKI, C. F. AND E. TAMER (2002): “Inference on regressions with interval data on a regressor or outcome,” *Econometrica*, 70, 519–546.
- MINCER, J. A. ET AL. (1974): “Schooling, Experience, and Earnings,” *NBER Books*.
- MOLINARI, F. (2008): “Partial identification of probability distributions with misclassified data,” *Journal of Econometrics*, 144, 81 – 117.
- OK, E. A. (2007): *Real analysis with economic applications*, vol. 10, Princeton University Press.
- PAKES, A., J. PORTER, K. HO, AND J. ISHII (2015): “Moment inequalities and their application,” *Econometrica*, 83, 315–334.
- ROCKAFELLAR, R. T. (1970): *Convex analysis*, vol. 36, Princeton university press.
- RUSSELL, T. M. (2017): “Sharp Bounds on Functionals of the Joint Distribution in the Analysis of Treatment Effects,” *Manuscript, University of Toronto*.
- SHAPIRO, A. (1991): “Asymptotic analysis of stochastic programs,” *Annals of Operations Research*, 30, 169–186.
- SHAPIRO, A., D. DENTCHEVA, AND A. RUSZCZYNSKI (2014): *Lectures on stochastic programming: modeling and theory*, vol. 16, SIAM.
- SHI, X. (2020): “Uniform Inference when Parameters Are Subject to Linear Inequality Constraints,” Ph.D. thesis, University of Washington.
- SHI, X., M. SHUM, AND W. SONG (2018): “Estimating Semi-Parametric Panel Multinomial Choice Models Using Cyclic Monotonicity,” *Econometrica*, 86, 737–761.
- STEWART, M. B. (1983): “On least squares estimation when the dependent variable is grouped,” *The Review of Economic Studies*, 50, 737–753.
- SYRGKANIS, V., E. TAMER, AND J. ZIANI (2017): “Inference on Auctions with Weak Assumptions on Information,” *arXiv preprint arXiv:1710.03830*.
- TORGOVITSKY, A. (2016): “Nonparametric inference on state dependence with applications to employment dynamics,” *Manuscript, University of Chicago*, 1.
- (2019): “Partial identification by extending subdistributions,” *Quantitative Economics*, 10, 105–144.
- TROSTEL, P., I. WALKER, AND P. WOOLLEY (2002): “Estimates of the economic return to schooling for 28 countries,” *Labour economics*, 9, 1–16.
- VAN DER VAART, A. AND J. WELLNER (1996): *Weak Convergence and Empirical Processes: With Applications to Statistics*, Springer Science & Business Media.
- WACHSMUTH, G. (2013): “On LICQ and the uniqueness of Lagrange multipliers,” *Operations Research Letters*, 41, 78–80.
- YE, Y. AND E. TSE (1989): “An extension of Karmarkar’s projective algorithm for convex quadratic programming,” *Mathematical programming*, 44, 157–179.
- YURINSKII, V. V. (1978): “On the error of the Gaussian approximation for convolutions,” *Theory of Probability & Its Applications*, 22, 236–247.

APPENDIX A: DETAILED PROOFS

A.1. Topological properties of optimal solutions

Consider any distribution P with support on $\mathbb{R}^{(k-2d) \times (d+1)}$ such that $(A_P, b_P) \triangleq \mathbb{E}_P W$ exist. Let $\mathcal{J}^a(\theta; P) \subset \{1, \dots, k\}$ be the set of indices of moment equality and inequality constraints active at θ , i.e. all j s.t. $m_j(\theta, P) \triangleq \mathbb{E}_P g_j(W, \theta) = 0$. $\mathcal{J}^a(\theta; P)$ can be empty.

LEMMA 1 (Characterization of the optimal solution) *Under Assumption 1 for any $\mu \geq 0$ any minimizer θ for Program (3.2) is a solution to the corresponding Karush–Kuhn–Tucker (KKT) optimality conditions for some finite $\lambda \in \mathbb{R}^k$,*

$$\begin{aligned} \text{(A.1)} \quad & (e_1 + 2\mu\theta)' = -\lambda' A_P, \\ \text{(A.2)} \quad & m_j(\theta, P) = 0 \quad j \in \mathcal{J}^{eq}, \\ \text{(A.3)} \quad & m_j(\theta, P) \leq 0, \lambda_j \geq 0, \lambda_j m_j(\theta, P) = 0 \quad j \in \mathcal{J}^{ineq}. \end{aligned}$$

PROOF: By Assumption 1, $\Theta(P) \subset \Theta$ is non-empty and closed, so the global optima for Program (3.2) exist. Program (3.2) is convex for any $\mu \geq 0$, i.e. the objective function is convex, the constraints are affine. Assumption 1 implies Slater's condition. Since the Program (3.2) is convex, any global optimum $\underline{\theta}(\mu, P)$ of Program 3.2 satisfies (A.1)-(A.3) for some finite vector of Lagrange multipliers λ (maybe non-unique) (see p.244 in Boyd and Vandenberghe (2004)). Q.E.D.

If we introduce notation $\underline{\mathcal{L}}(\lambda, \theta; \mu, P) \triangleq \theta_1 + \mu \|\theta\|^2 + \lambda' m(\theta, P)$, (A.1) becomes

$$\partial_{\theta} \underline{\mathcal{L}}(\lambda, \theta; \mu, P) = 0$$

Let $\underline{\xi}(\mu, P) \triangleq (\underline{\theta}(\mu, P), \underline{\lambda}(\mu, P))$ be a set of solutions to (A.1)-(A.3). In order to have a unique solution λ Program (3.2) need to meet a stronger constraint qualification condition, Assumption 3 defined below.

As before, I use symbols $\mathcal{J}^a(\theta; P)$, $\mathcal{J}^a(\mu, P)$ etc to denote the projectors on the coordinates with the corresponding indices. Let $\mathbb{J}_{d+1}^d \triangleq (e_1, \dots, e_d)$. Using this notation, Assumption 3 takes the following equivalent form

ASSUMPTION 8 (Linear Independence Constraint Qualification (LICQ)) *The matrix $\mathbb{J}^a(\theta; P) A_P$ has full row rank for any $\theta \in \Theta(P)$.*

LEMMA 2 (Sufficient condition for LICQ) *Assumptions 4-5 imply Assumption 3.*

PROOF: Assumption 5 implies that $\mathcal{J}^a(\theta; P)$ has at most d elements at any $\theta \in \Theta(P)$. Consider any point $\theta \in \Theta(P)$. The set \mathcal{J}^{ineq} includes (2.4), so $k \geq 2d \geq d+1$. It implies that there exists a set \mathcal{J} with $|\mathcal{J}| = d$ such that $\mathcal{J}^a(\theta; P) \subset \mathcal{J}$. By Assumption 4, $\text{rk}[\mathbb{J}\mathbb{E}_P W] = d$, so $M \triangleq \mathbb{J}^a(\theta; P)(A_P, b_P)$ has full row rank which is equal to $|\mathcal{J}^a|$. By definition, $\mathbb{J}^a(\theta; P) A_P \theta = \mathbb{J}^a b_P$. It implies by the Rouché–Capelli theorem that the matrices $M_{\theta} \triangleq \mathbb{J}^a(\theta; P) A_P$ and M have the same rank. This result implies Assumption 3. Q.E.D.

The inverse implication does not hold in general as the following remark shows.

REMARK 7 *Assumption 3 implies Assumption 5 and that for any $\theta \in \Theta(P)$*

$$\text{(A.4)} \quad \text{rk}[\mathbb{J}^a(\theta; P) \mathbb{E}_P W] = |\mathcal{J}^a(\theta; P)|.$$

Indeed, suppose that Assumption 3 holds. It immediately implies Assumption 5. To see (A.4) consider any point $\theta \in \Theta(P)$ such that $|\mathcal{J}^a(\theta; P)| \leq d$. By the Rouché–Capelli theorem and the full row rank property of M_{θ} correspondingly,

$$\text{rk}(M) = \text{rk}(M_{\theta}) = |\mathcal{J}^a(\theta; P)|.$$

LEMMA 3 (Uniqueness of the optimal solutions) *Suppose that both Assumptions 1 and 3 are satisfied. Then for any $\mu \geq 0$ the set of multipliers $\underline{\lambda}(\mu, P)$ is a singleton. Moreover if $\mu > 0$, then $\underline{\xi}(\mu, P)$ is a singleton.*

PROOF: By definition of $\mathbb{J}^a(\theta; P)$, any θ and λ satisfying (A.1) satisfy

$$(A.5) \quad \lambda'(\mathbb{J}^a(\theta; P))' \mathbb{J}^a(\theta; P) = \lambda'.$$

So (A.1) becomes

$$(A.6) \quad (e_1 + 2\mu\theta)' = -\gamma' \mathbb{J}^a(\theta; P) A_P,$$

where $\gamma' \triangleq \lambda'(\mathbb{J}^a(\theta; P))' \in \mathbb{R}^{|\mathcal{J}^a(\theta; P)|}$. By Assumption 3, for any $\theta \in \Theta(P)$ the matrix $A \triangleq \mathbb{J}^a(\theta; P) A_P$ has full rank. Hence for any θ there can be at most one $\gamma^* \in \mathbb{R}^{|\mathcal{J}^a(\theta; P)|}$ satisfying (A.6). If $e_1 + 2\mu\theta = 0$, then trivially λ is a zero vector. Otherwise it is given by

$$(A.7) \quad \gamma^* = -(AA')^{-1} A' (e_1 + 2\mu\theta).$$

Then $(\underline{\lambda}(\mu, P))' \triangleq (\gamma^*)' \mathbb{J}^a(\theta; P)$ is the unique solution to (A.1)-(A.3) for any solution θ .

Now consider the case $\mu > 0$. The second order derivative matrix of $\underline{\mathcal{L}}(\lambda, \theta; \mu, P)$ with respect to θ at any solution $\underline{\xi}(\mu, P)$ is $2\mu I_d$. It is positive definite for any $\mu > 0$, so the Second Order Sufficient Condition (SOSC) is satisfied at any point. By Theorem 3.63 from Bonnans and Shapiro (2000) the second order growth condition holds at $\underline{\theta}(\mu, P)$, i.e. $\exists \varepsilon > 0$ and $c > 0$ s.t. for $\forall \theta \in \Theta(P)$ s.t. $\|\theta - \underline{\theta}(\mu, P)\| < \varepsilon$ the following inequality holds

$$\theta_1 + \mu \|\theta\|^2 \geq e_1' \underline{\theta}(\mu, P) + \mu \|\underline{\theta}(\mu, P)\|^2 + c \|\theta - \underline{\theta}(\mu, P)\|^2.$$

So the value of the objective function at $\underline{\theta}(\mu, P)$ is strictly smaller than the value at any other point in a neighborhood of $\underline{\theta}(\mu, P)$. Since for the convex program the set of global optima is convex and connected, it implies that $\underline{\theta}(\mu, P)$ is the unique global minimizer. Q.E.D.

It is assumed (see eq. (2.4)) that the system of inequalities includes deterministic constraints

$$(A.8) \quad -\infty < -\underline{c}_\ell \leq \theta_\ell \leq \bar{c}_\ell < \infty \quad \text{for } \ell = 1, \dots, d.$$

These constraints define a compact set Θ . By construction, $\Theta(P) \subset \Theta$ for all measures P .

LEMMA 4 *Suppose that Assumptions 1,4-5 are satisfied and $\mu \leq 1/2$. Then*

$$(A.9) \quad \|\underline{\lambda}(\mu, P)\|^2 \leq C_\Lambda^2 \triangleq \frac{C_\Theta^3}{\eta^2} < \infty,$$

where $C_\Theta \triangleq (1 + \max_{\theta \in \Theta} \|\theta\|)$.

PROOF: Consider any point $\theta \in \underline{\theta}(\mu, P)$ and the corresponding $\mathcal{J}^a(\theta; P)$. Let $A \triangleq \mathbb{J}^a(\theta; P) A_P$ and $b \triangleq \mathbb{J}^a(\theta; P) b_P$. Let $\eta_A^2 \triangleq \text{eig}(AA')$ so that equation (A.7) implies

$$(A.10) \quad \|\underline{\lambda}(\mu, P)\| \leq \eta_A^{-1} \|e_1 + 2\mu\theta\|.$$

By the variational property of eigenvalues,

$$(A.11) \quad \eta_A^2 = \min_{v \in \mathbb{R}^\ell} \frac{v' AA' v}{v' v}.$$

By Assumption 4

$$\eta^2 \leq \text{eig}((A|b)(A|b)') \triangleq \min_{v \in \mathbb{R}^\ell} \frac{v'(AA' + bb')v}{v'v}.$$

Let v_A be any minimizer of the r.h.s. of (A.11) such that $v_A'v_A = 1$. Then

$$(A.12) \quad \eta^2 \leq v_A'(AA' + bb')v_A = (A'v_A)'(I_d + \theta\theta')(A'v_A)$$

where the last equality holds since by definition $b = A\theta$. Finally,

$$(A.13) \quad \frac{\eta^2}{\eta_A^2} \leq \frac{(A'v_A)'(I_d + \theta\theta')(A'v_A)}{(A'v_A)'(A'v_A)} \leq \|I_d + \theta\theta'\| \leq C_\Theta.$$

Result (A.9) then follows from (A.10) and (A.13) for any $\mu \leq 1/2$.

Q.E.D.

REMARK 8 Equation (A.13) provides bound for, A , a matrix with gradients of active moment conditions at any point $\theta \in \Theta$,

$$(A.14) \quad \left\| (AA')^{-1} \right\| \leq C_\Theta \eta^{-2}.$$

The function $\phi(a, b) \triangleq \sqrt{a^2 + b^2} + a - b$, considered in Fischer (1992), has the following property.

PROPOSITION 1

$$(A.15) \quad \phi(a, b) = 0 \text{ if and only if } a \leq 0, b \geq 0, ab = 0.$$

It can be used to replace (A.3) with an equivalent equality so that the KKT system becomes a system of equations. This result can be used to establish the continuity of the solutions in μ as the following lemma shows.

LEMMA 5 Under Assumptions 1,4-5 $\underline{\xi}(\mu, P)$ is u.h.c. in μ ; $\underline{v}(\mu, P)$ is continuous in μ for $\mu \geq 0$.

PROOF: By Proposition 1 equation (A.3) is equivalent to

$$(A.16) \quad \phi(m_j(\theta, P), \lambda_j) = 0 \text{ for } j \in \mathcal{J}^{ineq}.$$

Solutions to (A.1),(A.2),(A.16) coincide with solutions to

$$(A.17) \quad \Psi(\theta, \lambda; \mu, P) \triangleq \|\partial_\theta \underline{\mathcal{L}}(\lambda, \theta; \mu, P)\|_2^2 + \sum_{j \in \mathcal{J}^{eq}} (m_j(\theta, P))^2 + \sum_{j \in \mathcal{J}^{ineq}} (\phi(m_j(\theta, P), \lambda_j))^2 = 0.$$

Lemmas 3-4 imply that $\underline{\lambda}(\mu, P)$ is unique and satisfies (A.9) for any $\mu \in [0, 1/2]$. So the solution to (A.17) coincides with solutions of

$$(A.18) \quad \begin{aligned} & \min_{\theta, \lambda} \Psi(\theta, \lambda; \mu, P) \\ & \text{s.t. } \theta \in \Theta, \lambda \in \mathbb{R}^k, \|\lambda\| \leq C_\Lambda. \end{aligned}$$

The objective function of this program is continuous in μ and the domain is a compact valued continuous correspondence in μ . By the Maximum Theorem (see Ok (2007)) $\underline{\xi}(\mu, P)$ is u.h.c. function of $\mu \geq 0$.

Function $\underline{v}(\mu, P) = e_1' \underline{\theta}(\mu, P) + \mu \|\underline{\theta}(\mu, P)\|^2$ is a composition of u.h.c. functions and hence, by Theorem VI.2.1' from Berge (1963), is u.h.c. in $\mu \in \mathbb{R}_+$. Since by definition $\underline{v}(\mu, P)$ is a single-valued function, u.h.c. implies continuity in $\mu \geq 0$ for any fixed P .

Q.E.D.

A.2. Smoothness properties

In this section we will study the directional derivatives of the value and the optimal solutions of Program (3.2). We will pursue this goal by taking a limit of the perturbed program defined below as the size of the perturbation goes to zero. Consider a perturbation in parameters $\mathbb{E}_P W = (A_P | b_P)$ and μ in a direction $h' = (\text{vec}(h_W)', h_\mu) \in \mathbb{R}^{k(d+1)+1}$, where $h_W \triangleq (h_A, h_b) \in \mathbb{R}^{k \times (d+1)}$ and $h_\mu \in \mathbb{R}$. The corresponding perturbation in the constraints is $\dot{m}_h(\theta) \triangleq h_A \theta - h_b$. Given these directions, for any $t \geq 0$, $\mu > 0$ we can define a perturbed program,

$$(A.19) \quad \min_{\theta \in \Theta} \quad e_1' \theta + (\mu + t h_\mu) \|\theta\|^2,$$

$$\text{s.t.} \quad \begin{cases} m_j(\theta, P) + t \dot{m}_{h,j}(\theta) = 0 & \text{for } j \in \mathcal{J}^{eq}, \\ m_j(\theta, P) + t \dot{m}_{h,j}(\theta) \leq 0 & \text{for } j \in \mathcal{J}^{ineq}. \end{cases}$$

Perturbations of inequality constrained programs can lead to changes in the set of the constraints active at the optimum in response to an arbitrarily small perturbations. It is instructive to consider the following sets of constraints for the unperturbed program, i.e. with $h_A = 0, h_b = 0, h_\mu = 0$,

$$\begin{aligned} \mathcal{J}^+(\mu, P) &\triangleq \left\{ j \in \mathcal{J}^{ineq} \mid \underline{\lambda}_j(\mu, P) > 0 \right\} \cup \mathcal{J}^{eq}, \\ \mathcal{J}^-(\mu, P) &\triangleq \left\{ j \in \mathcal{J}^{ineq} \mid m_j(\underline{\theta}(\mu, P), P) > 0 \right\}, \\ \mathcal{J}^0(\mu, P) &\triangleq \left\{ j \in \mathcal{J}^{ineq} \mid \underline{\lambda}_j(\mu, P) = 0, m_j(\underline{\theta}(\mu, P), P) = 0 \right\}, \\ \mathcal{J}^a(\mu, P) &\triangleq \mathcal{J}^0(\mu, P) \cup \mathcal{J}^+(\mu, P). \end{aligned}$$

Set \mathcal{J}^+ contains active inequality constraints with positive Lagrange multipliers and the equality constraints. These constraints will remain active for small enough perturbations in any directions h_A, h_b (by continuity of the Lagrange multipliers). Set \mathcal{J}^- contains slack constraints. They will remain slack in response to sufficiently small perturbations (by continuity of the optimal solution and the constraints functions). Set \mathcal{J}^0 contains active inequality constraints with zero Lagrange multipliers. If we drop these constraints the optimal solution will not change, but they play an important role in the perturbed program. Constraints in \mathcal{J}^0 become inactive in response to perturbations in some directions no matter how small the perturbation is or remain active and acquire positive Lagrange multipliers for other directions. The optimal solution ξ would be fully differentiable iff \mathcal{J}^0 is empty as will be evident from the explicit formula for its directional derivative. Finally, \mathcal{J}^a contains all active constraints at the optimal solution.

Suppose that the perturbation size $t > 0$ is small enough such that Program (A.19) satisfies Assumptions 1,4-5. Then it has a unique solution $\underline{\xi}_h(t)$ for $0 \leq t < T$ which can be represented as $\underline{\xi}_h(t) = \underline{\xi} + t \dot{\underline{\xi}}_h$, as the following lemma shows. The directional derivative $\dot{\underline{\xi}}_h'(\mu, P) \triangleq (\dot{\underline{\theta}}'(\mu, P), \dot{\underline{\lambda}}'(\mu, P))$ will depend on the following objects,

$$\begin{aligned} \mathcal{J}^h(\mu, P) &\triangleq \left\{ j \in \mathcal{J}^0(\mu, P) \mid \dot{\underline{\lambda}}_{h,j}(\mu, P) > 0 \right\} \cup \mathcal{J}^+(\mu, P), \\ A_h(\mu, P) &\triangleq \mathbb{J}^h(\mu, P) A_P, \\ Q_h &\triangleq I_d - A_h' (A_h A_h')^{-1} A_h, \\ A^\dagger &\triangleq A' (A A')^{-1}. \end{aligned}$$

I suppress the argument (μ, P) from now on.

LEMMA 6 (Local linear representation) *Suppose that Assumptions 1,4-5 hold for P . There is a neighborhood $[0, T(\mu, h, P)]$ in which Program (A.19) has a unique solution $\underline{\xi}_h(t) = \underline{\xi} + t \dot{\underline{\xi}}_h$ with*

$$(A.20) \quad \dot{\underline{\xi}}_h = - \begin{pmatrix} (2\mu)^{-1} Q_h & A_h^\dagger \\ (\mathbb{J}^h)' (A_h^\dagger)' & -2\mu (\mathbb{J}^h)' (A_h A_h')^{-1} \end{pmatrix} \begin{pmatrix} (h_A)' \underline{\lambda} + 2h_\mu \underline{\theta} \\ \mathbb{J}^h(h_A \underline{\theta} - h_b) \end{pmatrix}.$$

PROOF: By Lemma 3, if $t = 0$ Program (A.19) has a unique solution ξ . Since this solution satisfies LICQ (Assumption 3) and SOSC, it is strongly regular by Proposition 5.38 from BS(2000). The remaining argument follows the proof of Theorem 5.60 from BS(2000), which uses an implicit function theorem for generalized equations (Theorem 5.13 in the same book) at a strongly regular solution. We are going to apply it to the KKT conditions for Program (A.19) at the strongly regular solution ξ ,

$$\begin{aligned} (A.21) \quad & \left\{ \begin{aligned} (e_1 + 2(\mu + h_\mu t)\theta)' &= -\lambda'(A_P + th_A), \\ m_j(\theta, P) + t\dot{m}_{h;j}(\theta) &= 0 \end{aligned} \right. & j \in \mathcal{J}^{eq}, \\ (A.22) \quad & \\ (A.23) \quad & \left\{ \begin{aligned} \phi(m_j(\theta, P) + t\dot{m}_{h;j}(\theta), \lambda_j) &= 0 \end{aligned} \right. & j \in \mathcal{J}^{ineq}. \end{aligned}$$

By Theorem 5.60 from BS(2000), $\xi_h(t)$ is analytic in t in some neighborhood $[0, T(\mu, h, P)]$, i.e. it can be represented as power series. First, let us compute the linear term. By the strong regularity and Theorem 5.13 in BS(2000), there exist a unique solution $(\dot{\theta}, \dot{\lambda})$ to the following system of equations (this system is the gradient of (A.21)-(A.23) with respect to t at point $t = 0$)

$$\begin{aligned} (A.24) \quad & \left\{ \begin{aligned} 2\mu\dot{\theta}'I_d + \dot{\lambda}'A_P &= -\lambda'h_A - 2h_\mu\dot{\theta}', \\ e'_j A_P \dot{\theta} + \dot{m}_{h;j}(\dot{\theta}) &= 0 \end{aligned} \right. & j \in \mathcal{J}^+(\mu, P), \\ (A.25) \quad & \\ (A.26) \quad & \left\{ \begin{aligned} \phi(e'_j A_P \dot{\theta} + \dot{m}_{h;j}(\dot{\theta}), \dot{\lambda}_j) &= 0 \end{aligned} \right. & j \in \mathcal{J}^0(\mu, P), \\ (A.27) \quad & \left\{ \begin{aligned} \dot{\lambda}_{h;j} &= 0 \end{aligned} \right. & j \in \mathcal{J}^-(\mu, P). \end{aligned}$$

This unique solution determines the set \mathcal{J}^h . System (A.24)-(A.27) can be represented in a matrix form:¹⁷

$$(A.28) \quad \begin{pmatrix} 2\mu I_d & A'_h \\ A_h & 0 \end{pmatrix} \begin{pmatrix} \dot{\theta} \\ \mathbb{J}^h \dot{\lambda} \end{pmatrix} = - \begin{pmatrix} (h_A)' \lambda + 2h_\mu \theta \\ \mathbb{J}^h (h_A \theta - h_b) \end{pmatrix}.$$

In addition to that, $\dot{\lambda} = (\mathbb{J}^h)' \mathbb{J}^h \dot{\lambda}$. One can check by direct computation that

$$\begin{pmatrix} 2\mu I_d & A'_h \\ A_h & 0 \end{pmatrix}^{-1} = \begin{pmatrix} (2\mu)^{-1} Q_h & A_h^\dagger \\ (A_h^\dagger)' & -2\mu (A_h A_h')^{-1} \end{pmatrix}.$$

Since the higher order derivatives of every constraint function and the objective function of Program (A.19) with respect to t are zero, the higher order directional derivatives of ξ_h are equal to zero at $t = 0$. Thus the power series expansion of ξ_h has only constant and linear terms. Q.E.D.

Now we can rewrite Program (A.19) in an explicit form assuming $h_\mu = 0$,

$$(A.29) \quad \begin{aligned} \min_{\theta \in \Theta} \quad & e'_1 \theta + \mu \|\theta\|^2, \\ \text{s.t.} \quad & \begin{cases} e'_j (A_P + th_A) \theta = b_P + th_b & \text{for } j \in \mathcal{J}^{eq}, \\ e'_j (A_P + th_A) \theta \leq b_P + th_b & \text{for } j \in \mathcal{J}^{ineq}. \end{cases} \end{aligned}$$

LEMMA 7 *Suppose that Assumptions 1,4-5 hold for P . There exist such $\delta > 0$ such that for any $h = (h_A, h_b) \in \mathbb{R}^{k \times (d+1)}$ with norm $\|h\| < \delta$ and any $\mu \in (0, 1/2]$ and $t \in [0, 1]$ the solution of Program (A.29) satisfies*

$$(A.30) \quad \|\theta_h(t) - \theta\| \leq L_\theta \frac{\|th\|}{\mu}$$

$$(A.31) \quad \|\lambda_h(t) - \lambda\| \leq L_\lambda \|th\|$$

¹⁷Compare Equation (A.28) with Equation (5.81) on page 186 in Shapiro et al. (2014).

$$(A.32) \quad |\underline{v}_h(t) - \underline{v} - t\underline{\lambda}'(h_A\underline{\theta} - h_b)| \leq L_v \frac{\|\underline{th}\|^2}{\mu},$$

where $L_\theta = \frac{\sqrt{2C_\Theta^3}}{(\eta/2)}$, $L_\lambda = \frac{(\|\mathbb{E}_P W\| + \delta)C_\Theta^4 + \eta^2 C_\Theta^2}{(\eta/2)^4}$, and $L_v = \frac{2C_\Theta^3}{(\eta/2)^2}$.

PROOF: First consider $t \in [0, T(\mu, h, P)]$ with $T(\mu, h, P)$ defined in Lemma 6. By Lemma 6 the value function $\underline{v}_h(t) \triangleq e_1' \underline{\theta}_h(t) + \mu \|\underline{\theta}_h(t)\|^2$ can be represented as

$$(A.33) \quad \underline{v}_h(t) = \underline{v} + t(e_1 + 2\mu\underline{\theta})' \underline{\dot{\theta}} + \mu t^2 \|\underline{\dot{\theta}}\|^2.$$

First, consider the second term. Since by definition $\mathcal{J}^+ \subseteq \mathcal{J}^h$, we have $\underline{\lambda}' = \underline{\lambda}'(\mathbb{J}^h)' \mathbb{J}^h$. Correspondingly, $\underline{\lambda}' A_P = \underline{\lambda}'(\mathbb{J}^h)' A_h$. By Lemma 3, $(e_1 + 2\mu\underline{\theta})' = -\underline{\lambda}' A_P$. So

$$(A.34) \quad (e_1 + 2\mu\underline{\theta})' Q_h = -\underline{\lambda}'(\mathbb{J}^h)'(A_h Q_h) = 0,$$

$$(A.35) \quad (e_1 + 2\mu\underline{\theta})' A_h^\dagger = -\underline{\lambda}'(\mathbb{J}^h)'(A_h A_h^\dagger) = -\underline{\lambda}'(\mathbb{J}^h)'.$$

Equations (A.34) and (A.35) imply that

$$(A.36) \quad (e_1 + 2\mu\underline{\theta})' \underline{\dot{\theta}} = -\underline{\lambda}' \dot{m}_h(\underline{\theta}).$$

Second, by Lemma 4 and Remark 8 for any $\mu \leq 1/2$,

$$(A.37) \quad \|(A_h A_h')^{-1}\| \leq C_\Theta \eta^{-2}(P) \text{ and } \|\underline{\lambda}\|^2 \leq C_\Theta^3 \eta^{-2}(P).$$

Then by the triangular inequality and inequalities (A.37) (and the fact that $C_\Theta \geq 1$)

$$(A.38) \quad \|\underline{\dot{\theta}}\|^2 = \frac{1}{(2\mu)^2} (\underline{\lambda}' h_A) Q_h (\underline{\lambda}' h_A)' + \dot{m}_h(\underline{\theta})' (\mathbb{J}^h)' (A_h A_h')^{-1} \mathbb{J}^h \dot{m}_h(\underline{\theta})$$

$$(A.39) \quad \leq \|h_W\|^2 \frac{C_\Theta^3}{\eta^2(P)} \left(\frac{1}{\mu^2} + 1 \right),$$

which implies (A.30). Equation (A.31) can be proven similarly,

$$(A.40) \quad \|\underline{\dot{\lambda}}\| = \left\| (\mathbb{J}^h)'(A_h^\dagger)' (\underline{\lambda}' h_A) - 2\mu (\mathbb{J}^h)' (A_h A_h')^{-1} \mathbb{J}^h \dot{m}_h(\underline{\theta}) \right\|$$

$$(A.41) \quad \leq \frac{\|\mathbb{E}_P W\| C_\Theta^4 + \eta^2(P) C_\Theta^2}{\eta^4(P)} \|h_W\|$$

Finally, the bound in (A.32) follows from equations (A.30), (A.33), and (A.36).

To extend the argument to the entire interval $t \in [0, 1]$, notice that $\eta(P)$ and $s(P)$ are Lipschitz-continuous. So there exist $\delta > 0$ such that $\eta(\mathbb{E}_P W + th) > \eta/2$ and $s(\mathbb{E}_P W + th) > s/2$ for any $h = (h_A, h_b) \in \mathbb{R}^{k \times (d+1)}$ with norm $\|h\| < \delta$ and any $t \in [0, 1]$. Such δ can be chosen uniformly for $P \in \mathcal{P}$. Now we can replace $\|\mathbb{E}_P W\|$ with $(\|\mathbb{E}_P W\| + \delta)$ and $\eta(P)$ with $\eta/2$ to obtain uniform constants L_θ , L_λ , and L_v .

Q.E.D.

A.3. Proof of Theorem 1

LEMMA 8 *Suppose that Assumptions 1,4-5 hold. There exist some $\bar{\mu}(P) > 0$ such that for any $\mu < \bar{\mu}(P)$ the solution to Program (3.2), $\underline{\theta}(\mu, P)$ is constant.*

PROOF: Consider a direction $h' = (\text{vec}(h_W)', h_\mu) \in \mathbb{R}^{k(d+1)+1}$ satisfying $h_W = 0$, $h_\mu = 1$ and any $\mu_0 > 0$ in a neighborhood of 0. By Lemma 6 we get

$$(A.42) \quad \dot{\underline{\theta}}(\mu_0, P) = -\frac{1}{\mu_0} Q_h(\mu_0, P) \underline{\theta}(\mu_0, P).$$

Substitute difference in eq.(A.1) (Lemma 1) between $\mu = 0$ and μ_0 for $\underline{\theta}$ in (A.42),

$$(A.43) \quad \dot{\underline{\theta}}(\mu_0, P) = (2\mu_0^2)^{-1} Q_h(\mu_0, P) A'_P (\underline{\lambda}(\mu_0, P) - \underline{\lambda}(0, P)).$$

Now we need to show that $\dot{\underline{\theta}}(\mu_0, P) = 0$. To establish this result, we need to study the behavior of the set of inequalities with positive Lagrange multipliers.

Consider any $j \in \mathcal{J}^{ineq}$. We know by Lemma 5 that $\underline{\lambda}_j(\mu, P)$ is continuous in μ . If $\underline{\lambda}_j(0, P) > 0$ then by continuity $\underline{\lambda}_j(\mu, P) > 0$ in some neighborhood $(0, \bar{\mu}_j(P)]$. If $\underline{\lambda}_j(0, P) = 0$ set $\bar{\mu}_j = 1$. Take $\bar{\mu}(P) \triangleq \min_{j \in \mathcal{J}^{ineq}} \bar{\mu}_j(P)$. WLOG suppose that $\mu_0 \in [0, \bar{\mu}(P)]$, so we get the inclusion

$$(A.44) \quad \mathcal{J}^+(0, P) \subseteq \mathcal{J}^+(\mu_0, P).$$

By definition of \mathcal{J}^h

$$(A.45) \quad \mathcal{J}^+(\mu_0, P) \subseteq \mathcal{J}^h(\mu_0, P).$$

By definition of the index matrices, inclusions (A.44) and (A.45) imply that

$$(A.46) \quad \underline{\lambda}(0, P) = (\mathbb{J}^h(\mu_0, P))' \mathbb{J}^h(\mu_0, P) \underline{\lambda}(0, P),$$

$$(A.47) \quad \underline{\lambda}(\mu_0, P) = (\mathbb{J}^h(\mu_0, P))' \mathbb{J}^h(\mu_0, P) \underline{\lambda}(\mu_0, P),$$

so

$$(A.48) \quad A'_P (\underline{\lambda}(\mu_0, P) - \underline{\lambda}(0, P)) = A'_h(\mu_0, P) (\mathbb{J}^h(\mu_0, P)) (\underline{\lambda}(\mu_0, P) - \underline{\lambda}(0, P)).$$

Since by definition $Q_h(\mu_0, P) A_h(\mu_0, P) = 0$ and Q_h is a symmetric matrix, equation (A.43) implies $\dot{\underline{\theta}}(\mu_0, P) = 0$. By Lemma 5 the single valued function $\underline{\theta}(\mu, P)$ is continuous for $\mu > 0$. So the r.h.s. directional derivative being equal to zero implies that $\underline{\theta}(\mu_0, P) = \underline{\theta}(\bar{\mu}(P), P)$ for any $\mu_0 \in (0, \bar{\mu}(P)]$. *Q.E.D.*

REMARK 9 Equation (A.1) from Lemma 1 with $\mu = 0$ and $\mu = \mu_0$ also implies

$$\underline{\lambda}(\mu_0, P) = \underline{\lambda}(0, P) - 2\mu_0 \underline{\theta}'(\bar{\mu}(P), P) A_h^\dagger(\mu_0, P) \mathbb{J}^h(\mu_0, P).$$

This implies that $\underline{\lambda}(\mu, P)$ is Lipschitz at $\mu = 0$. By Lemma 4 the Lipschitz constant can be taken equal to $2C_\Lambda$. So for any $j \in \mathcal{J}^{ineq}$ with $\underline{\lambda}_j(0, P) > 0$, we can take $\bar{\mu}_j(P) = \underline{\lambda}_j(0, P) / 2C_\Lambda$. On a top of that, Lemma 6 implies that $\underline{\lambda}(\mu, P)$ is Lipschitz in μ with the same constant for any $\mu \in [0, 1/2]$

PROOF OF THEOREM 1: Take any $\theta^* \in \underline{\theta}(0, P)$. Since θ^* is a feasible point of Program (3.2),

$$(A.49) \quad \underline{\theta}_1(\mu, P) + \mu \|\underline{\theta}(\mu, P)\|^2 \leq \theta_1^* + \mu \|\theta^*\|^2.$$

By definition $\underline{v}(P) = \theta_1^*$ which immediately implies the l.h.s. inequality in (3.13).

The first coordinate $\underline{\theta}_1(\mu, P)$ is increasing in μ while the norm $\|\underline{\theta}(\mu, P)\|^2$ is decreasing in μ . Since $\kappa \geq \mu \geq 0$, we get

$$(A.50) \quad \|\underline{\theta}(\mu, P)\|^2 \geq \|\underline{\theta}(\kappa, P)\|^2,$$

$$(A.51) \quad \underline{v}(P) \leq \underline{\theta}_1(\mu, P).$$

These two inequalities imply the r.h.s. inequality in (3.13),

$$(A.52) \quad \underline{v}(P) \leq \underline{\theta}_1(\mu, P) + \mu (\|\underline{\theta}(\mu, P)\|^2 - \|\underline{\theta}(\kappa, P)\|^2) = \underline{v}^{in}(\mu, \kappa, P).$$

The remaining part of the Theorem's assertion follows from Lemma 8.

Q.E.D.

A.4. Proof of Theorem 2

PROPOSITION 2 Suppose that $\mathbb{E}_P |\xi|^{1+\epsilon} \leq \infty$ for some $\epsilon > 0$. Then for any $r > 0$

$$\mathbb{E}_P [|\xi| I\{|\xi| \geq r\}] \leq \mathbb{E}_P |\xi|^{1+\epsilon} / r^\epsilon.$$

PROOF: The result follows from the monotonicity of integrals. Q.E.D.

Let $\pi(P, Q)$ denote the Prohorov distance between probability laws on P and Q , which induces the weak topology (see p. 456 in [van der Vaart and Wellner \(1996\)](#)). Let

$$G_n(P) \triangleq \sqrt{n} \left(\text{vec} \left(\frac{1}{n} \sum_{i=1}^n w_i \right) - \text{vec}(\mathbb{E}_P W) \right).$$

LEMMA 9 Consider \mathcal{P} , a class of distributions satisfying Assumptions 6-7, and any $\epsilon > 0$. Then

$$(A.53) \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P \left(\sup_{m \geq n} \left\| \frac{1}{m} \sum_{i=1}^m w_i - \mathbb{E}_P W \right\| \geq \epsilon \right) = 0,$$

$$(A.54) \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P \left(\sup_{m \geq n} \left\| \frac{1}{m} \sum_{i=1}^m w_i \otimes w_i - \mathbb{E}_P [W \otimes W] \right\| \geq \epsilon \right) = 0,$$

$$(A.55) \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \pi(G_n(P), N(0, \Omega_P)) = 0,$$

where $\Omega_P = \text{Cov}_P(\text{vec}(W))$.

PROOF: Consider any combination of indices r, ℓ, j, m . Assumption 7 together with the Schwarz inequality implies,

$$(A.56) \quad \mathbb{E}_P |W_{r,\ell} W_{j,m}|^{1+\epsilon/2} \leq (\mathbb{E}_P |W_{r,\ell}|^{2+\epsilon} \mathbb{E}_P |W_{j,m}|^{2+\epsilon})^{1/2} \leq \bar{M}.$$

So the random variables $|W_{r,\ell}|$ and $|W_{r,\ell} W_{j,m}|$ have correspondingly finite $1 + \epsilon/2$ and $2 + \epsilon$ moments. The bound (A.56) on the moments is independent of $P \in \mathcal{P}$, so these random variables are uniformly integrable on \mathcal{P} by Proposition 2. The limits (A.53) and (A.54) follow immediately from Proposition A.5.1 in [van der Vaart and Wellner \(1996\)](#). The result (A.55) follows from Proposition A.5.2 in the same book. Q.E.D.

LEMMA 10 Suppose that $P \in \mathcal{P}$. Then \mathbb{P}_n satisfies Assumptions 1 and 3 with probability approaching 1 uniformly in $P \in \mathcal{P}$.

PROOF: Consider any $P \in \mathcal{P}$. Since such a P satisfies Assumption 1, there exists a set of constraints \mathcal{J} with $|\mathcal{J}| = d$ and containing \mathcal{J}^{eq} such that

$$(A.57) \quad \theta_P^{\mathcal{J}} \triangleq ((A_P^{\mathcal{J}})' A_P^{\mathcal{J}})^{-1} A_P^{\mathcal{J}} b_P^{\mathcal{J}} \in \Theta(P),$$

where $A_P^{\mathcal{J}} \triangleq \mathbb{J} A_P$ and $b_P^{\mathcal{J}} \triangleq \mathbb{J} b_P$. By Assumption 5 for all $j \in \mathcal{J}^{ineq} \setminus \mathcal{J}$ we have

$$(A.58) \quad e_j(A_P \theta_P^{\mathcal{J}} - b_P) \geq \underline{s}.$$

The function $(A_P, b_P) = \mathbb{E}_P W$ is uniformly continuous on \mathcal{P} since this class of measures is uniformly integrable, as was shown in the proof of Lemma 9. The function $\theta_P^{\mathcal{J}}$ is uniformly continuous on \mathcal{P} since the product matrix $(A_P^{\mathcal{J}})' A_P^{\mathcal{J}}$ has eigenvalues uniformly bounded from below by η^2 . By Lemma 9 and the continuous mapping theorem,

$$(A.59) \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P \left\{ \inf_{m \geq n; j \in \mathcal{J}^{ineq} \setminus \mathcal{J}} \{e_j(\hat{A}_m \hat{\theta}_m^{\mathcal{J}} - \hat{b}_m)\} < \underline{s}/2 \right\} = 0,$$

where \hat{A}_m and \hat{b}_m are sample analog estimators of A_P and b_P based on a sample of size m ; $\hat{\theta}_m^{\mathcal{J}} \triangleq ((\hat{A}_m^{\mathcal{J}})' \hat{A}_m^{\mathcal{J}})^{-1} (\hat{A}_m^{\mathcal{J}})' \hat{b}_m^{\mathcal{J}}$. This result implies that $\Theta(\mathbb{P}_n)$ contains at least one element, $\hat{\theta}_n^{\mathcal{J}}$, with probability approaching 1 uniformly in $P \in \mathcal{P}$ as $n \rightarrow \infty$.

Analogously, the continuous mapping theorem implies

$$(A.60) \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P\left\{ \inf_{m \geq n} \eta(\mathbb{P}_m) < \underline{\eta}/2, \inf_{m \geq n} s(\mathbb{P}_m) < \underline{s}/2 \right\} = 0.$$

The result follows from Lemma 2.

Q.E.D.

LEMMA 11 *Suppose Assumption 6 holds and that $P \in \mathcal{P}$ and that $\mu_n \rightarrow 0$. Then*

$$(A.61) \quad \underline{v}(\mu_n, \mathbb{P}_n) = \underline{v}(\mu_n, P) + \frac{1}{n} \sum_{i=1}^n \underline{\lambda}(\mu_n, P)' g(w_i, \underline{\theta}(\mu_n, P)) + O_{\mathcal{P}}\left(\frac{1}{\mu_n n}\right).$$

PROOF: First note that by Lemma 10 the random variables $\underline{\theta}(\mu_n, \mathbb{P}_n)$ and $\underline{v}(\mu_n, \mathbb{P}_n)$ are well defined with probability approaching 1 uniformly in $P \in \mathcal{P}$. Consider $t = 1/\sqrt{n}$ and h such that $h_W = \sqrt{n}(\frac{1}{n} \sum_{i=1}^n w_i - \mathbb{E}_P W)$ and $h_\mu = 0$. The sequence of perturbations satisfies $\sqrt{n}(\frac{1}{n} \sum_{i=1}^n w_i - \mathbb{E}_P W) = O_{\mathcal{P}}(1)$ (i.e. it has uniformly tight measures) by (A.55) in Lemma 9 and the fact that $\|\Omega_P\| \leq \bar{M}$ (by Jensen's inequality). The assertion of the Lemma follows from equation (A.32) in Lemma 7. Indeed, by Lemma 9

$$(A.62) \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P\left(\sup_{m \geq n} \left\| \frac{1}{m} \sum_{i=1}^m w_i - \mathbb{E}_P W \right\| \geq \delta \right) = 0,$$

so the perturbation is small enough to preserve the results of Lemma 7, i.e. $\|th_W\| < \delta$, with probability approaching 1 uniformly as sample size grows.

Q.E.D.

Let's introduce the following definitions

$$\begin{aligned} \Sigma(\theta) &\triangleq \mathbb{E}_P [g(W, \theta)g(W, \theta)'] - m(\theta, P)m(\theta, P)', \\ \hat{\Sigma}_n(\theta) &\triangleq \frac{1}{n} \sum_{i=1}^n g(w_i, \theta)g(w_i, \theta)' - \frac{1}{n} \sum_{i=1}^n g(w_i, \theta) \frac{1}{n} \sum_{i=1}^n g(w_i, \theta)'. \end{aligned}$$

PROOF OF THEOREM 2: First note that Lemma 10 the random variables $\underline{\theta}(\mu_n, \mathbb{P}_n)$ and $\underline{\lambda}(\mu_n, \mathbb{P}_n)$ are well defined with probability approaching 1 uniformly in $P \in \mathcal{P}$. Consider t and h as in the proof of Lemma 11. Equation (3.18) follows from Lemma 9, Lemma 11, and Slutsky's theorem. The results (3.19) and (3.20) follow from Lemma 7. Finally, by the triangular inequality

$$(A.63) \quad \begin{aligned} \left| \underline{\sigma}(\mu, \mathbb{P}_n)^2 - \underline{\sigma}(\mu, P)^2 \right| &= \left| \underline{\lambda}(\mu, \mathbb{P}_n)' \hat{\Sigma}_n(\underline{\theta}(\mu, \mathbb{P}_n)) \underline{\lambda}(\mu, \mathbb{P}_n) - \underline{\lambda}(\mu, P)' \Sigma(\underline{\theta}(\mu, P)) \underline{\lambda}(\mu, P) \right| \leq \\ &\quad \left| \underline{\lambda}(\mu, \mathbb{P}_n)' \hat{\Sigma}_n(\underline{\theta}(\mu, \mathbb{P}_n)) \underline{\lambda}(\mu, \mathbb{P}_n) - \underline{\lambda}(\mu, P)' \hat{\Sigma}_n(\underline{\theta}(\mu, P)) \underline{\lambda}(\mu, P) \right| \\ &\quad + \left| \underline{\lambda}(\mu, P)' \hat{\Sigma}_n(\underline{\theta}(\mu, P)) \underline{\lambda}(\mu, P) - \underline{\lambda}(\mu, P)' \Sigma(\underline{\theta}(\mu, P)) \underline{\lambda}(\mu, P) \right|. \end{aligned}$$

Together with (3.19), (3.20) and Lemmas 4 and 9, it implies (3.21).

Q.E.D.

A.5. Proof of Theorem 3

LEMMA 12 *Suppose that Assumption 6 holds then for any $\epsilon > 0$ there exists $R \geq 0$ such that for any n the following uniform bound holds*

$$(A.64) \quad \sup_{P \in \mathcal{P}} P(\sqrt{n} \|\theta^*(\mathbb{P}_n) - \theta^*(P)\| \geq R) \leq \epsilon.$$

PROOF: The proof is based on the delta method applied to $\theta^*(P) = \theta^*(A_P, b_P)$, a composition of directionally differentiable functions. (Since the space of A_P, b_P is finite dimensional, Gâteaux and Hadamard directional derivatives coincide.)

First note that $\underline{v}(P) = \underline{v}(A_P, b_P)$ is directionally differentiable function. To see it, consider the minimax representation, which is valid since P satisfies Assumption 1,

$$(A.65) \quad \underline{v}(A_P, b_P) = \min_{\theta \in \Theta} \max_{\lambda \in \mathbb{R}^p \times \mathbb{R}_+^{k-p}, \|\lambda\| \leq C_\Lambda} \{\theta_1 + \lambda'(A_P \theta - b_P)\}.$$

Here we also used Lemma 4 to bound λ and make the domain compact. By Theorem 7.28 from Shapiro et al. (2014) for any direction $(h_A, h_b) \in \mathbb{R}^{k \times (d+1)}$ we have

$$(A.66) \quad \dot{\underline{v}}(P|h_A, h_b) \triangleq \lim_{t \rightarrow 0^+} \frac{\underline{v}(A_P + th_A, b_P + th_b) - \underline{v}(A_P, b_P)}{t} = \min_{\theta \in \underline{\theta}(P)} \{\underline{\lambda}(P)'(h_A \theta - h_b)\}.$$

Similarly,

$$(A.67) \quad \theta_i^\pm(A_P, b_P) = \left| \min_{\theta \in \Theta} \max_{\gamma \in \mathbb{R}^p \times \mathbb{R}_+^{k-p}, \|\gamma\| \leq C_\Lambda, \nu \geq 0} \{\pm \theta_i + \gamma'(A_P \theta - b_P) + \nu(\theta_1 - \underline{v}(P) - \mu_n)\} \right|.$$

Once again, by Theorem 7.28 from Shapiro et al. (2014) for any direction $(h_A, h_b) \in \mathbb{R}^{k \times (d+1)}$ we have

$$(A.68) \quad \dot{\theta}_i^\pm(P|h_A, h_b) = \min_{\theta \in \arg\min(A.67)} \{\underline{\gamma}(P)'(h_A \theta - h_b)\} + \underline{v}(P) \min_{\vartheta \in \underline{\theta}(P)} \{\underline{\lambda}(P)'(h_A \vartheta - h_b)\}.$$

Here we used Proposition 2.47 from Bonnans and Shapiro (2000), the chain rule for directional derivatives. By the same proposition, the vector with maximal components $\theta^*(P)$ is directionally differentiable.

Delta method (Theorem 7.67 in Shapiro et al. (2014)) and Lemma 9 imply

$$(A.69) \quad \sqrt{n}(\theta^*(\mathbb{P}_n) - \theta^*(P)) = \dot{\theta}^*(P|G_n(P)) + o_p(1).$$

By compactness of \mathcal{P} , the vanishing term $o_p(1)$ is bounded in probability uniformly in $P \in \mathcal{P}$. It is a routine to compute a uniform bound on the directional derivative, $\|\dot{\theta}^*(P|h_A, h_b)\| \leq L_{\mathcal{P}} < \infty$. The constant $L_{\mathcal{P}}$ provides a uniform asymptotic bound

$$(A.70) \quad \sqrt{n} \|\theta^*(\mathbb{P}_n) - \theta^*(P)\| \leq L_{\mathcal{P}} \|G_n(P)\| + O_{\mathcal{P}}(1).$$

By Lemma 9, this representation implies (A.64).

Q.E.D.

PROOF OF THEOREM 3: The proof is analogous for all CI. Consider, for example, $\text{CB}_{\alpha, n, \mathcal{P}}$. Pick an arbitrary measure $P \in \mathcal{P}$. Consider any $\delta > 0$ such that $z_{1-\alpha} \sigma^0 > 2\delta > 0$. Then by Lemma 12 and Theorem 2 correspondingly there exist $n(\delta, \epsilon)$ such that for any $n > n(\delta, \epsilon)$

$$(A.71) \quad \inf_{P \in \mathcal{P}} P\{\mu_n \sqrt{n} \left| \|\theta^*(\mathbb{P}_n)\|^2 - \|\theta^*(P)\|^2 \right| \leq \delta\} \geq 1 - \epsilon,$$

$$(A.72) \quad \inf_{P \in \mathcal{P}} P\{z_{1-\alpha} |\underline{\sigma}(\mu_n, \mathbb{P}_n) - \underline{\sigma}(\mu_n, P)| \leq \delta\} \geq 1 - \epsilon,$$

$$(A.73) \quad \sup_{P \in \mathcal{P}} \rho_n(P) \leq \delta$$

By construction $\|\theta^*(P)\| \geq \min_{\theta \in \underline{\theta}(P)} \|\theta\|$, so by Theorem 1

$$(A.74) \quad \underline{\theta}_1(\mu_n, P) + \mu_n(\|\underline{\theta}(\mu_n, P)\|^2 - \|\theta^*(P)\|^2) \leq \underline{v}(P).$$

Using this bound,

$$\begin{aligned} & P \left\{ \underline{v}(\mu_n, \mathbb{P}_n) - \mu_n \|\theta^*(\mathbb{P}_n)\|^2 - \underline{\sigma}(\mu_n, \mathbb{P}_n) z_{1-\alpha} n^{-1/2} \leq \underline{v}(P) \right\} \geq \\ & P \left\{ \sqrt{n}(\underline{v}(\mu_n, \mathbb{P}_n) - \underline{v}(\mu_n, P)) - \mu_n \sqrt{n}(\|\theta^*(\mathbb{P}_n)\|^2 - \|\theta^*(P)\|^2) \leq \underline{\sigma}(\mu_n, \mathbb{P}_n) z_{1-\alpha} \right\} \geq \text{(by (A.30))} \\ & P \left\{ \sqrt{n}(\underline{v}(\mu_n, \mathbb{P}_n) - \underline{v}(\mu_n, P)) \leq \underline{\sigma}(\mu_n, P) z_{1-\alpha} - 2\delta \right\} (1 - \epsilon)^2 \geq \\ & \quad (\Phi(z_{1-\alpha} - \frac{2\delta}{\sigma_0}) - \delta)(1 - \epsilon)^2 \end{aligned}$$

Since P is arbitrary, for any $n > n(\delta, \epsilon)$

$$(A.75) \quad \inf_{P \in \mathcal{P}} \min_{\theta \in \underline{\theta}(P)} P(\theta_1 \in \text{CB}_{\alpha, n, \mathcal{P}}) \geq (\Phi(z_{1-\alpha} - \frac{2\delta}{\sigma_0}) - \delta)(1 - \epsilon)^2.$$

Hence,

$$(A.76) \quad \liminf_{n \rightarrow \infty} \inf_{P_n \in \mathcal{P}} \min_{\theta \in \underline{\theta}(P_n)} P_n(\theta_1 \in \text{CB}_{\alpha, n, \mathcal{P}}) \geq (1 - \alpha).$$

Q.E.D.

APPENDIX B: ADDITIONAL RESULTS

B.1. Point-wise valid and non-conservative confidence intervals

How conservative is $\text{CB}_{\alpha, n, \mathcal{P}}$ in the nonregular case? We can evaluate it by comparing its average length with that of a confidence bound that has exact point-wise asymptotic coverage probability in a Monte Carlo study. As discussed in Section 2.3, the existing (point-wise valid) nonconservative procedures for inference on support functions in the nonregular case are based on simulation methods (for example, bootstrap for directionally differentiable functions). These procedures can be computationally costly in high-dimensional settings. Theorems 1 and 2 suggest alternative analytical confidence bounds on $\underline{v}(P)$ using the following inner estimator:

$$(B.1) \quad \underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) \triangleq \underline{v}(\mu_n, \mathbb{P}_n) - \mu_n \|\underline{\theta}(\kappa_n, \mathbb{P}_n)\|^2.$$

Theorem 1 guarantees that the corresponding population counterpart $\underline{v}^{in}(\mu, \kappa, P)$ coincides with $\underline{v}(P)$ if κ_n and μ_n are sufficiently small. Theorem 2 implies that if κ_n converges to zero slower than μ_n and both converge slower than $1/\sqrt{n}$, then $\underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n)$ is an asymptotically normal estimator with variance $\underline{\sigma}(\mu_n, P)$. This property suggests the following point-wise-valid CSs:

$$(B.2) \quad \begin{cases} \text{CB}_{\alpha, n} &= [\underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) - z_{1-\alpha} n^{-1/2} \underline{\sigma}(\mu_n, \mathbb{P}_n), \infty), \\ \text{CI}_{\alpha, n}^{\theta_1} &= [\underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) - z_{1-\alpha} n^{-1/2} \underline{\sigma}(\mu_n, \mathbb{P}_n); \bar{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) + z_{1-\alpha} n^{-1/2} \bar{\sigma}(\mu_n; \mathbb{P}_n)], \\ \text{CI}_{\alpha, n}^S &= \text{CI}_{\alpha/2, n}^{\theta_1}, \end{cases}$$

These CSs' properties are summarized in Theorem 4.

THEOREM 4 Suppose that Assumptions 1 and 4–7 hold (i.e. $P \in \mathcal{P}$), $0 < \alpha < 1/2$, $\kappa_n \rightarrow 0$ and $\mu_n \rightarrow 0$ are such that $\mu_n \sqrt{n} \rightarrow \infty$. Moreover, suppose that $\lim_{n \rightarrow \infty} \underline{\sigma}^2(\mu_n, P) > 0$ and $\lim_{n \rightarrow \infty} \bar{\sigma}^2(\mu_n, P) > 0$. Then

$$\lim_{n \rightarrow \infty} P(\mathcal{S}(P) \subset CB_{\alpha, n}) = \liminf_{n \rightarrow \infty} \inf_{\theta \in \Theta(P)} P(\theta_1 \in CB_{\alpha, n}) \geq 1 - \alpha,$$

$$\lim_{n \rightarrow \infty} P(\mathcal{S}(P) \subset CI_{\alpha, n}^{\mathcal{S}}) \geq 1 - \alpha, \quad \liminf_{n \rightarrow \infty} \inf_{\theta \in \Theta(P)} P(\theta_1 \in CI_{\alpha, n}^{\mathcal{S}}) \geq 1 - \alpha.$$

If in addition $\mu_n/\kappa_n \rightarrow 0$, then,

$$\lim_{n \rightarrow \infty} P(\mathcal{S}(P) \subset CB_{\alpha, n}) = \liminf_{n \rightarrow \infty} \inf_{\theta \in \Theta(P)} P(\theta_1 \in CB_{\alpha, n}) = 1 - \alpha.$$

If further, the model has no equality constraints—that is, if $p = 0$ —then

$$(B.3) \quad \liminf_{n \rightarrow \infty} \inf_{\theta \in \Theta(P)} P(\theta_1 \in CI_{\alpha, n}^{\theta_1}) = 1 - \alpha.$$

PROOF: See Appendix B.2.

Q.E.D.

The confidence band $CB_{\alpha, n}$ and confidence interval $CI_{\alpha, n}^{\theta_1}$ are asymptotically nonconservative for a fixed DGP satisfying the assumptions of the theorem; that is, they have coverage of exactly $1 - \alpha$. If $p > 0$, θ_1 can be point-identified if there are equality constraints in the model that are orthogonal to e_1 . In this case I recommend the Bonferroni-type CS $CI_{\alpha, n}^{\mathcal{S}}$, which remains valid under point identification. The shorter $CI_{\alpha, n}^{\theta_1}$ (compare [Imbens and Manski \(2004\)](#)) is valid only if θ_1 is not point-identified.

It is interesting to compare the uniformly valid confidence bound $CB_{\alpha, n, \mathcal{P}}$ with its nonconservative point-wise counterpart $CB_{\alpha, n}$. In the regular case, the two sets coincide asymptotically and have the same exact coverage of the support function $\underline{v}(P)$. In the presence of flat face (the nonregular case), $CB_{\alpha, n}$ will be shorter than $CB_{\alpha, n, \mathcal{P}}$ with probability 1. Moreover, the confidence level of $CB_{\alpha, n}$ will remain asymptotically exact (point-wise). Despite this, I recommend using the uniformly valid bound $CB_{\alpha, n, \mathcal{P}}$ since it would have better control of the confidence level in small samples (because of its validity under drifting DGP sequences; see Remark 10).

REMARK 10 Theorem 4 provides asymptotic coverage probability for a given DGP with measure P . The size of the sample required to achieve the nominal coverage of $1 - \alpha$ with a given precision in this result can depend on P through $\bar{\mu}(P)$ as defined in Theorem 1. The cutoff $\bar{\mu}(P)$ can be arbitrarily close to zero. It is possible to construct an example in which the sequence of measures P_n meets the assumptions of Theorem 4 but the asymptotic bias is larger than $1/\sqrt{n}$:

$$(B.4) \quad \sqrt{n} \mu_n (\|\underline{\theta}(\mu_n, P_n)\| - \|\underline{\theta}(\kappa_n, P_n)\|) \rightarrow +\infty.$$

In other words, there are examples of DGP with a measure P and some $\epsilon > 0$ such that for any n it is possible to find a measure Q in a neighborhood of P with

$$(B.5) \quad Q(\mathcal{S}(P) \subset CB_{\alpha, n}) < 1 - \alpha - \epsilon.$$

In practical terms, this means that the large-sample theory with a fixed P may provide a poor approximation for the true coverage probability of the point-wise-valid confidence set $CB_{\alpha, n}$. (A similar concern applies to other existing point-wise-valid inference procedures.)

B.2. Proof of Theorem 4

PROOF OF THEOREM 4: STEP 1. First, suppose that $\mu_n/\kappa_n \rightarrow 0$. Consider

$$(B.6) \quad \underline{\zeta}_n \triangleq \frac{\sqrt{n} \underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) - \underline{v}^{in}(\mu_n, \kappa_n, P) + \underline{v}^{in}(\mu_n, \kappa_n, P) - \underline{v}(P)}{\underline{\sigma}(\mu, \mathbb{P}_n)}$$

Since $\mu_n/\kappa_n \rightarrow 0$ and $\mu_n \rightarrow 0$, for all n large enough, such that $\mu_n \leq \kappa_n \leq \bar{\mu}(P)$, by Theorem 1 we get

$$(B.7) \quad \underline{v}^{in}(\mu_n, \kappa_n, P) = \underline{v}(P).$$

By Theorem 2 we get

$$(B.8) \quad \mu_n \sqrt{n} \left(\|\underline{\theta}(\kappa_n, \mathbb{P}_n)\|^2 - \|\underline{\theta}(\kappa_n, P)\|^2 \right) = \frac{\mu_n}{\kappa_n} O_p(1) = o_p(1).$$

By Lemma 5, $\underline{\theta}(\mu, P)$ and $\underline{\lambda}(\mu, P)$ are continuous for $\mu > 0$. The matrix function $\underline{\Sigma}(\theta, P)$ is continuous in θ and thus $\underline{\sigma}(\mu, P)$ is continuous in μ for $\mu > 0$. So the limit $\lim_{n \rightarrow \infty} \underline{\sigma}^2(\mu_n, P)$ exists and belongs to the set $\underline{\sigma}^2(0, P)$

which by assumptions implies $\lim_{n \rightarrow \infty} \underline{\sigma}^2(\mu_n, P) > 0$. Result (3.18) in Theorem 2 together with (B.7) and (B.8) imply by Slutsky's theorem that $\underline{\zeta}_n$ converges in distribution to $N(0, 1)$.

STEP 2. Consider the one-sided confidence band $\text{CB}_{\alpha, n}$.

$$\begin{aligned} & \lim_{n \rightarrow \infty} P \{ \mathcal{S}(P) \subset \text{CB}_{\alpha, n} \} \\ = & \lim_{n \rightarrow \infty} P \left\{ \underline{v}(P) \geq \underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) - \underline{\sigma}(\mu_n, \mathbb{P}_n) z_{1-\alpha} n^{-1/2} \right\} \\ = & \lim_{n \rightarrow \infty} P \left\{ \underline{\zeta}_n \leq z_{1-\alpha} \right\} \\ = & \Phi(z_{1-\alpha}) = 1 - \alpha. \end{aligned}$$

Proof for $\text{CI}_{\alpha, n}^S$ follows immediately from the Bonferroni inequality. Finally, consider the case $p = 0$. Then by Lemma 2 and Assumption 1, $\underline{v}(0, P) < -\bar{v}(0, P)$. So

$$\begin{aligned} & \lim_{n \rightarrow \infty} \min_{\theta \in \Theta(P)} P \left(\theta \in \text{CI}_{\alpha, n}^{\theta_1} \right) = \\ \min & \left\{ \lim_{n \rightarrow \infty} P \left\{ \underline{v}(\mu_n, \mathbb{P}_n) - \mu_n \|\underline{\theta}(\kappa_n, \mathbb{P}_n)\|^2 - \underline{\sigma}(\mu_n, \mathbb{P}_n) z_{1-\alpha} n^{-1/2} \leq \underline{v}(P) \right\}, 1, \dots \right. \\ & \left. \lim_{n \rightarrow \infty} P \left\{ -\bar{v}(\mu_n; \mathbb{P}_n) + \mu_n \|\bar{\theta}(\kappa_n, \mathbb{P}_n)\|^2 + z_{1-\alpha} \bar{\sigma}(\mu_n; \mathbb{P}_n) n^{-1/2} \geq \bar{v}(P) \right\} \right\} = \\ \text{(B.9)} \quad & \min \left\{ \lim_{n \rightarrow \infty} P \{ \mathcal{S}(P) \subset \text{CB}_{\alpha, n} \}, 1, \lim_{n \rightarrow \infty} P \{ \mathcal{S}(P) \subset \text{CB}_{\alpha, n}^R \} \right\} = \\ \text{(B.10)} \quad & \min \{ 1 - \alpha, 1, 1 - \alpha \} = 1 - \alpha. \end{aligned}$$

To understand the second equation, consider the following argument. Suppose that $\theta \in \Theta(P)$ is such that $\underline{v}(P) < \theta_1 < \bar{v}(P)$. Then such θ_1 will be covered with probability 1 since $\text{CI}_{\alpha, n}^{\theta}$ is the intersection of $\text{CB}_{\alpha, n}$ and $\text{CB}_{\alpha, n}^R$ which cover correspondingly $\underline{v}(P)$ and $\bar{v}(P)$.

STEP 3. Finally, suppose that $\mu_n/\kappa_n \rightarrow 0$ does not hold, i.e. sequence $\kappa_n \rightarrow 0$ can take any positive values. Consider any auxiliary sequence $\kappa_n^* \rightarrow 0$ such that $\mu_n/\kappa_n^* \rightarrow 0$ and $\kappa_n \leq \kappa_n^*$. By Step 2,

$$\lim_{n \rightarrow \infty} P \left\{ \underline{v}(P) \geq \underline{v}^{in}(\mu_n, \kappa_n^*, \mathbb{P}_n) - \underline{\sigma}(\mu_n, \mathbb{P}_n) z_{1-\alpha} n^{-1/2} \right\} = 1 - \alpha.$$

But $\underline{v}^{in}(\mu_n, \kappa, \mathbb{P}_n)$ is a non-increasing function in κ . Indeed, by (A.42) in proof of Lemma 8, the partial derivative $\|\underline{\theta}(\kappa, \mathbb{P}_n)\|^2$ w.r.t. κ at any point $\kappa_0 > 0$ is $2\underline{\theta}'(\kappa_0, \mathbb{P}_n)\underline{\theta}(\kappa_0, \mathbb{P}_n) \leq 0$. So $\|\underline{\theta}(\kappa, \mathbb{P}_n)\|^2$ is decreasing function of $\kappa > 0$.

The derivative at $\kappa = 0$ does not exist, so this case has to be considered separately. Take any $\theta^* \in \underline{\theta}(0, \mathbb{P}_n)$. Since θ^* is a feasible point of Program (3.2) with P replaced by \mathbb{P}_n ,

$$\text{(B.11)} \quad \underline{\theta}_1(\kappa, \mathbb{P}_n) + \kappa \|\underline{\theta}(\kappa, \mathbb{P}_n)\|^2 \leq \theta_1^* + \kappa \|\theta^*\|^2.$$

By Theorem 1, for $\kappa \leq \bar{\mu}(\mathbb{P}_n)$ we have $\underline{\theta}_1(\kappa, \mathbb{P}_n) = \theta_1^*$, so for such small values (B.11) becomes

$$\|\underline{\theta}(\kappa, \mathbb{P}_n)\|^2 \leq \|\theta^*\|^2.$$

This proves that $\|\underline{\theta}(\kappa, \mathbb{P}_n)\|^2$ is a non-increasing function of κ for all values $\kappa \geq 0$ including 0. So a.s. $\underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) \leq \underline{v}^{in}(\mu_n, \kappa_n^*, \mathbb{P}_n)$ and therefore

$$\begin{aligned} \lim_{n \rightarrow \infty} P \left\{ \underline{v}(P) \geq \underline{v}^{in}(\mu_n, \kappa_n, \mathbb{P}_n) - \underline{\sigma}(\mu_n, \mathbb{P}_n) z_{1-\alpha} n^{-1/2} \right\} & \geq \\ \lim_{n \rightarrow \infty} P \left\{ \underline{v}(P) \geq \underline{v}^{in}(\mu_n, \kappa_n^*, \mathbb{P}_n) - \underline{\sigma}(\mu_n, \mathbb{P}_n) z_{1-\alpha} n^{-1/2} \right\} & = 1 - \alpha. \end{aligned}$$

It means that if one uses a sequence κ_n that converges to 0 faster than μ_n , it makes all the confidence sets more conservative than the ones corresponding to κ_n^* . Q.E.D.

B.3. *A special case of over-identified moment conditions: Uniformly valid inference on intersection bounds*

The sufficient conditions for LICQ (Assumptions 4–5) considerably simplify the asymptotic analysis, but they can be violated in some interesting cases. In particular, the active inequality constraints can be linearly dependent. In those cases, the regularization approach can still be applied to the dual formulation of the linear program as in Equation (2.12). A full comprehensive analysis of dual regularization goes beyond the scope of this paper. In this section, I restrict my attention to a special case of the following linear program:

$$(B.12) \quad \underline{v}(P) = \min_{\theta_1 \in \mathbb{R}} \theta_1, \\ \text{s.t. } \theta_1 \geq \mathbb{E}_P W_j, \text{ for } j = 1, \dots, k.$$

This function appears in many econometric applications and has been studied in the literature (see (2.18) in Example 3; also see examples in Hall and Miller (2010), Chernozhukov et al. (2013)). The value of this program can be written explicitly as $\max_{j=1, \dots, k} \mathbb{E} W_j$ (also known as an *intersection-bound* parameter). Such a maximum over a finite set can in turn be equivalently represented as a linear program,

$$(B.13) \quad \underline{v}(P) = - \min_{\theta_1 \in \mathbb{R}, \lambda \in \mathbb{R}^k} \{-\theta_1\}, \\ \text{s.t. } \theta_1 - \sum_{j=1}^k \lambda_j \mathbb{E} W_j = 0, \quad \sum_{j=1}^k \lambda_j = 1, \quad \lambda_j \geq 0.$$

(Compare with the dual-linear-program representation (2.20).) The problem of overidentifying inequality constraints (multiple dual solutions) in formulation (B.12) becomes a “flat face” problem (multiple primal solutions) in formulation (B.13). Constraints of the program in (B.13) are almost surely linearly independent, so one can directly verify that Assumptions 1 and 4–5 are also satisfied. As a result, Theorems 1 and 2 can be applied to the following dual regularized formulation:

$$(B.14) \quad \underline{v}^{dual}(\mu_n, \mathbb{P}_n) = - \min_{\theta_1 \in \mathbb{R}, \lambda \in \mathbb{R}^k} \left\{ -\theta_1 + \mu_n(\theta_1^2 + \|\lambda\|^2) \right\}, \\ \text{s.t. } \theta_1 - \sum_{j=1}^k \lambda_j \mathbb{E} W_j = 0, \quad \sum_{j=1}^k \lambda_j = 1, \quad \lambda_j \geq 0.$$

This is completely analogous to (3.14) (with the exception of the negative sign in front of the minimum operator). Estimator $\underline{v}^{dual}(\mu_n, \mathbb{P}_n)$ has the following bias-corrected version:

$$(B.15) \quad \underline{v}^{dual, in}(\mu_n, \kappa_n, \mathbb{P}_n) = \underline{v}^{dual}(\mu_n, \mathbb{P}_n) + \mu_n(\underline{\theta}_1^2(\kappa_n, \mathbb{P}_n) + \|\underline{\lambda}(\kappa_n, \mathbb{P}_n)\|^2).$$

Here, $\underline{\theta}_1(\kappa_n, \mathbb{P}_n)$ and $\underline{\lambda}(\kappa_n, \mathbb{P}_n)$ are solutions to (B.14) with a tuning parameter κ_n instead of μ_n .

By implication of Theorem 1, the estimator $\underline{v}^{dual, in}(\mu_n, \kappa_n, \mathbb{P}_n)$ is either asymptotically unbiased for $\underline{v}(P)$ from (B.13) (in the regular case, in which only one inequality is binding) or biased downward (in the nonregular case, in which multiple inequalities are binding). As a result, $\underline{v}^{dual, in}(\mu_n, \kappa_n, \mathbb{P}_n)$ can be used in a delta-method CS for $\underline{v}(P)$,

$$(B.16) \quad \text{CB}_{\alpha, n, \mathcal{P}}^{dual} \triangleq \left[\underline{v}^{dual, in}(\mu_n, \kappa_n, \mathbb{P}_n) - z_{1-\alpha} n^{-1/2} \underline{\sigma}(\mu_n, \mathbb{P}_n), \infty \right),$$

where $\underline{\sigma}(\mu_n, \mathbb{P}_n)$ corresponds to the asymptotic variance estimator for $\underline{v}^{dual}(\mu_n, \mathbb{P}_n)$. Following the steps of the proof of Theorem 3, one can establish the correct uniform asymptotic coverage of $\underline{v}(P)$ with $\text{CB}_{\alpha, n, \mathcal{P}}^{dual}$. Moreover, in this case the coverage will be exactly $1 - \alpha$ for any fixed DGP even in the nonregular case (analogous to the results in Theorem 4).

To conclude, the regularization approach proposed can be applied to systems of moment-inequality conditions that violate Assumptions 4 and 5 after an appropriate dual reformulation of the program. A comprehensive study of this dual regularization is a topic for another paper. The intersection-bounds problem is a specific empirically relevant example in which the method can be applied with minimal modification to the convex reformulation in Equation (B.13).

APPENDIX C: DISCUSSION OF COMPUTATIONAL PROPERTIES

C.1. *Fast convergence to a minimum*

The existing uniform methods of AS, BCS, and KMS are based on standardized moment conditions that are nonconvex even if the original inequalities are affine in θ . Example 2 in the previous section illustrates this feature.

The estimator $\hat{\theta}(\mu_n, \mathbb{P}_n)$ is a solution to a strictly convex quadratic program for any affine-moment-inequality model. For convex programs, the Karush–Kuhn–Tucker (KKT) conditions provide necessary and sufficient conditions for the global optimum (see Lemma 3 in Appendix). Moreover, convex quadratic programs can be solved using interior-point algorithms with a polynomial rate of convergence. (See, for example, Ye and Tse (1989).) This strict convexity gives a dramatically faster rate of obtaining the optimum than the ones used in BCS and KMS. These methods are based on nonconvex-constraint optimization problems, which are NP-hard. Section 4 compares computational time in specific examples.

C.2. Uniqueness of a global optimum

The KKT system for strictly convex optimization problems has a unique solution. The number of KKT points of the optimization problems in the KMS, BCS, and AS procedures in affine moment inequality models can be large and typically grows exponentially with the dimension d and number of inequalities k . The following example illustrates this point.

EXAMPLE 6 Consider a set of moment inequalities with coefficients that have expectation

$$\mathbb{E}_P W = \begin{pmatrix} -I_d & -\iota \\ I_d & -\iota \end{pmatrix}.$$

Suppose that components of W are independent and have the same variance s^2 . $\Theta(P)$ is a box $[-1, 1]^d$. The standardized moment conditions take the form

$$(C.1) \quad \frac{\pm\theta_j + 1}{s\sqrt{1 + \|\theta\|^2}} \leq 0, \quad j = 1, \dots, d.$$

The KMS procedure adds slack $c(\theta)$ to the right-hand side of every standardized moment inequality. Consider, for example, $j = 1$, where

$$(C.2) \quad \theta_1 \geq 1 - c(\theta) s\sqrt{1 + \|\theta\|^2}.$$




FIGURE 7.— The identified set and the corresponding optimization domain of a KMS-type procedure for $d = 3$ in Example 6.

The slack function $c(\theta)$ is computed using a resampling on a grid of points. Assume, for simplicity, that $c(\theta)$ is a constant—for example, one provided using the Bonferroni approach. Figure 7 shows the identified set and the corresponding expansion with $c(\theta) = \text{const}$. The optimization domain of the E-A-M algorithm in KMS is similar to the nonconvex set on the right side of Figure 7. Every vertex of the $[-1, 1]^d$ with $\theta_1 = -1$ corresponds to an isolated local minimum of the optimization procedure in KMS. Correspondingly, the number of local minima grows exponentially with the dimension d . For example, the number of local minima for $d = 10$ is 512. The growth in the number of local optima is even faster in models with more than two inequalities per coordinate. \square

Multiplicity of KKT points both makes the procedures of KMS, AS, and BCS computationally costly and does not guarantee convergence to a global optimum for large d .

C.3. Implications for validity of multiplier bootstrap

The multiplier bootstrap provides a numerical way to implement delta-method inference. Main advantage of this procedure when compared to conventional analytical formulas is the fact that one can use simulations to obtain joint variance-covariance matrix of several asymptotically Gaussian estimators (for example, bounds on projections of an identified set for a set of directions) and/or directly make draws from their joint limiting Gaussian distribution without necessity to derive the corresponding matrix formulas explicitly. In comparison to the standard non-parametric bootstrap methods, multiplier bootstrap eliminates the need to recompute estimators based on potentially costly optimization routines.¹⁸ The multiplier-bootstrap approach would be particularly appealing in subvector inference on more than one component as discussed in Section 3.3.2.

The proposed estimators of the regularized support functions for fixed direction have a Bahadur-Kiefer representation with explicit influence functions given in Equation (3.16). One can use this property to justify multiplier bootstrap for inference on the support of the identified set along the lines of proving validity of the delta-method. Proof of Theorem 3 can be adapted by replacing estimators of standard deviation $\underline{\sigma}(\mu_n, \mathbb{P}_n)$ with the multiplier bootstrap estimator of asymptotic standard deviation of $\sqrt{n}(\underline{v}(\mu_n, \mathbb{P}_n) - \underline{v}(\mu_n, P))$. Indeed, Theorem 2 shows that the unknown parameters in the influence function are (uniformly over \mathcal{P}) consistently estimated by their sample analog. Multiplier bootstrap procedure could also be generalized to obtain non-standard critical values of sup t statistics as discussed earlier in Subsection 3.3.2. I leave full analysis of such an extension for future work.

C.4. Choice of solvers


The point-wise CIs in (B.2) can be computed using any Newton-type optimization software that provides accurate Lagrange multipliers. I use the *fmincon* function of MATLAB software. I recommend using the active-set or SQP option since the interior-point option does not provide accurate Lagrange multipliers. The estimator $\|\theta^*(\mathbb{P}_n)\|^2$, which enters the uniformly valid CS described in Theorem 3, is based on $2d$ linear programs. Linear programs typically scale very well.

A benchmarking of results for the state-of-the-art commercial LP solvers can be found, for example, at <http://plato.asu.edu/bench.html>. The commercial solvers can tackle LP with tens of thousands of constraints and variables in a matter of minutes. Matlab's *linprog* solver tends to underperform in the time comparison.

¹⁸For example, FH and KMS solve mathematical programs repeatedly for every bootstrap sample. One of the recent papers in the moment inequality literature that made use of multiplier bootstrap is Chernozhukov et al. (2023). That paper however uses multiplier bootstrap to simulate critical values of GMM-type statistics at a hypothetical parameter values for subsequent test inversion.

APPENDIX D: ADDITIONAL MONTE CARLO RESULTS

FIGURE 8.— Coverage frequency for $CB_{\alpha,n}$ as function of ω in the 2-dimensional design with $\kappa_n = \mu_n$.



Note: The dashed lines correspond to the asymptotic 95% confidence interval (point-wise) for the parameter $p = 0.95$ of Bernoulli random variable based on a random sample of 1000 observations. Some values of the estimated frequency can be slightly outside of the confidence interval as a result of multiple hypothesis testing. Values of ω close to zero result in negligible under-coverage. As sample size grows, the problematic area shrinks.

FIGURE 9.— MC average excess lengths of $CB_{\alpha,n}$ and $CB_{\alpha,n,\mathcal{P}}$ as function of ω in the 2-dimensional design with $\kappa_n = \mu_n$.





FIGURE 10.— Coverage frequency for $CB_{\alpha,n}$ as function of ω in the 2-dimensional design with $\kappa_n = 0$.



Note: The dashed lines correspond to the asymptotic 95% confidence interval (point-wise) for the parameter $p = 0.95$ of Bernoulli random variable based on a random sample of 1000 observations. Some values of the estimated frequency can be slightly outside of the confidence interval as a result of multiple hypothesis testing. Values of ω close to zero result in negligible conservative coverage.

FIGURE 11.— MC average excess lengths of $CB_{\alpha,n}$ and $CB_{\alpha,n,\mathcal{P}}$ as function of ω in the 2-dimensional design with $\kappa_n = 0$.

