

Sensitivity Analysis using Approximate Moment Condition Models*

Timothy B. Armstrong[†]

Yale University

Michal Kolesár[‡]

Princeton University

July 29, 2020

Abstract

We consider inference in models defined by approximate moment conditions. We show that near-optimal confidence intervals (CIs) can be formed by taking a generalized method of moments (GMM) estimator, and adding and subtracting the standard error times a critical value that takes into account the potential bias from misspecification of the moment conditions. In order to optimize performance under potential misspecification, the weighting matrix for this GMM estimator takes into account this potential bias, and therefore differs from the one that is optimal under correct specification. To formally show the near-optimality of these CIs, we develop asymptotic efficiency bounds for inference in the locally misspecified GMM setting. These bounds may be of independent interest, due to their implications for the possibility of using moment selection procedures when conducting inference in moment condition models. We apply our methods in an empirical application to automobile demand, and show that adjusting the weighting matrix can shrink the CIs by a factor of 3 or more.

*We thank Isaiah Andrews, Gary Chamberlain, Tim Christensen, Mikkel Plagborg-Møller, Jesse Shapiro, Martin Weidner and participants at several conferences and seminars for helpful comments and suggestions, and Soonwoo Kwon for research assistance. All remaining errors are our own. Armstrong acknowledges support by National Science Foundation Grant SES-1628939. Kolesár acknowledges support by the Sloan Research Fellowship, and by the National Science Foundation Grant SES-1628878.

[†]email: timothy.armstrong@yale.edu

[‡]email: mcolesar@princeton.edu

1 Introduction

Economic models are typically viewed as approximations of reality. However, conventional approaches to estimation and inference assume that a model holds exactly. In this paper, we weaken this assumption, and consider inference in a class of models characterized by moment conditions which are only required to hold in an approximate sense. The failure of the moment conditions to hold exactly may come from failure of exclusion restrictions (e.g. through omitted variable bias or because instruments enter the structural equation directly in an IV model), functional form misspecification, or other sources such as measurement error, or data contamination.

We assume that we have a model characterized by a set of population moment conditions $g(\theta)$. In the generalized method of moments (GMM) framework, for instance, $g(\theta) = E[g(w_i, \theta)]$, which can be estimated by the sample analog $\frac{1}{n} \sum_{i=1}^n g(w_i, \theta)$, based on the sample $\{w_i\}_{i=1}^n$. When evaluated at the true parameter value θ_0 , the population moment condition lies in a known set specified by the researcher,

$$g(\theta_0) = c/\sqrt{n}, \quad c \in \mathcal{C}.$$

The set \mathcal{C} formalizes the way in which the moment conditions may fail, and it can then be varied as a form of sensitivity analysis, with $\mathcal{C} = \{0\}$ reducing to the correctly specified case.

We focus on local misspecification: the scaling by \sqrt{n} implies that the specification error and the sampling error are of the same order of magnitude, and it ensures that the asymptotic approximation captures the fact that it may not be clear from the sample at hand whether the model is correctly specified. It also leads to increased tractability, allowing us to deliver a simple method for inference on a structural parameter of interest $h(\theta_0)$, rather than a pseudo-true parameter. This tractability has made local misspecification a popular tool for sensitivity analysis in applied work, especially following the recent influential paper by [Andrews et al. \(2017\)](#).¹ As with any asymptotic device, our modeling of misspecification as local should not be taken to mean that we literally believe that the model would be closer to correct if we had more data. Rather, its usefulness should be judged by whether it yields accurate approximations to the finite-sample behavior of estimators and confidence intervals, which in our case requires that the set \mathcal{C}/\sqrt{n} be small relative to sampling uncertainty.

We propose a simple method for constructing asymptotically valid confidence intervals (CIs) under this setup: one takes a standard estimator, such as the GMM estimator, and adds and subtracts its standard error times a critical value that takes into account the potential

¹For recent empirical examples using local sensitivity analysis, see [Gayle and Shephard \(2019\)](#), or [Duffo et al. \(2018\)](#).

asymptotic bias of the estimator, in addition to its variance. A key insight of this paper is that because the CIs must be widened to take into account the potential bias, the optimal weighting matrix for the correctly specified case (the inverse of the variance matrix of the moments) is generally no longer optimal under local misspecification. Rather, the optimal weighting matrix takes into account potential misspecification in the moments in addition to the variance of their estimates: it places less weight on moments that are allowed to be further from zero according to the researcher’s specification of the set \mathcal{C} . We also show that an analogous result holds for other performance criteria, such as estimation under the mean-squared error: the optimal weighting matrix again trades off the potential misspecification of the moments against their precision, although the optimal tradeoff is different.

To illustrate the practical importance of this result, we apply our methods to form misspecification-robust CIs in an empirical model of automobile demand based on [Berry et al. \(1995\)](#). We consider sets \mathcal{C} motivated by the forms of local misspecification considered in [Andrews et al. \(2017\)](#), who calculate the asymptotic bias of the usual GMM estimator in this model. We find that adjusting the weighting matrix to account for potential misspecification substantially reduces the potential bias of the estimator and, as a result, leads to large efficiency improvements of the optimal CI relative to a CI based on the GMM estimator that is optimal under correct specification: it shrinks the CI by up to a factor of 3 or more in our main specifications. As a result, we obtain informative CIs in this model even under moderate amounts of misspecification.

We show that the CIs we propose are near-optimal when the set \mathcal{C} is convex and centrosymmetric ($c \in \mathcal{C}$ implies $-c \in \mathcal{C}$). To this end, we argue that the relevant “limiting experiment” for the locally misspecified GMM model is isomorphic to an approximately linear model of [Sacks and Ylvisaker \(1978\)](#), which falls under a general framework studied by, among others, [Donoho \(1994\)](#), [Cai and Low \(2004\)](#) and [Armstrong and Kolesár \(2018\)](#). We derive asymptotic efficiency bounds for CIs in the locally misspecified GMM model that formally translate bounds from the approximately linear limiting experiment to the locally misspecified GMM setting. In particular, these bounds imply that the scope for improvement over our CIs by optimizing expected length at a particular value of θ_0 and $c = 0$ (while still maintaining coverage over the whole parameter space for θ and \mathcal{C}) is limited, even if one optimizes expected length at the true values of θ_0 and c .

These efficiency bounds address an important criticism of our CIs: they require a priori specification of the set \mathcal{C} that defines misspecification, including both the *magnitude* of misspecification and *which* moments are misspecified. In particular, one cannot substantively improve upon our CI by, say, trying to use data-driven methods that gauge misspecification magnitude or try to determine which moments are misspecified. These bounds have implica-

tions for procedures proposed by [Andrews and Guggenberger \(2009\)](#), [DiTraglia \(2016\)](#) and [McCloskey \(2019\)](#), who consider the case where some moments are known to be correctly specified and no a priori bound is placed on the magnitude of misspecification of the remaining moments. As we discuss in [Section 4.3.2](#), in this case our CI reduces to the usual CI based on the k_1 correctly specified moments, and our efficiency bounds show that CIs proposed in these papers cannot substantively improve upon it.

Because we cannot use the data to determine the magnitude M of the set \mathcal{C} , we recommend plotting our CIs as a function of the potential misspecification magnitude M , or reporting the smallest value of M for which a particular finding breaks down. Such sensitivity analysis is easy to conduct under our proposed implementation. In particular, we show that, when the set \mathcal{C} is characterized by ℓ_p constraints, the class of weightings that trace out the optimal bias-variance tradeoff as a function of how much relative weight we put on the bias can be easily computed by recasting the problem as a penalized regression problem. By exploiting this analogy, we develop a simple algorithm for computing this class under ℓ_∞ constraints that is similar to the LASSO/LAR algorithm ([Efron et al., 2004](#); [Rosset and Zhu, 2007](#)); under ℓ_2 constraints, the solution admits a closed form.² Furthermore, as we discuss in [Section 3](#), this class of weightings is entirely determined by the shape of \mathcal{C} ; its magnitude M only determines the optimal relative weight we should put on the bias. Thus, tracing out the optimal weighting as a function of M can be done at essentially no additional computational cost. Furthermore, to avoid having to reoptimize the objective function with respect to the new weighting matrix, one can also form the CIs by adding and subtracting our critical value from a one-step estimator (see [Newey and McFadden, 1994](#), [Section 3.4](#)) based on any initial estimate that is \sqrt{n} -consistent under correct specification. We illustrate this approach in our empirical application in [Section 6](#).

Our paper is related to several strands of literature. Our efficiency results are related to those in [Chamberlain \(1987\)](#) for point estimation in the correctly specified setting (see also [Hansen, 1985](#)) and, more broadly, semiparametric efficiency theory in correctly specified settings (see, e.g., [Chapter 25 in van der Vaart, 1998](#)). As we discuss in [Section 4.3](#), some of our efficiency results are novel even in the correctly specified case, and may be of independent interest. [Kitamura et al. \(2013\)](#) consider efficiency of point estimators satisfying certain regularity conditions when the misspecification is bounded by the Hellinger distance. As we discuss in more detail in [Section 4.3.4](#), our results imply that under this form of misspecification, the optimal weighting matrix remains the same as under correct specification; both the usual GMM estimator and the estimator proposed by [Kitamura et al. \(2013\)](#) can

²An R package implementing our CIs under ℓ_p constraints is available at <https://github.com/kolesarm/GMMSensitivity>.

thus be used to form near-optimal CIs, and both estimators have the same local asymptotic minimax properties.

Local misspecification has been used in a number of papers, which include, among others, [Newey \(1985\)](#), [Berkowitz et al. \(2012\)](#), [Conley et al. \(2012\)](#), [Guggenberger \(2012\)](#), and [Bugni and Ura \(2019\)](#), and has antecedents in the literature on robust statistics (see [Huber and Ronchetti, 2009](#), and references therein). [Andrews et al. \(2017\)](#) consider this setting and note that asymptotic bias of a regular estimator can be calculated using influence function weights, which they call the sensitivity, and show how such calculations can be used for sensitivity analysis in applications (see also extensions of these ideas in [Andrews et al. 2018](#) and [Mukhin 2018](#)). Our results imply that, if one is interested in inference, conclusions of such sensitivity analysis may be substantially sharpened by using the misspecification-optimal weighting matrix, or, equivalently, the misspecification-optimal sensitivity. In independent work, [Bonhomme and Weidner \(2020\)](#) provide a framework for estimation and inference in misspecified likelihood models when the misspecification set \mathcal{C} is defined with respect to a larger class of models using statistical notions of distance. Our focus is on overidentified moment condition models, as in [Andrews et al. \(2017\)](#), and we are agnostic about how \mathcal{C} is determined. The proposal to use estimators that optimize an asymptotic bias-variance tradeoff using the influence function is common to both papers. The efficiency bounds in [Section 4](#) are unique to the present paper.

The rest of this paper is organized as follows. [Section 2](#) presents our misspecification robust CIs. [Section 3](#) gives step-by-step instructions for computing our CIs, along with discussion of other practical issues. [Section 4](#) presents efficiency bounds for CIs in locally misspecified models; it can be skipped by readers interested only in implementing the methods. [Section 5](#) discusses applications to particular moment condition models. [Section 6](#) presents an empirical application. Additional results and proofs are collected in appendices and an online supplement.

2 Misspecification-robust CIs

We have a model that maps a vector of parameters $\theta \in \Theta \subseteq \mathbb{R}^{d_\theta}$ to a d_g -dimensional population moment condition $g(\theta)$ that restricts the distribution of the observed data $\{w_i\}_{i=1}^n$. We allow the moment condition model to be locally misspecified, so that at the true value θ_0 , the population moment condition is not necessarily zero, but instead lies in a \sqrt{n} -neighborhood of 0:

$$g(\theta_0) = c/\sqrt{n}, \quad c \in \mathcal{C}, \quad (1)$$

where $\mathcal{C} \subseteq \mathbb{R}^{d_g}$ is a known set. The set \mathcal{C} may allow for misspecification in potentially all moment conditions; we do not require that some elements of c are zero. Our goal is to construct a CI for a scalar $h(\theta_0)$, where $h: \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}$ is a known function. For example, if we are interested in one of the elements θ_j of θ , we would take $h(\theta) = \theta_j$. More generally, the function h will be nonlinear, as is, for example, generally the case when θ is a vector of supply or demand parameters, and $h(\theta)$ is an elasticity, or some counterfactual.

This setup allows (but does not require) both θ_0 and $h(\theta_0)$ to have the same interpretation as in the correctly specified case, so that our CIs may still be interpreted as CIs for the structural parameter, elasticity, or counterfactual of interest. For this interpretation, one typically needs to rule out forms of misspecification that affect the mapping $\theta \mapsto h(\theta)$. While we do not formally consider cases in which this mapping itself is misspecified, such cases are covered under a mild generalization of our framework, in which h is a function of both θ and c .

Note that the interpretation of $h(\theta_0)$, and the conceptual framework defining θ_0 is not affected by our modeling of the misspecification as local: given a set $\tilde{\mathcal{C}}_n = \mathcal{C}/\sqrt{n}$, the moment conditions $g(\theta_0) \in \tilde{\mathcal{C}}_n$ describe the restrictions that the data generating process and the researcher’s modeling assumptions place on θ_0 .³ The plausibility of these restrictions is evaluated for a given sample size at hand; it doesn’t depend on assumptions about how $\tilde{\mathcal{C}}_n$ changes with n . While we focus on sequences $\tilde{\mathcal{C}}_n = \mathcal{C}/\sqrt{n}$, we discuss in Remark 3.3 how our insights can be used to construct CIs that are valid global misspecification, when $\tilde{\mathcal{C}}_n$ is fixed with n .

To formalize the notion of asymptotic validity and efficiency of CIs, we will need to allow the true parameter value θ_0 as well as the vector c and the data generating process (and hence the map $\theta \mapsto g(\theta)$) to vary with the sample size. For clarity of exposition, we focus here on the case in which these parameters are fixed. See Theorem 4.1 and Appendix C for the general case. Under some forms of misspecification, such as functional form misspecification, there may be additional higher-order terms on the right-hand side of (1); our results remain unchanged if this is the case. Again, for clarity of exposition, we focus on the case in which (1) holds exactly.

We assume that the sample moment condition $\hat{g}(\theta)$, constructed using the data $\{w_i\}_{i=1}^n$, satisfies

$$\sqrt{n}(\hat{g}(\theta_0) - g(\theta_0)) \xrightarrow{d} \mathcal{N}(0, \Sigma), \quad (2)$$

³Formally, for a given sample size n , θ_0 may be set identified, and the identified set under a distribution P is defined as the set of parameters θ_0 that satisfy the moment conditions $E_P g(w_i, \theta_0) \in \tilde{\mathcal{C}}_n$ where E_P denotes expectation under P . We construct CIs that cover $h(\theta_0)$ for points θ_0 in the identified set (see Imbens and Manski, 2004, for a discussion of this notion of coverage). See Section 4 and Appendix C for formal definitions of coverage and optimality of our CIs.

where \xrightarrow{d} denotes convergence in distribution as $n \rightarrow \infty$. In the GMM model, the population and sample moment conditions are given by $g(\theta) = E[g(w_i, \theta)]$ and $\hat{g}(\theta) = \frac{1}{n} \sum_{i=1}^n g(w_i, \theta)$, respectively, where $g(\cdot, \cdot)$ is a known function. However, to cover other minimum distance problems, we do not require that the moment conditions necessarily take this form. We further assume that the moment condition is smooth enough so that

$$\text{for any } \theta_n = \theta_0 + \mathcal{O}_P(1/\sqrt{n}), \quad \hat{g}(\theta_n) - \hat{g}(\theta_0) = \Gamma(\theta_n - \theta_0) + o_P(1/\sqrt{n}), \quad (3)$$

where Γ is the $d_g \times d_\theta$ derivative matrix of g at θ_0 . Conditions (2) and (3) are standard regularity conditions in the literature on linear and nonlinear estimating equations; see [Newey and McFadden \(1994\)](#) for primitive conditions. Finally, we also assume that h is continuously differentiable with the $1 \times d_\theta$ derivative matrix at θ_0 given by H .

2.1 CIs based on asymptotically linear estimators

Under correct specification, when $\mathcal{C} = \{0\}$, standard estimators \hat{h} of $h(\theta)$ are asymptotically linear in $\hat{g}(\theta_0)$. This will typically extend to our locally misspecified case, so that for some vector $k \in \mathbb{R}^{d_g}$,

$$\sqrt{n}(\hat{h} - h(\theta_0)) = k' \sqrt{n} \hat{g}(\theta_0) + o_P(1) \xrightarrow{d} \mathcal{N}(k'c, k'\Sigma k), \quad (4)$$

where the convergence in distribution follows by (1) and (2). If in addition, the estimator is regular (so that equality in Eq. (4) holds uniformly for θ in a \sqrt{n} -neighborhood of θ_0), then k will satisfy (see e.g. Section 2 in [Newey, 1990](#))

$$H = -k'\Gamma. \quad (5)$$

For example, in a GMM model, if we take $\hat{h} = h(\hat{\theta}_W)$ where

$$\hat{\theta}_W = \underset{\theta}{\operatorname{argmin}} \hat{g}(\theta)' W \hat{g}(\theta), \quad (6)$$

is the GMM estimator with weighting matrix W , Eqs. (4) and (5) will hold with $k' = -H(\Gamma'W\Gamma)^{-1}\Gamma'W$ (see [Newey, 1985](#)). Because the vector k determines the local asymptotic bias of the estimator, we follow [Andrews et al. \(2017\)](#), and refer to k as the *sensitivity* of \hat{h} .

We now show how to construct misspecification-robust CIs based on an asymptotically linear estimator \hat{h} with a given sensitivity k . In Section 2.2, we show how to choose this sensitivity optimally, to achieve the shortest CI among those based on regular asymptotically linear estimators. In Section 4, we will show that, under this choice of k , the resulting CI

is (near) optimal not only within the class of CIs based on regular asymptotically linear estimators, but among *all* CIs that satisfy the asymptotic coverage requirement.

Let \hat{k} and $\hat{\Sigma}$ be consistent estimates of k and Σ . Then by Slutsky's theorem,

$$\frac{\sqrt{n}(\hat{h} - h(\theta_0))}{\sqrt{\hat{k}'\hat{\Sigma}\hat{k}}} \xrightarrow{d} \mathcal{N}\left(\frac{k'c}{\sqrt{k'\Sigma k}}, 1\right).$$

Under correct specification, the right-hand side corresponds to a standard normal distribution, and we can form a CI with asymptotic coverage $100 \cdot (1 - \alpha)\%$ as $\hat{h} \pm z_{1-\alpha/2} \sqrt{\hat{k}'\hat{\Sigma}\hat{k}/n}$, where $z_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of a $\mathcal{N}(0, 1)$ distribution; this is the usual Wald CI.

When we allow for misspecification, the Wald CI will no longer be valid. However, note that the asymptotic bias $k'c/\sqrt{k'\Sigma k}$ is bounded in absolute value by $\overline{\text{bias}}_c(k)/\sqrt{k'\Sigma k}$ where $\overline{\text{bias}}_c(k) \equiv \sup_{c \in \mathcal{C}} |k'c|$. Therefore, given c , the z -statistic in the preceding display is asymptotically $\mathcal{N}(t, 1)$ where $|t| \leq \overline{\text{bias}}_c(k)/\sqrt{k'\Sigma k}$. This leads to the CI

$$\hat{h} \pm \text{cv}_\alpha\left(\frac{\overline{\text{bias}}_c(\hat{k})}{\sqrt{\hat{k}'\hat{\Sigma}\hat{k}}}\right) \cdot \sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\sqrt{n}, \quad (7)$$

where $\text{cv}_\alpha(\bar{t})$ is the $1 - \alpha$ quantile of $|Z|$, with $Z \sim \mathcal{N}(\bar{t}, 1)$. In particular, $\text{cv}_\alpha(0) = z_{1-\alpha/2}$, so that in the correctly specified case, (7) reduces to the usual Wald CI. As we discuss in Section 4, in the limiting experiment, this CI becomes equivalent to the fixed-length CI proposed by Donoho (1994).

To form a one-sided CI based on an estimator \hat{h} with sensitivity k , one can simply subtract its maximum bias, in addition to the standard error:

$$[\hat{h} - \overline{\text{bias}}_c(\hat{k})/\sqrt{n} - z_{1-\alpha} \sqrt{\hat{k}'\hat{\Sigma}\hat{k}/n}, \infty). \quad (8)$$

One could also form a valid two-sided CI by adding and subtracting the worst-case bias $\overline{\text{bias}}_c(\hat{k})/\sqrt{n}$ from \hat{h} , in addition to adding and subtracting $z_{1-\alpha/2} \sqrt{\hat{k}'\hat{\Sigma}\hat{k}/n}$; however, since \hat{h} cannot simultaneously have a large positive and a large negative bias, such CI will be conservative, and longer than the CI in (7).

2.2 Optimal CIs

The asymptotic length of the CI in Eq. (7) is given by

$$2 \cdot \text{cv}_\alpha\left(\overline{\text{bias}}_c(k)/\sqrt{k'\Sigma k}\right) \cdot \sqrt{k'\Sigma k}/\sqrt{n}. \quad (9)$$

To attain the shortest possible CI, we therefore need to use an estimator with sensitivity that minimizes this expression. We restrict attention to asymptotically linear estimators that are regular, so that we need to minimize (9) subject to (5). The CI length in Eq. (9) depends on θ only through Σ . Furthermore, it depends on the sensitivity only through the maximum bias, $\overline{\text{bias}}_{\mathcal{C}}(k)$, and the variance $k'\Sigma k$. Therefore, rather than minimizing (9) directly over all sensitivities k , one can first minimize the variance subject to a bound \overline{B} on the worst-case bias,

$$\min_k k'\Sigma k \quad \text{s.t. (5) and} \quad \sup_{c \in \mathcal{C}} |k'c| \leq \overline{B}, \quad (10)$$

and then vary the bound \overline{B} to find the bias-variance trade-off that leads to the shortest CI. In our implementation in Section 3, we focus on the case where \mathcal{C} is characterized by ℓ_p constraints, in which case a closed-form expression for the worst-case bias $\sup_{c \in \mathcal{C}} |k'c|$ is available, and it is computationally trivial to solve (10) directly or in Lagrange multiplier form. In general, when the set \mathcal{C} is convex, one can reformulate (10) as a convex optimization problem, leading to a computationally tractable solution (see Section 4). One can also use (10) to determine the optimal sensitivity for constructing one-sided CIs, if we use quantiles of excess length as the criterion for choosing a CI. We provide details in Appendix C.

Once the optimal sensitivity has been determined, we can implement an estimator with this sensitivity as a one-step estimator. In particular, let $\hat{\theta}_{\text{initial}}$ be an initial \sqrt{n} -consistent estimator of θ_0 , let $\hat{k} = k + o_P(1)$ be a consistent estimator of the desired sensitivity k . Then the one-step estimator

$$\hat{h} = h(\hat{\theta}_{\text{initial}}) + \hat{k}'\hat{g}(\hat{\theta}_{\text{initial}})$$

will have the desired sensitivity. This follows from the Taylor expansion

$$\begin{aligned} \sqrt{n}(\hat{h} - h(\theta_0)) &= H\sqrt{n}(\hat{\theta}_{\text{initial}} - \theta_0) + \hat{k}'\sqrt{n}\hat{g}(\hat{\theta}_{\text{initial}}) + o_P(1) \\ &= (H + \hat{k}'\Gamma)\sqrt{n}(\hat{\theta}_{\text{initial}} - \theta_0) + \hat{k}'\sqrt{n}\hat{g}(\theta_0) + o_P(1), \end{aligned}$$

where the second line follows from (3). It then follows from (5) that the first term converges in probability to zero, and \hat{h} satisfies (4).

3 Practical implementation

We now give step-by-step instructions for computing our CI. To make it easy to determine the sensitivity of the CI to the magnitude of misspecification, we consider sets of the form $\mathcal{C} = \mathcal{C}(M) = \{Mc: c \in \mathcal{C}(1)\}$, where the scalar M measures the magnitude of misspecification. We discuss the exact specification of the set $\mathcal{C}(M)$ in Remark 3.1 below.

The fact that M simply scales the potential magnitude of misspecification leads to a simplification when tracing out the optimal CI as a function of M . In particular, let $\{k_\lambda\}_{\lambda \geq 0}$ be the bias-variance optimizing class of sensitivities that traces out the solutions to Eq. (10) as we vary the bound \bar{B} when $\mathcal{C} = \mathcal{C}(1)$. The index λ determines the relative weight on the bias; it can correspond to the Lagrange multiplier in a Lagrangian formulation of (10), or we can simply take $\lambda = \bar{B}$ if we are solving (10) directly. Let $\bar{B}_\lambda = \overline{\text{bias}}_{\mathcal{C}(1)}(k_\lambda)$. It then follows by a change-of-variables argument that $\overline{\text{bias}}_{\mathcal{C}(M)}(k_\lambda) = M\bar{B}_\lambda$, and that k_λ minimizes the asymptotic variance subject to this bound on worst-case bias over $\mathcal{C}(M)$. Thus, $\{k_\lambda\}_{\lambda \geq 0}$ is also a bias-variance optimizing class of sensitivities for $\mathcal{C}(M)$. We therefore only need to compute the class $\{k_\lambda\}_{\lambda \geq 0}$ only once, even when a range of values M is considered.

With this simplification, we can construct CIs for a range of values of M as follows:

1. Obtain an initial estimate $\hat{\theta}_{\text{initial}}$ and estimates \hat{H} , $\hat{\Gamma}$ and $\hat{\Sigma}$ of H , Γ and Σ .

In particular, for the GMM model, when $\hat{g}(\theta) = \frac{1}{n} \sum_{i=1}^n g(w_i, \theta)$, we can take $\hat{\theta}_{\text{initial}}$ to be the GMM estimator $\hat{\theta}_W = \text{argmin}_\theta \hat{g}(\theta)' W \hat{g}(\theta)$ for some weight matrix W . The remaining objects are the usual quantities used to estimate the asymptotic variance of this estimator: $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n g(w_i, \hat{\theta}_{\text{initial}}) g(w_i, \hat{\theta}_{\text{initial}})'$ (or, in the case of dependent observations, an autocorrelation robust version of this estimate), $\hat{\Gamma} = \frac{d}{d\theta'} \hat{g}(\theta) \Big|_{\theta=\hat{\theta}_{\text{initial}}}$ (or, if $g(\theta, w)$ is nonsmooth, a numerical derivative as in [Hong et al. 2015](#), or Section 7.3 of [Newey and McFadden 1994](#)) and $\hat{H} = \frac{d}{d\theta'} h(\theta) \Big|_{\theta=\hat{\theta}_{\text{initial}}}$.

2. Compute the bias-variance optimizing class $\{\hat{k}_\lambda\}_{\lambda \geq 0}$ that solves (10) with $\mathcal{C} = \mathcal{C}(1)$ and with $\hat{\Sigma}$ in place of Σ . Algorithms and closed-form solutions for computing $\{\hat{k}_\lambda\}_{\lambda \geq 0}$ for particular choices of $\mathcal{C}(M)$ are discussed in [Remark 3.1](#). Let $\bar{B}_\lambda = \sup_{c \in \mathcal{C}(1)} |\hat{k}'_\lambda c|$. For each M , let λ_M^* minimize the CI length⁴ $2 \text{cv}_\alpha(M\bar{B}_\lambda / \sqrt{\hat{k}'_\lambda \hat{\Sigma} \hat{k}_\lambda}) \cdot \sqrt{\hat{k}'_\lambda \hat{\Sigma} \hat{k}_\lambda}$ over λ .
3. For each M , construct the one-step estimator $\hat{h}_{\lambda_M^*} = h(\hat{\theta}_{\text{initial}}) + \hat{k}'_{\lambda_M^*} \hat{g}(\hat{\theta}_{\text{initial}})$, and report the misspecification-robust CI under $\mathcal{C}(M)$

$$\hat{h}_{\lambda_M^*} \pm \text{cv}_\alpha \left(M\bar{B}_{\lambda_M^*} / \sqrt{\hat{k}'_{\lambda_M^*} \hat{\Sigma} \hat{k}_{\lambda_M^*}} \right) \cdot \sqrt{\hat{k}'_{\lambda_M^*} \hat{\Sigma} \hat{k}_{\lambda_M^*} / n}. \quad (11)$$

Remark 3.1 (Choice of $\mathcal{C}(M)$). A simple and flexible way of forming the set \mathcal{C} is to take

$$\mathcal{C} = \mathcal{C}(M) = \{B\gamma: \|\gamma\| \leq M\}, \quad (12)$$

⁴The critical value $\text{cv}_\alpha(b)$ can easily be computed in statistical software as the square root of the $1 - \alpha$ quantile of a non-central χ^2 distribution with 1 degree of freedom and non-centrality parameter b^2 .

where B is a $d_g \times d_\gamma$ matrix and $\|\cdot\|$ is some norm. The matrix B can be used to standardize moments, account for their correlations, or to pick out which moments are believed to be misspecified. For instance, setting B to the last d_γ columns of the $d_g \times d_g$ identity matrix allows for misspecification in the last d_γ moments, while maintaining that the first $d_g - d_\gamma$ moments are valid.

In light of our result in Section 4 that it is not possible to determine the set \mathcal{C} in a data-driven way, the normalizing matrix B and the baseline misspecification magnitude M used should be chosen to reflect application-specific arguments about which forms of misspecification are plausible; we can then vary M over other plausible choices as a form of sensitivity analysis. We illustrate this in the context of our application in Section 6, and we refer the reader to Conley et al. (2012) and Andrews et al. (2017) for additional examples and discussion. Alternatively, one can also use measures of statistical distance such as the probability of detecting that the model is misspecified to aid with interpretation of M , as suggested in Hansen and Sargent (2008) or Bonhomme and Weidner (2020).

While it is not possible to determine M automatically, it is possible to obtain a lower CI $[M_{\min}, \infty]$ for M , which can be used as a diagnostic check verifying that the values of M considered are not too small. We develop such tests by generalizing the J -test of overidentifying restrictions in Appendix B. We recommend reporting the lower bound M_{\min} along with the plot of the optimal CI as a function of M .

The norm $\|\cdot\|$ determines how the researcher's bounds on each element of γ interact. With the ℓ_∞ norm, one places separate bounds on each element of γ , which leads to a simple interpretation: no single element of γ can be greater than M . Under an ℓ_p norm with $1 \leq p < \infty$, the bounds on each element of γ interact with each other, so that larger amounts of misspecification in one element is allowed if other elements are correctly specified. Depending on whether such interactions are desirable, we recommend setting $p = 2$, or $p = \infty$.

For these choices of the norm, computing the class of optimal sensitivities $\{\hat{k}_\lambda\}_{\lambda \geq 0}$ is particularly simple. In particular, when $\|\cdot\|$ corresponds to an ℓ_p norm, the worst-case bias has a closed form, since by Hölder's inequality, $\overline{\text{bias}}_{\mathcal{C}(M)}(k) = \sup_{\|\gamma\|_p \leq 1} M|k' B \gamma| = M \|B' k\|_{p'}$, where p' is the Hölder complement of p ($p' = 1$ if $p = \infty$, while $p' = 2$ if $p = 2$), and the optimal sensitivities $\{\hat{k}_\lambda\}_{\lambda \geq 0}$ can be computed by casting the problem as a penalized regression problem. We explain the connection to penalized regression, and provide details in Appendix A.2.

When $p = 2$, so that $\|\cdot\|$ corresponds to the Euclidean norm, the problem is analogous to ridge regression, and the optimal sensitivities in Step 2 of the implementation take the form $\hat{k}'_\lambda = -H(\Gamma' W_\lambda \Gamma)^{-1} \Gamma' W_\lambda$, where $W_\lambda = (\lambda B B' + \hat{\Sigma})^{-1}$, with $\bar{B}_\lambda = \|B' \hat{k}_\lambda\|_2$. As an alternative

to using the one-step estimator in Step 3 of the implementation, one can implement this sensitivity directly as a GMM estimator with weighting matrix W_λ (see also Remark 3.2 below). Relative to the optimal weighting matrix Σ^{-1} under correct specification, the matrix W_λ trades off precision of the moments against their potential misspecification.

When $p = \infty$, the penalized regression analogy leads a simple algorithm for computing the optimal sensitivities $\{\hat{k}_\lambda\}_{\lambda \geq 0}$ that is similar to the LASSO/LAR algorithm (Efron et al., 2004). We give details on the algorithm in Appendix A.2. It follows from this algorithm if B corresponds to columns of the identity matrix, as M grows, the optimal sensitivity successively drops the “least informative” moments, so that in the limit, if $d_g \leq d_\gamma + d_\theta$, the optimal sensitivity corresponds to that of an exactly identified GMM estimator based on the d_θ “most informative” moments only, where “informativeness” is given by both the variability of a given moment, and its potential misspecification. If $d_g > d_\gamma + d_\theta$, one simply drops all invalid moments in the limit.

Remark 3.2. In Step 3 of our implementation, we use a one-step estimator \hat{h}_λ to compute a CI that is asymptotically valid and optimal. Due to concerns about finite-sample behavior (analogous to concerns about finite sample behavior of one-step estimators in the correctly specified case), one may prefer using a different estimator that is asymptotically equivalent to \hat{h}_λ . In general, one can implement an estimator with sensitivity k as a GMM or minimum distance estimator by using an appropriate weighting matrix, so that one can in particular replace \hat{h}_λ by $h(\hat{\theta}_W)$, with the weighting matrix W appropriately chosen. To give the formula for the weighting matrix, let Γ_\perp denote a $d_g \times (d_g - d_\theta)$ matrix that’s orthogonal to Γ , so that $\Gamma'_\perp \Gamma = 0$, and let $\hat{\Gamma}_\perp$ denote a consistent estimate. Let S denote a $d_g \times d_\theta$ matrix that satisfies $S'\hat{\Gamma} = -I$ and $\hat{k}_\lambda = S\hat{H}'$. Then we can set $W = SW_1S' + \hat{\Gamma}_\perp W_2 \hat{\Gamma}'_\perp$ for some non-singular matrix W_1 , and an arbitrary conformable matrix W_2 . It can be verified by simple algebra that $\hat{\theta}_W$ will have sensitivity k_λ .

Remark 3.3 (Global misspecification). While we focus on the local misspecification setting, in which the set \mathcal{C}/\sqrt{n} shrinks with n at a $1/\sqrt{n}$, one can use our insights about optimal weighting to construct a CI that retains the near-optimality properties of the above CI under local misspecification, while having correct coverage under asymptotics in which this set shrinks more slowly or stays fixed with the sample size (the latter is termed “global misspecification” in the literature). Let W be a weighting matrix that leads to the optimal sensitivity, as described in Remark 3.2 above, and let $\mathcal{I}_{\tilde{c}}$ be a CI constructed from the GMM estimator with moment conditions $\theta \mapsto g(w_i, \theta) - \tilde{c}$. Let $\mathcal{I} = \cup_{\tilde{c} \in \mathcal{C}/\sqrt{n}} \mathcal{I}_{\tilde{c}}$ be the union of these CIs over possible values of \tilde{c} in the set \mathcal{C}/\sqrt{n} . Such an approach was suggested in the context of misspecified linear IV by Conley et al. (2012), although they did not consider adjusting

the weighting matrix. The resulting CI has correct asymptotic coverage under both local and global misspecification, and, for one-sided CI construction, is asymptotically equivalent under local misspecification to the CI discussed above. We provide further details in Appendix D. In the appendix, we also discuss a second approach to constructing CIs valid under global misspecification based on misspecification-robust standard errors (Hall and Inoue, 2003), which is applicable if the estimate of the worst-case bias under global misspecification is asymptotically normal. The resulting one- and two-sided CIs are asymptotically equivalent under local misspecification to the optimal CIs discussed above.

Remark 3.4 (Other performance criteria). In addition to constructing a CI, one may be interested in a point estimate of $h(\theta_0)$, using mean squared error (MSE) as the criterion. The steps to forming the MSE optimal point estimate are exactly the same as above, except that, rather than minimizing CI length in Step 2, we choose λ to minimize $\overline{\text{bias}}_{\mathcal{C}}(\hat{k}_\lambda)^2 + \hat{k}'_\lambda \hat{\Sigma} \hat{k}_\lambda = M\overline{B}_\lambda + \hat{k}'_\lambda \hat{\Sigma} \hat{k}_\lambda$. Similar ideas apply to other criteria, such as mean absolute deviation or quantiles of excess length of one-sided CIs (discussed in Appendix C). If λ is chosen differently in Step 2 the CI computed in Step 3 will be longer than the one computed at λ_M^* , but it will still have correct coverage.

4 Efficiency bounds and near optimality

The CI given in Eq. (11) has the apparent defect that the local misspecification vector c is reflected in the length of the CI only through the a priori restriction \mathcal{C} imposed by the researcher. Thus, if the researcher is conservative about misspecification, the CI will be wide, even if it turns out that c is in fact much smaller than the a priori bounds defined by \mathcal{C} . Moreover, this approach requires the researcher to explicitly specify the set \mathcal{C} , including any tuning parameters such as the parameter M if the set takes the form $\mathcal{C} = \mathcal{C}(M)$ that we considered in Section 3. One may therefore seek to improve upon this CI by estimating the magnitude of c , or by estimating the tuning parameters, and constructing a CI that is shorter if these estimates indicate that misspecification is mild. Similarly, it may be restrictive to require that the CI be centered at an asymptotically linear estimator: this rules out, for example, using a J -test to decide which moments to use.

The main result of this section shows that, when \mathcal{C} is convex and centrosymmetric ($c \in \mathcal{C}$ implies $-c \in \mathcal{C}$), the scope for improving on the CI in (11) is nonetheless limited: no sequence of CIs that maintain coverage under all local misspecification vectors $c \in \mathcal{C}$ can be substantially tighter, even under correct specification. This result can be interpreted as translating results from a “limiting experiment” that is an extension of the linear regression model. We first give a heuristic derivation of this limiting experiment and explain our result

in the context of this limiting experiment. We then present the formal asymptotic result, and discuss its implications in some familiar settings. Readers who are interested only in implementing the methods, rather than efficiency results, can skip this section.

We restrict attention in this section to the GMM model, in which $\hat{g}(\theta) = \frac{1}{n} \sum_{i=1}^n g(w_i, \theta)$, and we further restrict the data $\{w_i\}_{i=1}^n$ to be independent and identically distributed (i.i.d.). Similar to semiparametric efficiency theory in the standard, correctly specified case, this facilitates parts of the formal statements and proofs, such as the definition of the set of distributions under which coverage is required and the construction of least favorable submodels. We expect that analogous results could be obtained in other settings.

4.1 Limiting experiment

As discussed in Section 2.1, we can form CIs based on linear estimators with asymptotic distribution $\mathcal{N}(k'c, k'\Sigma k)$. This suggests that the problem of constructing an asymptotically valid CI for $h(\theta)$ in the model (1) is asymptotically equivalent to the problem of constructing a CI for the parameter $H\theta$ in the approximately linear model

$$Y = -\Gamma\theta + c + \Sigma^{1/2}\varepsilon, \quad c \in \mathcal{C}, \quad \varepsilon \sim \mathcal{N}(0, I), \quad (13)$$

where Γ , H and $\Sigma^{1/2}$ are known, and we observe Y . One can think of this model as an “approximately” linear regression model, with $-\Gamma$ playing the role of the design matrix of the (fixed) regressors, and c giving the approximation error. This model dates back at least to [Sacks and Ylvisaker \(1978\)](#), who considered estimation in this model when \mathcal{C} is a rectangular set and Σ is diagonal. The analog of the asymptotically linear estimator \hat{h} in (4) is the linear estimator $k'Y$. To see the analogy, note that $k'Y - H\theta$ is distributed $\mathcal{N}((-k'\Gamma - H)\theta + k'c, k'\Sigma k)$, and restricting ourselves to estimators that do not have infinite worst-case bias when θ is unrestricted gives the condition (5). In the limiting experiment, the analog of the CI (7) is given by the CI $k'Y \pm cv_\alpha(\overline{\text{bias}}_c(k)/\sqrt{k'\Sigma k}) \cdot \sqrt{k'\Sigma k}$. Finding weights that minimize the length of this CI is isomorphic to the problem of finding the sensitivity that minimizes the asymptotic CI length in (9).

For a general convex set \mathcal{C} , the bias-variance optimization problem in Eq. (10) can be reformulated as a convex programming problem, as shown in [Low \(1995\)](#). In particular, when the set \mathcal{C} is centrosymmetric (see Appendix C for the general case), the bias-variance optimizing class of weights $\{k_\lambda\}_{\lambda>0}$ is given by the class $\{k_\delta\}_{\delta>0}$, where

$$k'_\delta = k'_{\delta, \Sigma, \Gamma, H, \mathcal{C}} = \frac{-(c_\delta - \Gamma\theta_\delta)' \Sigma^{-1}}{(c_\delta - \Gamma\theta_\delta)' \Sigma^{-1} \Gamma H' / H H'}, \quad (14)$$

and, for each δ , c_δ, θ_δ are the solutions to the convex program

$$\sup_{\theta, c} H\theta \quad \text{s.t.} \quad c \in \mathcal{C}, \quad (c - \Gamma\theta)' \Sigma^{-1} (c - \Gamma\theta) \leq \delta^2/4. \quad (15)$$

It then follows from [Donoho \(1994\)](#) that among fixed-length CIs based on linear estimators (CIs that take the form $k'Y \pm \chi$ for some constant χ), the shortest CI in the limiting experiment takes the form

$$k'_{\delta^*} Y \pm \text{cv}_\alpha(\overline{\text{bias}}_{\mathcal{C}}(k_{\delta^*}) / \sqrt{k'_{\delta^*} \hat{\Sigma} k_{\delta^*}}) \cdot \sqrt{k'_{\delta^*} \Sigma k_{\delta^*}}, \quad (16)$$

where $\overline{\text{bias}}_{\mathcal{C}}(k_\delta) = -k'_\delta c_\delta$, and $\delta^* = \text{argmin}_{\delta > 0} 2 \text{cv}_\alpha(\overline{\text{bias}}_{\mathcal{C}}(k_\delta) / \sqrt{k'_\delta \Sigma k_\delta}) \cdot \sqrt{k'_\delta \Sigma k_\delta}$ is chosen to minimize the CI length. The CI in (11) is an analog of this CI, with δ playing the role of the index λ .

The CI in (16) takes a familiar form in the special case in which \mathcal{C} is a linear subspace of \mathbb{R}^{d_g} , so that for some $d_g \times d_\gamma$ full-rank matrix B with $d_\gamma \leq d_g - d_\theta$, $\mathcal{C} = \{B\gamma : \gamma \in \mathbb{R}^{d_\gamma}\}$. Let B_\perp denote a $d_g \times (d_g - d_\gamma)$ matrix that's orthogonal to B . Then for any $\delta > 0$, $k'_\delta = k'_{LS,B}$, where

$$k'_{LS,B} = -H(\Gamma' B_\perp (B'_\perp \Sigma B_\perp)^{-1} B'_\perp \Gamma)^{-1} \Gamma' B_\perp (B'_\perp \Sigma B_\perp)^{-1} B'_\perp \quad (17)$$

is the sensitivity of the GLS estimator after pre-multiplying (13) by B'_\perp , (which effectively picks out the observations with zero misspecification). Since this estimator is unbiased, the CI in (16) becomes $k'_{LS,B} Y \pm z_{1-\alpha/2} \sqrt{k'_{LS,B} \Sigma k_{LS,B}}$.

Like the asymptotic CI (11), the CI in (16) has the potential drawback that its length is determined by the worst possible misspecification in \mathcal{C} , leaving open the possibility of efficiency improvements when c turns out to be close to zero. As a best-case scenario for such improvements, consider the problem: among confidence sets with coverage at least $1 - \alpha$ for all $\theta \in \mathbb{R}^{d_\theta}$ and $c \in \mathcal{C}$, minimize expected length when $\theta = \theta^*$ and $c = 0$. Note that this setup is even more favorable for potential improvements on our CI, since it allows the researcher to guess correctly that θ is equal to some θ^* , and it allows for confidence sets that are not intervals (in this case, length is defined as Lebesgue measure). Let $\kappa_*(H, \Gamma, \Sigma, \mathcal{C})$ denote the ratio of this optimized expected length relative to the length of the CI in (16) (it can be shown that this ratio does not depend on θ^*).

If \mathcal{C} is convex, a formula for $\kappa_*(H, \Gamma, \Sigma, \mathcal{C})$ follows from applying the general results in Corollary 3.3 in [Armstrong and Kolesár \(2018\)](#) to the limiting model. If \mathcal{C} is also centrosymmetric, then

$$\kappa_*(H, \Gamma, \Sigma, \mathcal{C}) = \frac{(1 - \alpha) E[\omega(2(z_{1-\alpha} - Z)) | Z \leq z_{1-\alpha}]}{2 \min_\delta \text{cv}_\alpha \left(\frac{\omega(\delta)}{2\omega'(\delta)} - \frac{\delta}{2} \right) \omega'(\delta)}, \quad (18)$$

where $Z \sim \mathcal{N}(0, 1)$ and $\omega(\delta)$ is two times the optimized value of (15). Furthermore, we show in Theorem C.7 that the right-hand side is lower-bounded by $(z_{1-\alpha}(1-\alpha) - \tilde{z}_\alpha \Phi(\tilde{z}_\alpha) + \phi(z_{1-\alpha}) - \phi(\tilde{z}_\alpha))/z_{1-\alpha/2}$, where $\tilde{z}_\alpha = z_{1-\alpha} - z_{1-\alpha/2}$ for any H, Γ, Σ and \mathcal{C} , where $\phi(\cdot)$ denotes the standard normal density. For $\alpha = 0.05$, this universal lower bound evaluates to 71.7%. Evaluating κ_* for particular choices of H, Γ, Σ , and \mathcal{C} often yields even higher efficiency.

If \mathcal{C} is a linear subspace, then $\omega(\delta)$ is linear, and

$$\kappa_*(H, \Gamma, \Sigma, \mathcal{C}) = \frac{(1-\alpha)z_{1-\alpha} + \phi(z_{1-\alpha})}{z_{1-\alpha/2}} \geq \frac{z_{1-\alpha}}{z_{1-\alpha/2}}, \quad (19)$$

where the lower bound follows since $\phi(z_{1-\alpha}) \geq \alpha z_{1-\alpha}$ by the Gaussian tail bound $1 - \Phi(x) \leq \phi(x)/x$ for $x > 0$. This bound corresponds to that in Pratt (1961) for the case of a univariate normal mean. The potential efficiency improvement essentially comes from using prior knowledge of θ^* to turn a two-sided critical value into a one-sided critical value. Furthermore, it follows from Joshi (1969) that the CI $k'_{LS,B}Y \pm z_{1-\alpha/2}\sqrt{k'_{LS,B}\Sigma k_{LS,B}}$ is the unique CI that achieves minimax expected length. Thus, not only is the scope for improvement at a particular θ^* bounded by (19), any CI with shorter expected length at some θ^* must necessarily perform worse elsewhere in the parameter space.

For the one-sided CI (8), the analogous CI in the limiting experiment is $[k'Y - \overline{\text{bias}}_c(k) - z_{1-\alpha}\sqrt{k'\Sigma k}, \infty)$, and, as we discuss in Appendix C, to choose the optimal sensitivity k , one can consider optimizing a given quantile of its worst-case excess length. The results in Armstrong and Kolesár (2018) again give an efficiency bound for improvement at $c = 0$ and a particular θ^* , analogous to (18) for the two-sided case. See Appendix C for details. If \mathcal{C} is a linear subspace, then optimizing quantiles of worst-case excess length yields the CI $[k'_{LS,B}Y - z_{1-\alpha}\sqrt{k'_{LS,B}\Sigma k_{LS,B}}, \infty)$, independently of the quantile one is optimizing. Furthermore, the efficiency bound implies that this one-sided CI is in fact fully optimal over all quantiles of excess length and all values of θ, c in the local parameter space.

These efficiency results for the CI (16) in the limiting experiment suggest that the scope for improvement over the CI in (11) should be limited in large samples. Theorem 4.1, stated in the next section, uses the analogy with the approximately linear model (13) along with Le Cam-style arguments involving least favorable submodels to show that this bound indeed translates to the locally misspecified GMM model. For one-sided CIs, we state an analogous result in Appendix C. We discuss the implications of these results in Section 4.3.

4.2 Asymptotic efficiency bound

To make precise our statements about coverage and efficiency, we need the notion of uniform (in the underlying distribution) coverage of a confidence interval. This requires additional notation, which we now introduce. Let \mathcal{P} denote a set of distributions P of the data $\{w_i\}_{i=1}^n$, and let $\Theta_n \subseteq \mathbb{R}^{d_\theta}$ denote the parameter space for θ . We require coverage for all pairs $(\theta, P) \in \Theta_n \times \mathcal{P}$ such that $\sqrt{n}g_P(\theta) \in \mathcal{C}$, where the subscript P on the population moment condition makes it explicit that it depends on the distribution of the data.⁵ Letting $\mathcal{S}_n = \{(\theta, P) \in \Theta_n \times \mathcal{P} : \sqrt{n}g_P(\theta) \in \mathcal{C}\}$ denote this set, the condition for coverage at confidence level $1 - \alpha$ can be written

$$\liminf_{n \rightarrow \infty} \inf_{(\theta, P) \in \mathcal{S}_n} P(h(\theta) \in \mathcal{I}_n) \geq 1 - \alpha. \quad (20)$$

We say that a confidence set \mathcal{I}_n is asymptotically valid (uniformly over \mathcal{S}_n) at confidence level $1 - \alpha$ if this condition holds.⁶

Among two-sided CIs of the form $\hat{h} \pm \hat{\chi}$ that are asymptotically valid, we prefer CIs with shorter expected length. To avoid issues with convergence of moments, we use truncated expected length, and define the asymptotic expected length of a two-sided CI at $P_n \in \mathcal{P}$ as $\liminf_{T \rightarrow \infty} \liminf_{n \rightarrow \infty} E_{P_n} \min\{\sqrt{n} \cdot 2\hat{\chi}, T\}$, where E_P denotes expectation under P .

We are now ready to state the main efficiency result.

Theorem 4.1. *Suppose that \mathcal{C} is convex and centrosymmetric. Let \hat{h}_{λ^*} and $\hat{\chi}_{\lambda^*}$ be formed as in Section 3. Suppose that Assumptions C.2, C.3, C.5, C.6 and C.7 in Appendix C hold. Suppose that the data $\{w_i\}_{i=1}^n$ are i.i.d. under all $P \in \mathcal{P}$. Let (θ^*, P_0) be correctly specified (i.e. $g_{P_0}(\theta^*) = 0$) such that \mathcal{P} contains a submodel through P_0 satisfying Assumption C.1. Then:*

- (i) *The CI $\hat{h}_{\lambda^*} \pm \hat{\chi}_{\lambda^*}$ is asymptotically valid, and its half-length $\hat{\chi}_{\lambda^*}$ satisfies $\sqrt{n}\hat{\chi}_{\lambda^*} = \chi(\theta, P) + o_P(1)$ uniformly over $(\theta, P) \in \mathcal{S}_n$ where*

$$\chi(\theta, P) = \min_k \text{cv}_\alpha(\overline{\text{bias}}_{\mathcal{C}}(k) / \sqrt{k' \Sigma_{\theta, P} k}) \sqrt{k' \Sigma_{\theta, P} k}$$

with $\overline{\text{bias}}_{\mathcal{C}}(k)$ calculated with $\Gamma = \Gamma_{\theta, P}$ and $H = H_\theta$.

⁵To be precise, we should also subscript all other quantities such as Γ and Σ by P . To prevent notational clutter, we drop this index in the main text unless it causes confusion.

⁶In general, θ_0 and $h(\theta_0)$ may be set identified for a given sample size n (although our assumptions imply that the identified set will shrink at a root- n rate). The coverage requirement (20) states that the CI must cover points in the identified set for $h(\theta)$, as in Imbens and Manski (2004); see Appendix C.

(ii) For any other asymptotically valid CI $\hat{h} \pm \hat{\chi}$,

$$\frac{\liminf_{T \rightarrow \infty} \liminf_{n \rightarrow \infty} E_{P_0} \min\{\sqrt{n} \cdot 2\hat{\chi}, T\}}{2\chi(\theta^*, P_0)} \geq \kappa_*(H_{\theta^*}, \Gamma_{\theta^*, P_0}, \Sigma_{\theta^*, P_0}, \mathcal{C}),$$

where $\kappa_*(H, \Gamma, \Sigma, \mathcal{C})$ is defined in (18). Furthermore, for any H, Σ, Γ , and \mathcal{C} , κ_* admits the universal lower bound $(z_{1-\alpha}(1-\alpha) - \tilde{z}_\alpha \Phi(\tilde{z}_\alpha) + \phi(z_{1-\alpha}) - \phi(\tilde{z}_\alpha))/z_{1-\alpha/2}$, where $\tilde{z}_\alpha = z_{1-\alpha} - z_{1-\alpha/2}$ and $\phi(\cdot)$ denotes the standard normal density.

The proof for this theorem is given in Appendix C, which also gives an analogous result for one-sided confidence intervals. Assumptions C.2, C.3, C.5, C.6 and C.7, stated in Appendix C, require that the conditions in Section 2 hold in a uniform sense over the class \mathcal{P} . In the supplemental materials, we give primitive conditions for these assumptions in the misspecified linear IV model. Assumption C.1, also stated in the appendix, requires that the class \mathcal{P} be rich enough to contain a submodel that is least favorable for the GMM problem, so that the class doesn't implicitly impose any other conditions that could be used to make inference easier. In the supplemental materials, we provide a general way of constructing a submodel satisfying these conditions.

The universal lower bound on κ_* is new and may be of independent interest. For $\alpha = 0.05$, it evaluates to 71.7%. The universal lower bound is sharp in the sense that there exist Γ, Σ, H and \mathcal{C} for which κ_* equals this lower bound. In particular applications, the efficiency bound κ_* can be computed at estimates of Γ, Σ and H , and often, this gives much higher efficiencies. We illustrate these bounds in the empirical application in Section 6.

4.3 Discussion

To help build intuition for the efficiency bound in Theorem 4.1, and to relate this result to the literature, we now consider some special cases. We first discuss the (standard) correctly specified case. Second, we consider the case in which some moments are known to be valid, and the misspecification in the remaining moments is unrestricted. This case may be of interest in its own right. We then discuss the general case. Finally, we discuss the connection to certain statistical measures of distance considered in the literature.

4.3.1 Correctly specified case

Suppose that $\mathcal{C} = \{0\}$. This is in particular a linear subspace of \mathbb{R}^{d_g} , with $B = 0$, and $B_\perp = I$, the $d_g \times d_g$ identity matrix. Thus, in the limiting experiment, the optimal CI uses the GLS estimator $k'_{LS,0} Y$, with $k_{LS,0}$ given in (17) (with $B = 0$). For testing the null hypothesis $H\theta = h_0$ against the one-sided alternative $H\theta \geq h_0$, the one-sided z -statistic based on

$k'_{LS,0}Y$ is uniformly most powerful (van der Vaart, 1998, Proposition 15.2). Inverting these tests yields the CI $[k'_{LS,0}Y - z_{1-\alpha}\sqrt{k'_{LS,0}\Sigma k_{LS,0}}, \infty)$. Since the underlying tests are uniformly most powerful, this CI achieves the shortest excess length, simultaneously for all quantiles and all possible values of the parameter θ . For two-sided CIs, the results described in Section 4.1 imply that the CI $h'_{LS,0}Y \pm z_{1-\alpha/2}\sqrt{k'_{LS,0}\Sigma k_{LS,0}}$ is the unique CI that achieves minimax expected length, and the efficiency of this CI relative to a CI that optimizes its expected length at a single value θ^* of θ when indeed $\theta = \theta^*$ is given in Eq. (19). It evaluates to 84.99% at $\alpha = 0.05$.

Applying Theorem 4.1 to the case $\mathcal{C} = \{0\}$ gives an asymptotic version of the two-sided efficiency bound. Furthermore, the CI in Theorem 4.1 reduces to the usual two-sided CI based on $\hat{\theta}_{\Sigma^{-1}}$. Thus, in this case, Theorem 4.1 shows that very little can be gained over the usual two-sided CI by optimizing the CI relative to a particular distribution P_0 . Results in the appendix give an analogous result for one-sided CIs. In the one-sided case, this asymptotic result is essentially a version of a classic result from the semiparametric efficiency literature for one-sided tests, applied to CIs (see Chapter 25.6 in van der Vaart, 1998). In the two-sided case, the result is, to our knowledge, new.

4.3.2 Some valid and some invalid moments

Consider now the case in which the first $d_g - d_\gamma$ moments are known to be valid, with the potential misspecification for the remaining d_γ moments unrestricted. Then $\mathcal{C} = \{(0', \gamma)'\} : \gamma \in \mathbb{R}^{d_\gamma}\}$ corresponds to a linear subspace with B given by the last d_γ columns of the identity matrix, and B_\perp given by the first $d_g - d_\gamma$ columns. Optimal CIs in the limiting experiment therefore use the estimator $k'_{LS,B}Y$, which is the GLS estimator based only on the observations with no misspecification.

The one-sided CI based on $k'_{LS,B}Y$ achieves the shortest excess length, simultaneously for all quantiles and all possible values of the parameter θ . The two-sided CI $k'_{LS,B}Y \pm z_{1-\alpha/2}\sqrt{k'_{LS,B}\Sigma k_{LS,B}}$ is optimal in the same sense as the usual CI in Section 4.3.1: it achieves minimax expected length, and its efficiency, relative to a CI that optimizes its length at a single θ^* and $\gamma = 0$, is lower-bounded by $z_{1-\alpha}/z_{1-\alpha/2}$. Theorem 4.1 formally translates the efficiency bound from the limiting model to the GMM model, so that the usual two-sided CI based on $h(\hat{\theta}_{W(B)})$ is asymptotically efficient in the same sense as the usual CI based on $h(\hat{\theta}_{\Sigma^{-1}})$ discussed in Section 4.3.1 under correct specification. Just as with the results in Section 4.3.1, this asymptotic result is, to our knowledge, new. The one-sided analog follows from the results in Appendix C. These results stand in sharp contrast to the results for estimation, where the MSE improvement at small values of γ may be substantial.

An important consequence of these results is that asymptotically valid one-sided CIs

based on shrinkage or model-selection procedures, such as one-sided versions of the CIs proposed in Andrews and Guggenberger (2009), DiTraglia (2016) or McCloskey (2019) must have *worse* excess length performance than the usual one-sided CI based on the GMM estimator $h(\hat{\theta}_{W(B)})$ that uses valid moments only. While it is possible to construct two-sided CIs that improve upon the usual CI based on $h(\hat{\theta}_{W(B)})$ at particular values of θ and γ , the scope for such improvement is smaller than the ratio of one- to two-sided critical values. Furthermore, any such improvement must come at the expense of worse performance at other points in the parameter space.⁷ Therefore, in order to tighten CIs based on valid moments only, it is *necessary* to make a priori restrictions on the potential misspecification of the remaining moments.

4.3.3 General case

According to the results in Section 4.3.2, one must place a priori bounds on the amount of misspecification in order to use misspecified moments. This leads us to the general case, where we place the local misspecification vector c in some set \mathcal{C} that is not necessarily a linear subspace. One can then form a CI centered at an estimate formed from these misspecified moments using the methods in Section 3. In the case where \mathcal{C} is convex and centrosymmetric, Theorem 4.1 shows that this CI is near optimal, in the sense that no other CI can improve upon it by more than a factor of κ_* , even in the favorable case of correct specification. Since the width of the CI is asymptotically constant under local parameter sequences $\theta_n \rightarrow \theta^*$ and sufficiently regular probability distributions $P_n \rightarrow P_0$ (for example, $P_n \rightarrow P_0$ along submodels satisfying Assumption C.1), this also shows that the CI is near optimal in a local minimax sense. In the general case, Theorem 4.1, as well as the analogous results for one-sided CIs in Appendix C are, to our knowledge, new.

As we discuss in Section 3, we recommend reporting results for a range of sets $\mathcal{C}(M)$ indexed by a scalar M that bounds the magnitude of misspecification. One may instead wish to report a single CI based on a data-driven estimate of M , for example, by using a first-stage J test to assess plausible magnitudes of misspecification. Formally, one would seek a CI that is valid over $\mathcal{C}(\overline{M})$ while improving length when in fact $\|\gamma\| \ll \overline{M}$, where \overline{M} is some initial conservative bound. When \mathcal{C} is convex and centrosymmetric, Theorem 4.1 shows that the scope for such improvements is limited: the average length of any such CI cannot be much smaller than the CI that uses the most conservative choice \overline{M} , even when $c = 0$. The impossibility of choosing M based on the data is related to the impossibility of

⁷Consistently with these results, in a simulation study considered in DiTraglia (2016), the post-model selection CI that he proposes is shown to be wider on average than the usual CI around a GMM estimator that uses valid moments only.

using specification tests to form an upper bound for M . On the other hand, it is possible to obtain a lower bound for M using such tests. We develop lower CIs for M in Appendix B.

4.3.4 Cressie-Read divergences

Andrews et al. (2018) have shown that defining misspecification in terms of the magnitude of any divergence in the Cressie and Read (1984) family leads to a set \mathcal{C} that asymptotically takes the form $\mathcal{C} = \{\Sigma^{1/2}\gamma : \|\gamma\|_2 \leq M\} = \{c: c'\Sigma^{-1}c \leq M^2\}$. The Cressie-Read family includes the Hellinger distance used by Kitamura et al. (2013), who consider minimax point estimation among estimators satisfying certain regularity conditions. Since this set \mathcal{C} takes the form discussed in Remark 3.1 with $p = 2$ and $B = \Sigma^{1/2}$, it follows from the discussion in Remark 3.1 that the optimal sensitivity corresponds to the GMM estimator with weighting matrix $(\lambda BB' + \Sigma)^{-1} = (\lambda + 1)^{-1}\Sigma^{-1}$. Since this is proportional to the weighting matrix Σ^{-1} that is optimal under correct specification, we obtain the same optimal sensitivity $k'_{LS,0} = -H(\Gamma'\Sigma^{-1}\Gamma)^{-1}\Gamma'\Sigma^{-1}$ as in the correctly specified case discussed in Section 4.3.1. As we show in Appendix A.1, this form of \mathcal{C} leads to a closed form solution for the efficiency bound κ_* .

The results above imply that any estimator with sensitivity $k_{LS,0}$ is near-optimal for CI construction. In line with these results, the estimator in Kitamura et al. (2013) has sensitivity $k_{LS,0}$. Thus, the usual GMM estimator $h(\hat{\theta}_{\Sigma^{-1}})$ and the estimator in Kitamura et al. (2013) are both near-optimal for CI construction, even if one allows for arbitrary CIs that are not necessarily centered at estimators that satisfy the regularity conditions in Kitamura et al. (2013). Also, because they have the same sensitivity, under this form of misspecification, the usual GMM estimator $h(\hat{\theta}_{\Sigma^{-1}})$ and the estimator in Kitamura et al. (2013) have the same local asymptotic minimax properties.

4.4 Extensions: asymmetric constraints and constraints on θ

If the set \mathcal{C} is convex but asymmetric (such as when \mathcal{C} includes bounds on a norm as well as sign restrictions, or when \mathcal{C} includes equality and sign restrictions, as in Moon and Schorfheide (2009)), one can still apply bounds from Armstrong and Kolesár (2018) to the limiting model described in Section 4.1. Our general asymptotic efficiency bounds in Appendix C translate these results to the locally misspecified GMM model so long as \mathcal{C} is convex. Since the negative implications for efficiency improvements under correct specification use centrosymmetry of \mathcal{C} , introducing asymmetric restrictions, such as sign restrictions, is one possible way of getting efficiency improvements at some smaller set $\mathcal{D} \subseteq \mathcal{C}$ while maintaining coverage over \mathcal{C} . We derive efficiency bounds and optimal CIs for this problem in Appendix C. Interestingly, the scope for efficiency improvements can be different for one-

and two-sided CIs, and can depend on the direction of the CI in this case. To get some intuition for this, note that, in the instrumental variables model with a single instrument and single endogenous regressor, sign restrictions on the covariance of an instrument with the error term can be used to sign the direction of the bias of the instrumental variables estimator, which is useful for forming a one-sided CI only in one direction.

Finally, while we focus on restrictions on c , one can also incorporate local restrictions on θ . Our general results in Appendix C give efficiency bounds that cover this case. Similar to the discussion above, these results have implications for using prior information about θ to determine the amount of misspecification, or to shrink the width of a CI directly. In particular, while it is possible to use prior information on θ (say, an upper bound on $\|\theta\|$ for some norm $\|\cdot\|$) to shrink the width of the CI, the width of the CI and the estimator around which it is centered must depend on the a priori upper bounds on the magnitude of θ and c when this prior information takes the form of a convex, centrosymmetric set for $(\theta', c)'$. This rules out, for example, choosing the moments based on whether the resulting estimate for θ is in a plausible range.

5 Applications

This section describes particular applications of our approach, along with a discussion of implementation details appropriate to each application.

5.1 Instrumental variables

The single equation linear instrumental variables (IV) model is given by

$$y_i = x_i' \theta_0 + \varepsilon_i \tag{21}$$

where, in the correctly specified case, $E\varepsilon_i z_i = E(y_i - x_i' \theta_0) z_i = 0$, with z_i a d_g -vector of instruments. This is an instance of a GMM model with $g(\theta) = E(y_i - x_i' \theta) z_i$ and $\hat{g}(\theta) = \frac{1}{n} \sum_{i=1}^n z_i (y_i - x_i' \theta)$.

One common reason for misspecification in this model is that the instruments do not satisfy the exclusion restriction, because they appear directly in the structural equation (21), so that $\varepsilon_i = z_{I_i}' \gamma / \sqrt{n} + \eta_i$, where $E[z_i \eta_i] = 0$, and z_{I_i} corresponds to a subset I of the instruments, the validity of which one is worried about. This form of misspecification has previously been considered in a number of papers, including [Hahn and Hausman \(2005\)](#), [Conley et al. \(2012\)](#), and [Andrews et al. \(2017\)](#), among others. Bounding the norm of γ using some norm $\|\cdot\|$ then leads to the set given in Eq. (12), with $B = E[z_i z_{I_i}']$.

Although the matrix B is unknown, for the purposes of estimating the optimal sensitivity and constructing asymptotically valid CIs, it can be replaced by the sample analog $\hat{B} = n^{-1} \sum_{i=1}^n z_i z_i'$. This does not affect the asymptotic validity or coverage properties of the resulting CI. Under this setup, the parameter M bounds that magnitude of γ , the direct effect of the instruments on the outcome. Therefore, the appropriate choice of M will depend on the plausible magnitude of these direct effects—see, for example, [Conley et al. \(2012\)](#) for examples and a discussion.

The linearity of the moment condition leads to simplifications in our implementation in [Section 3](#). In [Step 1](#), as the initial estimator, one can use the two-stage least squares (2SLS) estimator

$$\hat{\theta}_{\text{initial}} = \left[\left(\sum_{i=1}^n z_i x_i' \right)' \left(\sum_{i=1}^n z_i z_i' \right)^{-1} \left(\sum_{i=1}^n z_i x_i' \right) \right]^{-1} \left(\sum_{i=1}^n z_i x_i' \right)' \left(\sum_{i=1}^n z_i z_i' \right)^{-1} \sum_{i=1}^n z_i y_i.$$

This leads to the estimates $\hat{\Gamma} = -\frac{1}{n} \sum_{i=1}^n z_i x_i'$ and $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \hat{\theta}_{\text{initial}})^2 z_i z_i'$. Alternatively, if we assume homoskedasticity, we can use the estimator $\hat{\Sigma}_H = \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \hat{\theta}_{\text{initial}})^2 \cdot \frac{1}{n} \sum_{i=1}^n z_i z_i'$. In the correctly specified case, the 2SLS estimator is only optimal under homoskedasticity, while the GMM estimator with weighting matrix $\hat{\Sigma}^{-1}$ is optimal in general. Due to concerns with finite sample performance, however, it is common to use the 2SLS estimator along with standard errors based on a robust variance estimate, even when heteroskedasticity is suspected. Mirroring this practice, one can use $\hat{\Sigma}_H$ when forming the optimal sensitivity $\hat{k}_{\lambda_M^*}$ and worst-case bias in [Step 2](#), but use the robust variance estimate $\hat{\Sigma}$ in [Step 3](#) when forming the final CI in [Eq. \(11\)](#). The CI will then be optimal under homoskedasticity, but it will remain valid under heteroskedasticity, just like the usual CI based on 2SLS with robust standard errors in the correctly specified case.

If the parameter of interest linear in θ , $h(\theta) = H\theta$, then the one-step estimator $\hat{h}_{\lambda_M^*}$ in [Step 3](#) does not depend on the choice of the initial estimator (except possibly through the estimate of Σ when forming the desired sensitivity):

$$\begin{aligned} \hat{h}_{\lambda_M^*} &= H \hat{\theta}_{\text{initial}} + \hat{k}'_{\lambda_M^*} \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \hat{\theta}_{\text{initial}}) z_i = \hat{k}'_{\lambda_M^*} \frac{1}{n} \sum_{i=1}^n y_i z_i + \left(H - \hat{k}'_{\lambda_M^*} \frac{1}{n} \sum_{i=1}^n z_i x_i' \right) \hat{\theta}_{\text{initial}} \\ &= \hat{k}'_{\lambda_M^*} \frac{1}{n} \sum_{i=1}^n y_i z_i, \end{aligned}$$

where the second line follows since the sensitivities $\hat{k}_{\lambda_M^*}$ satisfy $H = -\hat{k}'_{\lambda_M^*} \hat{\Gamma} = \hat{k}'_{\lambda_M^*} \frac{1}{n} \sum_{i=1}^n z_i x_i'$. Since the estimator $\hat{h}_{\lambda_M^*}$ is linear, the worst-case bias calculations are the same under global misspecification, when the magnitude M of γ in [Eq. \(12\)](#) grows at the rate \sqrt{n} . By using

variance estimates that are valid under global misspecification in place of the variance estimate $\hat{k}'_{\lambda^*} \hat{\Sigma} \hat{k}_{\lambda^*}$ in the CI construction, one can ensure that the resulting CI also remains valid under global misspecification. See Appendix D.2 for details.

Remark 5.1. This framework can also be used to incorporate a priori restrictions on the magnitude of coefficients on control variables in an instrumental variables regression. Suppose that we have a set of controls w_i , that appear in the structural equation (21), so that $y_i = x_i'\theta + w_i'\gamma/\sqrt{n} + \epsilon_i$, and ϵ_i is uncorrelated with w_i as well as vector of instruments \tilde{z}_i . If one is willing to restrict the magnitude of the coefficient vector γ , so that $\|\gamma\| \leq M$, then one can add w_i to the original vector of instruments \tilde{z}_i , $z_i = (\tilde{z}_i', w_i')'$. For example, if one is concerned with functional form misspecification, one can define the control variables to be higher order series terms. We then obtain the misspecified IV model with the set \mathcal{C} given by (12), with $B = E[z_i w_i']$. Thus, we can interpret this model as a locally misspecified version of a model with w_i used as an excluded instrument.

Remark 5.2. Instead of bounding the coefficient vector γ , one can alternatively bound the magnitude of the direct effect $z_{Ii}'\gamma$. If all instruments are potentially invalid, $z_{Ii} = z_i$, and one sets $\mathcal{C} = \{\gamma: E[(z_i'\gamma)^2] \leq M\}$, then under homoscedasticity, this corresponds to the case discussed in Section 4.3.4, where the uncertainty from potential misspecification is exactly proportional to the asymptotic sampling uncertainty in $\hat{g}(\theta)$. Consequently, in this case the optimal sensitivity is the same as that given by the 2SLS estimator.

5.2 Omitted variables bias in linear regression

Specializing to the case where $z_i = x_i$, the misspecified IV model of Section 5.1 gives a misspecified linear regression model as a special case. This can be used to assess sensitivity of regression results to issues such as omitted variables bias. In particular, consider the linear regression model

$$y_i = x_i'\theta + w_i^* + \tilde{\epsilon}_i, \quad E x_i \tilde{\epsilon}_i = 0$$

where x_i and y_i are observed and w_i^* is a (possibly unobserved) omitted variable. Correlation between w_i^* and x_i will lead to omitted variables bias in the OLS regression of y_i on x_i . If w_i^* is unobserved, then we obtain our framework by making the assumption $\sqrt{n} E w_i^* x_i \in \mathcal{C}$, for some set \mathcal{C} , and letting $\hat{g}(\theta) = \frac{1}{n} \sum_{i=1}^n x_i (y_i - x_i'\theta)$. This setup can also cover choosing between different sets of control variables. Suppose that $w_i^* = w_i'\gamma$, where w_i is a vector of observed control variables that the researcher is considering not including in the regression. If γ is unrestricted, then by the results in Section 4.3.2, the long regression of y_i on both x_i and w_i yields nearly optimal CIs. If one is willing to restrict the magnitude of γ , it is possible

to tighten these CIs, with the setting reducing to that in Remark 5.1, with $z_i = (x_i, w_i)'$. The same framework can be used to incorporate selection bias by defining w_i^* to be the inverse Mills ratio term in the formula for $E[y_i | x_i, i \text{ observed}]$ in Heckman (1979).

5.3 Functional form misspecification

Our setup allows for misspecification in moment conditions arising from functional form misspecification. To apply our setup, one must relate this misspecification to the bounds \mathcal{C} on the moment conditions at the true parameter value. One approach to bounding functional form misspecification is to use smoothness conditions from the nonparametric statistics literature, such as bounds on derivatives (see, for example, Tsybakov, 2009, for an introduction to this literature). Since these sets are typically convex (taking a convex combination of two functions that satisfy a given bound on a given derivative gives a function that also satisfies this bound), they typically lead to convex sets \mathcal{C} , so that our framework can be applied.

As a simple example, consider a nonparametric IV model with discrete covariates:

$$E[y_i - m(x_i) | z_i] = 0.$$

Suppose x takes values in the finite set $\mathcal{X} = \{\tilde{x}_1, \dots, \tilde{x}_{N_x}\}$ and z_i takes values in the finite set $\mathcal{Z} = \{\tilde{z}_1, \dots, \tilde{z}_{N_z}\}$. This setting was considered by Freyberger and Horowitz (2015), who place only nonparametric smoothness or shape restrictions on the unknown function m . To see the connection with our setting, we note that such restrictions can be interpreted as bounds on specification error from a parametric model. If one models these restrictions as local to a parametric family, one obtains our setting. In particular, let $m(x_i) = f(x_i, \theta_0) + n^{-1/2}r(x_i)$, $r \in \mathcal{R}$, where \mathcal{R} is a nonparametric smoothness class. For example, if x_i is univariate, we can let $f(x_i, \theta) = \theta_1 + \theta_2 x_i$ and define \mathcal{R} to be the class of functions with $r(0) = r'(0) = r''(0)$ and second derivative bounded by some constant M . This is equivalent to placing the bound $n^{-1/2}M$ on the second derivative of $m(\cdot)$, which corresponds to a Hölder smoothness class. We can then map this to a misspecified GMM model, with the j th element of the moment function given by $g_j(x_i, y_i, \theta) = (y_i - f(x_i, \theta_0))I(z_i = \tilde{z}_j)$ and j th element of the misspecification vector c given by $Er(x_i)I(z_i = \tilde{z}_j) = \sum_{\tilde{x} \in \mathcal{X}} r(\tilde{x})P(x_i = \tilde{x}, z_i = \tilde{z}_j)$. Stacking these equations, we see that $c = B\gamma$ where B is a matrix composed of the elements $P(x_i = \tilde{x}, z_i = \tilde{z}_j)$ and $\gamma = (r(\tilde{x}_1), \dots, r(\tilde{x}_{N_x}))'$. As with the IV setting in Section 5.1, B is unknown, but can be replaced by a consistent estimate based on the sample analogue. So long as the set \mathcal{R} is convex, we obtain convex restrictions on γ and therefore c , so that our framework applies.

This example brings up an important point about the interpretation of $h(\theta)$. If the

object of interest is a functional of $m(x) = f(x, \theta_0) + n^{-1/2}r(x)$, then we will need to allow the object of interest $h(\cdot)$ to depend on the misspecification vector directly, as well as on θ . As discussed at the beginning of Section 2, this falls into a mild extension of our framework. Alternatively, under a suitable parametrization of f and r , it is often possible to define the object of interest to be function of θ alone. For example, if we are interested in the derivative $m'(x_0)$ at a particular point x_0 under a bound on the second derivative of $m(\cdot)$, we can let $f(x, \theta) = \theta_1 + \theta_2 x$ and define \mathcal{R} to be the class of functions with $r(x_0) = r'(x_0) = r''(x_0) = 0$ and second derivative bounded by M . Then $m'(x_0) = \theta_2$.

5.4 Treatment effect extrapolation

Often, the average effect of a counterfactual policy on a particular subset of a population is of interest, and we would like to weaken the assumptions under which this effect is point-identified. We have available estimates $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_m)'$ of the parameter τ , with $\hat{g}(\theta) = \hat{\tau} - A\theta$, and $g(\theta) = \tau - A\theta$ for a known matrix A . We would like to extrapolate from these estimates to learn about the parameter of interest $h(\theta_0) = H\theta_0$. The potential extrapolation bias is captured by the assumption that $g(\theta_0) \in \tilde{\mathcal{C}}_n$, some convex set.

Note that, because the moment condition is linear in the parameter of interest, asymptotic validity of our CIs does not require that the set $\tilde{\mathcal{C}}_n$ takes the form $\tilde{\mathcal{C}}_n = \mathcal{C}/\sqrt{n}$. CIs given in Eq. (7) based on linear estimators of the form $\hat{h} = k'\hat{\tau}$ (such as minimum distance estimators), with $\mathcal{C} = \sqrt{n}\tilde{\mathcal{C}}_n$ are valid under both local and global misspecification (i.e. under the assumption that the set $\tilde{\mathcal{C}}_n$ is fixed as $n \rightarrow \infty$).

One example that falls into this setup are differences-in-differences designs when the parallel trends assumption is violated. Here there are m time periods, with treatment taking place in period T_0 . The $(m - T_0)$ -vector θ corresponds to a vector of dynamic treatment effects on the treated, $A\theta = (0', \theta)'$, and $g(\theta_0)$ is a vector of trend differences between the treated and untreated, with $g(\theta_0) = 0$ if the parallel trends assumption holds. [Rambachan and Roth \(2019\)](#) build on the framework in this paper to develop CIs in this setting.

Another example that has been of recent interest involves nonseparable models with endogeneity. Under conditions in [Imbens and Angrist \(1994\)](#) and [Heckman and Vytlacil \(2005\)](#), instrumental variables estimates $\hat{\tau}_m$ with different instruments are consistent for average treatment effects for different subpopulations. A recent literature ([Kowalski, 2016](#); [Brinch et al., 2017](#); [Mogstad et al., 2018](#)) has focused on using assumptions on treatment effect heterogeneity to extrapolate these estimates to other populations. Our framework applies if these assumptions amount to placing the differences between the estimated treatment effects and the effect of interest in a known convex set.

6 Empirical application

This section illustrates the confidence intervals developed in Section 2 in an empirical application to automobile demand based on the data and model in [Berry et al. \(1995, BLP hereafter\)](#). We use the version of the model as implemented by [Andrews et al. \(2017\)](#), who calculate the asymptotic bias of the GMM estimator with weighting matrix Σ^{-1} under local misspecification in this setting.⁸

6.1 Model description and implementation

In this model, the utility of consumer i from purchasing a vehicle j , relative to the outside option, is given by a random-coefficient logit model $U_{ij} = \sum_{k=1}^K x_{jk}(\beta_k + \sigma_k v_{ik}) - \alpha p_j / y_i + \xi_j + \epsilon_{ij}$, where p_j is the price of the vehicle, x_{jk} the k th observed product characteristic, ξ_j is an unobserved product characteristic, and ϵ_{ij} is has an i.i.d. extreme value distribution. The income of consumer i is assumed to be log-normally distributed, $y_i = e^{m+\varsigma v_{i0}}$, where the mean m and the variance ς of log-income are assumed to be known and set to equal to estimates from the Current Population Survey. The unobservables $v_i = (v_{i0}, \dots, v_{iK})$ are i.i.d. standard normal, while the distribution of the unobserved product characteristic ξ_j is unrestricted.

The marginal cost mc_j for producing vehicle j is given by $\log(mc_j) = w'_j \nu + \omega_j$, where w_j are observable characteristics, and ω_j is an unobservable characteristic. The full vector of model parameters is given by $\theta = (\sigma', \alpha, \beta', \nu)'$. Given this vector, and given a vector of unobservable characteristics, one can compute the market shares implied by utility maximization, which can be inverted to yield the unobservable characteristic as a function of θ , $\xi_j(\theta)$. One can similarly invert the unobserved cost component, writing it as a function of θ , $\omega_j(\theta)$, under the assumption that firms set prices to maximize profits in a Bertrand-Nash equilibrium. Given a vector z_{dj} of demand-side instruments, and a vector z_{sj} of supply-side instruments, this yields the moment condition $g(\theta) = E[\hat{\gamma}(\theta)]$, where

$$\hat{g}(\theta) = \frac{1}{n} \sum_{j=1}^n \begin{pmatrix} z_{dj} \xi_j(\theta) \\ z_{sj} \omega_j(\theta) \end{pmatrix}.$$

The BLP data spans the period 1971 to 1990, and includes information on essentially all $n = 999$ models sold during that period (for simplicity, we have suppressed the time dimension in the description above). There are 5 observable characteristics x_j : a constant,

⁸The dataset for this empirical application has been downloaded from the [Andrews et al. \(2017\)](#) replication files, available at <https://doi.org/10.7910/DVN/LLARSN>.

horsepower per 10 pounds of weight (HPWt), a dummy for whether air-conditioning is standard (Air), mileage per 10 dollars (MP\$) defined as MPG over average gas price in a given year, and car size (Size), defined as length times width. The vector z_{dj} consists of x_j , plus the sum of x_j across models other than j produced by the same firm, and for rival firms. There are 6 cost variables w_j : a constant, log of HPWt, Air, log of MPG, log of Size, and a time trend. The vector z_{sj} consists of these variables, MP\$, and the sums of w_j for own-firm products other than j , and for rival firms. After excluding collinear instruments, this gives a total of $d_g = 31$ instruments, 25 of which are excluded to identify $d_\theta = 17$ model parameters. The parameter of interest is average markup, $h(\theta) = \frac{1}{n} \sum_j (p_j - mc_j(\theta))/p_j$.

One may worry that some of these instruments are invalid, because elements of z_{dj} or z_{sj} may appear directly in the utility or cost function with the coefficient on the ℓ th element given by $\delta_{d\ell}\gamma_{d\ell}/\sqrt{n}$ or $\delta_{s\ell}\gamma_{s\ell}/\sqrt{n}$, respectively. Here $\delta_{d\ell}$ and $\delta_{s\ell}$ are scaling constants so that, given the sample size $n = 999$ at hand, $\gamma_{d\ell}$ has the interpretation that the consumer willingness to pay for one standard deviation change in the ℓ th demand-side instrument $z_{dj\ell}$ is $\gamma_{d\ell}\%$ of the average 1980 car price, and changing the ℓ th supply-side instrument $z_{sj\ell}$ by one standard deviation changes the marginal cost by $\gamma_{s\ell}\%$ of the average car price. [Andrews et al. \(2017\)](#) use this scaling in their sensitivity analysis, and they discuss economic motivation for concerns about this form of misspecification. By way of comparison, the estimates of the parameters β and ν in the utility and cost function imply that consumers are on average willing to pay between 2.2 and 10.0% of the average car price for a standard deviation change in one of the included car characteristics, and that a standard deviation change in the included cost characteristics changes the marginal cost by between 3.8 and 11.1% of the average car price. We therefore interpret specifications of the set \mathcal{C} that allow for $|\gamma_{s\ell}| \approx 1-2$ (or $|\gamma_{d\ell}| \approx 1-2$) as allowing for moderate amounts of misspecification in the ℓ th supply-side (or demand-side) instrument.

We follow the implementation in Section 3. Given a set I of potentially invalid instruments, we follow Remark 3.1 and consider sets \mathcal{C} of the form (12), with $\|\cdot\|$ corresponding to an ℓ_p norm with $p \in \{2, \infty\}$, and $B = \tilde{B}_I \cdot \#I^{1/p}$, where \tilde{B}_I is given by the columns of

$$\tilde{B} = \begin{pmatrix} E[z_{dj}z'_{dj}] \text{diag}(\delta_{d\ell}) & 0 \\ 0 & E[z_{sj}z'_{sj}] \text{diag}(\delta_{d\ell}) \end{pmatrix},$$

and $\#I$ is the number of potentially invalid instruments. The scaling by $(\#I)^{1/p}$ ensures that the vector $\gamma = M(1, \dots, 1)'$ is always included in the set. [Andrews et al. \(2017\)](#) report the sensitivity of the usual GMM estimator under this form of misspecification, considering misspecification in each instrument individually (so that I contains a single element), and setting $M = 1$. However, if one is concerned about the validity of several instruments, it is

natural to allow I to contain all instruments the validity of which is questionable. In our analysis, we vary the set of potentially misspecified instruments. We also vary M in order to assess the sensitivity of conclusions to different amounts of misspecification. As we will see below, different choices of \mathcal{C} lead to different sensitivities for the optimal estimator, and using the optimal sensitivity can reduce the width of the CI substantially relative to CIs based on the usual GMM estimator.

We use the estimate $\hat{\theta}_{\text{initial}}$ that corresponds to the GMM estimator based on the weight matrix that's optimal under correct specification, as reported in [Andrews et al. \(2017\)](#), and the estimates $\hat{\Gamma}$, \hat{H} , and $\hat{\Sigma}$ are computed following Step 1 of the implementation.

6.2 Results

To illustrate that using the sensitivity that is optimal under local misspecification can yield substantially tighter CIs, Figure 1 plots the confidence intervals based on the optimal sensitivity, as well as those based on $\hat{\theta}_{\text{initial}}$ under different sets I of potentially invalid instruments and ℓ_2 constraints on γ . It is clear from the figure that using the optimal sensitivity yields substantially tighter confidence intervals, relative to simply adjusting the usual CI by using the critical value $\text{cv}_\alpha(\cdot)$ to take into account the potential bias of $h(\hat{\theta}_{\text{initial}})$, by as much as a factor of 3.4. The intuitive reason for this is that by adjusting the sensitivity of the estimator, it is possible to substantially reduce its bias at little cost in terms of an increase in variance. Thus, for example, while the CI for the average markup based on the estimate $\hat{\theta}_{\text{initial}}$ is essentially too wide to be informative when the set of potentially invalid instruments corresponds to all excluded instruments, the CI based on the optimal sensitivity, [46.0, 66.0]%, is still quite tight.

If a researcher is ex ante unsure what form of misspecification one should worry about, as a sensitivity check, it is useful to consider the effects of different forms of misspecification. In Figure 2, we plot the optimal confidence intervals for different subsets of invalid instruments, under both ℓ_2 and ℓ_∞ norms for γ . Although the choice of norm matters when the number of potentially misspecified instruments is greater than one, the results are qualitatively similar. Comparing the results for different choices of the set of potentially invalid instruments suggests that allowing supply-side instruments to be invalid generally increases the average markup estimate, while allowing demand-side instruments to be invalid has the opposite effect.

As it may be ex ante unclear what magnitude of misspecification is reasonable to allow for, as discussed in Section 3, it is useful to plot the optimal CI for multiple choices of M . We do this in Figure 3 for $p = 2$, and we allow all excluded instruments to be potentially invalid.

One can see that while the CI is unstable for values of M smaller than about 0.4, for larger values of M , the estimate is quite stable and equal to about 50%. Even at $M = 2$, one rejects the hypothesis that the optimal markup is equal to the initial estimate $h(\hat{\theta}_{\text{initial}}) = 32.7\%$. This suggests that ignoring misspecification in the BLP model likely leads to a downward bias in the estimate of the average markup. At the same time, it is possible to obtain reasonably tight CIs for the average markup even under a moderate amount of misspecification.

The J -statistic for testing the hypothesis that all moments are correctly specified equals 426.7. Consequently, the hypothesis is rejected at the usual significance levels. Furthermore, it can be seen from Figure 2 that the CIs for “all excluded” (that allow all excluded instruments to be invalid at $M = 1$), and “all excluded demand” (that assume validity of supply-side instruments) do not overlap. This implies that either the misspecification in the demand-side instruments must be greater than 1% of the average care price ($M = 1$), or else the supply-side instruments must also be invalid. Table 1 implements the specification test from Appendix B that gives lower CI $[M_{\min}, \infty]$ for M . The results suggest that if one assumes only a subset of the instruments is invalid, the misspecification in the potentially invalid instruments must be quite large. For example, if we assume that all instruments are valid except potentially the demand-side instruments based on rival firms’ product characteristics, then the misspecification in these instruments must be greater than $M = 5.36$. If we allow all instruments to be invalid, then $M \geq 1.13$.

Finally, to illustrate the implication of Theorem 4.1 that one cannot substantively improve upon the CIs that we construct, we calculate the efficiency bound κ_* for these CIs in Table 2. The table shows that the bound is at least as high as the efficiency bound for the usual CI under correct specification (given in (19) and equal to 84.99% at $\alpha = 0.05$). Thus, the asymptotic scope for improvement over the CIs reported in Figure 2 at particular values of θ and $c = 0$ is even smaller than the scope for improvement over the usual CI at particular values of θ under correct specification.

A Details of calculations

This appendix contains the details of calculations for particular sets \mathcal{C} discussed in Section 4.3.4 and Remark 3.1.

A.1 Cressie-Read divergences

Consider the problem (15) under constraints of the form $\{c: c\Sigma^{-1}c \leq M^2\}$. The Lagrangian for this problem can be written as

$$2H\theta + \lambda_1(\delta^2/4 - (c - \Gamma\theta)'\Sigma^{-1}(c - \Gamma\theta)) + \lambda_2(M^2 - c'\Sigma^{-1}c).$$

(we multiply the objective function by 2 so that its optimized value equals $\omega(\delta)$). The first-order conditions imply that at optimum, $c = \frac{\lambda_1}{\lambda_1 + \lambda_2}\Gamma\theta$, and $\theta = \frac{\lambda_1 + \lambda_2}{\lambda_1\lambda_2}(\Gamma'\Sigma^{-1}\Gamma)^{-1}H'$. Plugging these expressions into the constraints yields $M^2 = H(\Gamma'\Sigma^{-1}\Gamma)^{-1}H'/\lambda_2^2$ and $\delta^2/4 = H(\Gamma'\Sigma^{-1}\Gamma)^{-1}H'/\lambda_1^2$. Since $H(\Gamma'\Sigma^{-1}\Gamma)^{-1}H' = k'_{LS,0}\Sigma k_{LS,0}$, solving for λ_1 and λ_2 , and plugging into the expression for θ yields

$$\theta = \frac{\delta/2 + M}{\sqrt{k'_{LS,0}\Sigma k_{LS,0}}} \cdot (\Gamma'\Sigma^{-1}\Gamma)^{-1}H'.$$

Thus, $\omega(\delta) = 2H\theta = (\delta + 2M)\sqrt{k'_{LS,0}\Sigma k_{LS,0}}$. With this form of ω , the bound in (18) becomes

$$\kappa_*(H, \Gamma, \Sigma, \mathcal{C}) = \frac{(1 - \alpha)(z_{1-\alpha} + M) + \phi(z_{1-\alpha})}{\text{cv}_\alpha(M)}.$$

This efficiency equals at least $\min\{\kappa_{*,\alpha}^L, 1 - \alpha\}$, where $\kappa_{*,\alpha}^L = ((1 - \alpha)z_{1-\alpha} + \phi(z_{1-\alpha}))/z_{1-\alpha/2}$ denotes the efficiency given in (19) when \mathcal{C} is a linear subspace. To show this, observe that $\text{cv}'_\alpha(M) \leq 1$ for all $M \geq 0$. Therefore, the derivative of

$$(1 - \alpha)(z_{1-\alpha} + M) + \phi(z_{1-\alpha}) - \min\{1 - \alpha, \kappa_{*,\alpha}^L\} \text{cv}_\alpha(M)$$

with respect to M , given by $1 - \alpha - \min\{1 - \alpha, \kappa_{*,\alpha}^L\} \text{cv}'_\alpha(M)$, is always non-negative. Since the expression in the above display equals $(\kappa_{*,\alpha}^L - \min\{\kappa_{*,\alpha}^L, 1 - \alpha\})z_{1-\alpha/2} \geq 0$ at $M = 0$, it follows that it is always non-negative. Rearranging it then yields $\kappa_*(H, \Gamma, \Sigma, \mathcal{C}) \geq \min\{\kappa_{*,\alpha}^L, 1 - \alpha\}$. Furthermore, it follows from Eq. (50) below that the efficiency of one-sided CIs at $c = 0$ is given by $\kappa_*^{\text{OCI},\beta} = 1$.

A.2 ℓ_p Bounds

We now consider the form of the optimal sensitivity when the set $\mathcal{C} = \mathcal{C}(M)$ takes the form in Eq. (12), and $\|\cdot\|$ corresponds to an ℓ_p norm, as discussed in Remark 3.1. First, we explain the connection with penalized estimation. Since $c = B\gamma$, one can write the approximately

linear model (13) as

$$Y = -\Gamma\theta + B\gamma + \Sigma^{1/2}\varepsilon,$$

which one can think of as a regression model with correlated errors, design matrix $(-\Gamma, B)$, and coefficient vector $(\theta', \gamma)'$. With this interpretation, it is clear that if the number of regressors $d_\theta + d_\gamma$ is greater than the number of observations d_g , the constraint on the norm of γ is necessary to make the model informative. Using the observation from Remark 3.1 that $\overline{\text{bias}}_{\mathcal{C}(1)}(k) = \|B'k\|_{p'}$, it follows that the optimization problem (10) under $\mathcal{C} = \mathcal{C}(1)$ is equivalent to

$$\min_k k'\Sigma k \quad \text{s.t.} \quad H = -k'\Gamma \quad \text{and} \quad \|B'k\|_{p'} \leq \overline{B}. \quad (22)$$

We now specialize the results to the cases $p = 2$ and $p = \infty$. We discuss the case $p = 1$ in a working paper version of this paper (Armstrong and Kolesár, 2020).

A.2.1 $p = 2$

In this case, the Lagrangian form of (22) becomes

$$\min_k k'(\Sigma + \lambda BB')k \quad \text{s.t.} \quad H = -k'\Gamma,$$

with the Lagrange multiplier λ giving the relative weight on bias. Optimizing this objective is isomorphic to deriving the minimum variance unbiased estimator of $H\theta$ in a regression model with design matrix $-\Gamma$ and variance $\Sigma + \lambda BB'$, so the Gauss-Markov theorem implies that the optimal sensitivities are $k'_\lambda = -H(\Gamma'W_\lambda\Gamma)^{-1}\Gamma'W_\lambda$ where $W_\lambda = [\Sigma + \lambda BB']^{-1}$.

A.2.2 $p = \infty$

Write the Lagrangian form of (22) as

$$\min_k k'\Sigma k/2 + \lambda\|B'k\|_1 \quad \text{s.t.} \quad H = -k'\Gamma. \quad (23)$$

It will be convenient to transform the problem so that the ℓ_1 constraint only involves d_γ elements of k . Let

$$T = \begin{pmatrix} B'_\perp \\ (B'B)^{-1}B' \end{pmatrix}, \quad T^{-1} = \begin{pmatrix} B_\perp & B \end{pmatrix}, \quad (24)$$

where B_\perp is an orthonormal matrix that's orthogonal to B . Then, since $TB = (0, I_{d_\gamma})'$, the above minimization problem is equivalent to the problem

$$\min_\kappa \kappa'S\kappa/2 + \lambda \sum_{i \in I} |\kappa_i| \quad \text{s.t.} \quad H' = -G'\kappa,$$

where $\kappa = T'^{-1}k$, $S = T\Sigma T'$, $G = T\Gamma$, and $I = \{d_g - d_\gamma, \dots, d_g\}$ indexes the last d_γ elements of κ .

To minimize the above display and give the solution path as λ varies, we use arguments similar to those in Theorem 2 of [Rosset and Zhu \(2007\)](#). For $i \in I$, write $\kappa_i = \kappa_{+,i} - \kappa_{-,i}$, where $\kappa_{+,i} = \max\{\kappa_i, 0\}$ and $\kappa_{-,i} = -\min\{\kappa_i, 0\}$. We minimize the objective function in the preceding display over $\{\kappa_{+,i}, \kappa_{-,i}, \kappa_j : i \in I, j \notin I\}$ subject to the constraints $\kappa_{+,i} \geq 0$ and $\kappa_{-,i} \geq 0$. Let μ denote a vector of Lagrange multipliers on the restriction $-H' = G'\kappa$. Then the Lagrangian can be written as

$$\kappa' S \kappa / 2 + \lambda \sum_{i \in I} (\kappa_{+,i} + \kappa_{-,i}) + \mu' (H' + G' \kappa) - \sum_{i \in I} (\lambda_{+,i} \kappa_{+,i} + \lambda_{-,i} \kappa_{-,i}).$$

The first-order conditions are given by

$$e'_i S \kappa + e'_i G \mu = 0 \quad i \in I^C, \quad (25)$$

$$e'_i S \kappa + e'_i G \mu + \lambda = \lambda_{+,i} \quad i \in I, \quad (26)$$

$$-(e'_i S \kappa + e'_i G \mu) + \lambda = \lambda_{-,i} \quad i \in I. \quad (27)$$

The complementary slackness conditions are given by $\lambda_{+,i} \kappa_{+,i} = 0$ and $\lambda_{-,i} \kappa_{-,i} = 0$ for $i \in I$, and the feasibility constraints are $\lambda_{+,i} \geq 0$, $\lambda_{-,i} \geq 0$ for $i \in I$ and $-H' = G' \kappa$.

Let $\mathcal{A}^C = \{i : i \in I, \kappa_i = 0\}$, and let $\mathcal{A} = \{i : i \notin \mathcal{A}^C\}$ denote the set of active constraints. Let s denote a vector of length $|\mathcal{A}|$ with elements $s_i = \text{sign}(\kappa_i)$ if $i \in I$ and $s_i = 0$ otherwise.

The slackness and feasibility conditions imply that if for $i \in I$, $\kappa_i > 0$, then $\lambda_{+,i} = 0$, and if $\kappa_i < 0$ or $\lambda_{-,i} = 0$. It therefore follows from (26) and (27) that $e'_i S \kappa + e'_i G \mu = -\text{sign}(\kappa_i) \lambda = -s_i \lambda$. We can combine this condition with (25) and write

$$e'_i S \kappa + e'_i G \mu = -s_i \lambda, \quad i \in \mathcal{A}. \quad (28)$$

On the other hand, if $i \in \mathcal{A}^C$, then since $\lambda_{+,i}$ and $\lambda_{-,i}$ are non-negative, it follows from (26) and (27) that

$$|e'_i S \kappa + e'_i G \mu| \leq \lambda = |e'_j S \kappa + e'_j G \mu|, \quad i \in \mathcal{A}^C, j \in \mathcal{A}. \quad (29)$$

Let $\kappa_{\mathcal{A}}$ denote the subset of κ corresponding to the active moments, $G_{\mathcal{A}}$ denote the corresponding rows of G , and $S_{\mathcal{A}\mathcal{A}}$ the corresponding submatrix of S . Then we can write the condition (28) together with the feasibility constraint $G' \kappa = -H'$ compactly as

$$\begin{pmatrix} 0 & G'_{\mathcal{A}} \\ G_{\mathcal{A}} & S_{\mathcal{A}\mathcal{A}} \end{pmatrix} \begin{pmatrix} \mu \\ \kappa_{\mathcal{A}} \end{pmatrix} = \begin{pmatrix} -H' \\ -s\lambda \end{pmatrix}.$$

Using the block matrix inverse formula, this implies

$$\begin{aligned}\mu &= (G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}G_{\mathcal{A}})^{-1} (H' - G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}s\lambda) \\ \kappa_{\mathcal{A}} &= -S_{\mathcal{A}\mathcal{A}}^{-1}G_{\mathcal{A}}\mu - S_{\mathcal{A}\mathcal{A}}^{-1}s\lambda \\ &= S_{\mathcal{A}\mathcal{A}}^{-1}G_{\mathcal{A}}(G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}G_{\mathcal{A}})^{-1} (G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}s\lambda - H') - S_{\mathcal{A}\mathcal{A}}^{-1}s\lambda\end{aligned}$$

Consequently, if we're in a region in where the solution path is differentiable with respect to λ , we have

$$\frac{\partial \kappa_{\mathcal{A}}}{\partial \lambda} = S_{\mathcal{A}\mathcal{A}}^{-1}G_{\mathcal{A}}(G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}G_{\mathcal{A}})^{-1}G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}s - S_{\mathcal{A}\mathcal{A}}^{-1}s. \quad (30)$$

The differentiability of path is violated if either (a) the constraint (29) is violated for some $i \in \mathcal{A}^C$ if $\kappa(\lambda)$ keeps moving in the same direction, and we add i to \mathcal{A} at a point at which (29) holds with equality; or else (b) the sensitivity $\kappa_i(\lambda)$ for some $i \in \mathcal{A}$ reaches zero. In this case, drop i from \mathcal{A} . In either case, we need to re-calculate the direction (30) using the new definition of \mathcal{A} .

Based on the arguments above and the fact that $\kappa(0) = -S^{-1}G(G'S^{-1}G)^{-1}H'$, we can derive the following algorithm, similar to the LAR-LASSO algorithm, to generate the path of optimal sensitivities $\kappa(\lambda)$:

1. Initialize $\lambda = 0$, $\mathcal{A} = \{1, \dots, d_g\}$, $\mu = (G'S^{-1}G)^{-1}H'$, $\kappa = -S^{-1}G\mu$. Let s be a vector of length d_g with elements $s_i = \mathbb{I}\{i \in I\} \text{sign}(\kappa_i)$, and calculate initial directions as $\mu_{\Delta} = -(G'S^{-1}G)^{-1}G'S^{-1}s$, $\kappa_{\Delta} = -S^{-1}(G\mu_{\Delta} + s)$
2. While ($|\mathcal{A}| > \max\{d_g - d_{\gamma}, d_{\theta}\}$):
 - (a) Set step size to $d = \min\{d_1, d_2\}$, where

$$\begin{aligned}d_1 &= \min\{d > 0: \kappa_i + d\kappa_{\Delta,i} = 0, i \in \mathcal{A} \cap \mathcal{I}\} \\ d_2 &= \min\{d > 0: |e'_i(S\kappa + G\mu) + de'_i(S\kappa_{\Delta} + G\mu_{\Delta})| = \lambda + d, i \in \mathcal{A}^C\}\end{aligned}$$

Take step of size d : $\kappa \mapsto \kappa + d\kappa_{\Delta}$, $\mu \mapsto \mu + d\mu_{\Delta}$, and $\lambda \mapsto \lambda + d$.

- (b) If $d = d_1$, drop $\text{argmin}(d_1)$ from \mathcal{A} , and if $d = d_2$, then add $\text{argmin}(d_2)$ to \mathcal{A} . Let s be a vector of length d_g with elements $s_i = -\mathbb{I}\{i \in I\} \text{sign}(e'_i S\kappa + e'_i G\mu)$, and calculate new directions as

$$\begin{aligned}\mu_{\Delta} &= -(G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}G_{\mathcal{A}})^{-1}G'_{\mathcal{A}}S_{\mathcal{A}\mathcal{A}}^{-1}s_{\mathcal{A}} \\ (\kappa_{\Delta})_{\mathcal{A}} &= -S_{\mathcal{A}\mathcal{A}}^{-1}(G_{\mathcal{A}}\mu_{\Delta} + s_{\mathcal{A}}) \\ (\kappa_{\Delta})_{\mathcal{A}^C} &= 0\end{aligned}$$

The solution path $k(\lambda)$ is then obtained as $k(\lambda) = T'\kappa(\lambda)$.

Finally, we show that in the limit $M \rightarrow \infty$, the optimal sensitivity corresponds to a method of moments estimator based on the most informative set of d_θ moments, with the remaining $d_g - d_\theta$ moments dropped. The optimal sensitivity as $M \rightarrow \infty$ obtains by solving (23) as $\lambda \rightarrow \infty$. If B corresponds to columns of the identity matrix, then this is equivalent to minimizing $\|k_I\|_1$ subject to $H = -k'\Gamma$. This can be written as a linear program $\min k_{I,+} + k_{I,-}$ st $-H' = \Gamma'(k_+ - k_-)$, $k_+, k_- \geq 0$. The minimization problem is done on a d_θ -dimensional hyperplane, and solution must occur at a boundary point of the feasible set, where only d_θ variables are non-zero. So the optimal k has d_θ non-zero elements.

B Specification test

One can test the null hypothesis of correct specification (i.e. the null hypothesis that $c = 0$) using the J statistic

$$J = n \min_{\theta} \hat{g}(\theta)' \hat{\Sigma}^{-1} \hat{g}(\theta) = n \hat{g}(\hat{\theta})' \hat{\Sigma}^{-1} \hat{g}(\hat{\theta}),$$

where $\hat{\theta} = \operatorname{argmin}_{\theta} \hat{g}(\theta)' \hat{\Sigma}^{-1} \hat{g}(\theta)$. Alternatively, letting $\hat{\Sigma}^{-1/2}$ denote the symmetric square root of $\hat{\Sigma}^{-1}$, one can project $\hat{\Sigma}^{-1/2} \hat{g}(\tilde{\theta})$, where $\tilde{\theta}$ is some consistent estimate, onto the complement of the space spanned by $\hat{\Sigma}^{-1/2} \hat{\Gamma}$,

$$S = n \hat{g}(\tilde{\theta})' \hat{\Sigma}^{-1/2} \hat{R} \hat{\Sigma}^{-1/2} \hat{g}(\tilde{\theta}),$$

where $\hat{R} = I - \hat{\Sigma}^{-1/2} \hat{\Gamma} (\hat{\Gamma}' \hat{\Sigma}^{-1} \hat{\Gamma})^{-1} \hat{\Gamma}' \hat{\Sigma}^{-1/2}$. If the model is correctly specified, so that $c = 0$, S and J are asymptotically equivalent (Newey and McFadden, 1994, p. 2231), and distributed $\chi_{d_g - d_\theta}^2$.

Under local misspecification, the J statistic has a noncentral χ^2 distribution, with non-centrality parameter depending on c (Newey, 1985), and the asymptotic equivalence of J and S still holds. In this section, we use this observation to form a test of the null hypothesis $H_0: c \in \mathcal{C}$. When \mathcal{C} takes the form in Eq. (12) for some norm $\|\cdot\|$, inverting these tests gives a lower CI for M . We begin with a lemma deriving the asymptotic distribution of S and J under local misspecification.

Lemma B.1. *Suppose that Eqs. (1), (2) and (3) hold, and that $\hat{\theta}$ and $\tilde{\theta}$ satisfy, for some K and $K'_{opt} = -(\Gamma'\Sigma^{-1}\Gamma)^{-1}\Gamma'\Sigma^{-1}$,*

$$\sqrt{n}(\hat{\theta} - \theta_0) = K'_{opt} \sqrt{n} \hat{g}(\theta_0), \quad \text{and} \quad \sqrt{n}(\tilde{\theta} - \theta_0) = K' \sqrt{n} \hat{g}(\theta_0).$$

Suppose that $\hat{\Sigma}$ and $\hat{\Gamma}$ are consistent estimates of Σ and Γ , and that Σ and Γ are full rank.

Then $S = J + o_P(1)$ and S and J converge in distribution to a noncentral chi-square distribution with $d_g - d_\theta$ degrees of freedom and noncentrality parameter $c'\Sigma^{-1/2}R\Sigma^{-1/2}c$ where $R = I - \Sigma^{-1/2}\Gamma(\Gamma'\Sigma^{-1}\Gamma)^{-1}\Gamma\Sigma^{-1/2}$.

Proof. By Eqs. (1), (2) and (3), $\sqrt{n}\hat{g}(\tilde{\theta}) = (I + \Gamma K')\Sigma^{1/2}(\Sigma^{-1/2}c + Z_n) + o_P(1)$ where $Z_n = \Sigma^{-1/2}[\sqrt{n}\hat{g}(\theta_0) - c] \xrightarrow{d} \mathcal{N}(0, I_{d_g})$, so that

$$\begin{aligned} S &= (\Sigma^{-1/2}c + Z_n)'\Sigma^{1/2}(\Sigma^{-1/2} + \Sigma^{-1/2}\Gamma K')'R(\Sigma^{-1/2} + \Sigma^{-1/2}\Gamma K')\Sigma^{1/2}(\Sigma^{-1/2}c + Z_n) + o_P(1) \\ &= (\Sigma^{-1/2}c + Z_n)'R(\Sigma^{-1/2}c + Z_n) + o_P(1) \xrightarrow{d} (\Sigma^{-1/2}c + Z)'R(\Sigma^{-1/2}c + Z) \end{aligned}$$

where $Z \sim \mathcal{N}(0, I_{d_g})$ and we use the fact that $R(I + \Sigma^{-1/2}\Gamma K'\Sigma^{1/2}) = R$. Similarly,

$$\sqrt{n}\hat{g}(\hat{\theta}) = (I - \Gamma(\Gamma'\Sigma^{-1}\Gamma)\Gamma'\Sigma^{-1})(c + \Sigma^{1/2}Z_n) + o_P(1) = \Sigma^{1/2}R(\Sigma^{-1/2}c + Z_n) + o_P(1),$$

so that $J = (\Sigma^{-1/2}c + Z_n)'R(\Sigma^{-1/2}c + Z_n) + o_P(1) = S + o_P(1)$. To prove the second claim, decompose $R = P_1P_1'$, where $P_1 \in \mathbb{R}^{d_\theta \times (d_g - d_\theta)}$ corresponds to the eigenvectors associated with non-zero eigenvalues of R . Then

$$(\Sigma^{-1/2}c + Z)'R(\Sigma^{-1/2}c + Z) = (P_1'\Sigma^{-1/2}c + P_1'Z)'(P_1'\Sigma^{-1/2}c + P_1'Z).$$

Since $P_1'Z \sim \mathcal{N}(0, I_{d_g - d_\theta})$, it follows that the random variable in the preceding display has a non-central χ^2 distribution with $d_g - d_\theta$ degrees of freedom and non-centrality parameter $c'\Sigma^{-1/2}R\Sigma^{-1/2}c$. \square

Lemma B.1 can be interpreted in using the limiting experiment described in Section 4.1. In particular, the asymptotic distribution of the S and J statistics is given by the distribution of the statistic $Y'\Sigma^{-1/2}R\Sigma^{-1/2}Y$ in the limiting experiment.

The quantiles of a non-central chi-square distribution are increasing in the noncentrality parameter (Sun et al., 2010). Thus, to test the null hypothesis $H_0: c \in \mathcal{C}$, the appropriate critical value for tests based on the J or S statistic is based on a non-central chi-squared distribution, with non-centrality parameter

$$\bar{\lambda} = \sup_{c \in \mathcal{C}} c'\Sigma^{-1/2}R\Sigma^{-1/2}c.$$

If $\mathcal{C} = \{B\gamma: \|\gamma\|_p \leq M\}$, then this becomes

$$\bar{\lambda} = \sup_{\|t\|_p \leq M} t'B'\Sigma^{-1/2}R\Sigma^{-1/2}Bt = \sup_{\|t\|_p \leq 1} M^2\|R\Sigma^{-1/2}Bt\|_2^2 = M^2\|A\|_{p,2}^2.$$

where the second equality uses the fact that R is idempotent, $A = R\Sigma^{-1/2}B$, and $\|A\|_{p,q} = \max_{\|x\|_p \leq 1} \|Ax\|_q$ is the (p, q) operator norm. For $p = 2$, the operator norm has a closed form, which gives $\bar{\lambda} = M \max \text{eig}(B'\Sigma^{-1/2}R\Sigma^{-1/2}B)$.

C Asymptotic coverage and efficiency

This appendix contains the asymptotic coverage and efficiency results discussed in Section 4. In particular, we prove Theorem 4.1. In order to allow for stronger statements, we state upper and lower bounds separately. Theorem 4.1 then follows by combining these results. Theorem 4.1 focuses on two-sided CIs in the case where \mathcal{C} is centrosymmetric, in addition to being convex. In this appendix, we also prove analogous results for one-sided CIs, and we generalize these results to the case where \mathcal{C} is a convex but asymmetric set. When \mathcal{C} is convex but asymmetric, the negative results about the scope for improvement when c is close to zero no longer hold. Therefore, we consider the general problem of optimizing quantiles of excess length over a set $\mathcal{D} \subseteq \mathcal{C}$, which may be a strict subset of \mathcal{C} .

The remainder of this appendix is organized as follows. Appendix C.1 presents notation and definitions, as well as an overview of the results. Appendix C.2 contains results on least favorable submodels as well as a two-point testing lemma used in later proofs. We then use this to obtain efficiency bounds for one-sided CIs in Appendix C.3, and for two-sided CIs in Appendix C.4. Appendix C.5 shows that our CIs achieve (or, for two-sided CIs, nearly achieve) these bounds. Appendix C.6 shows how Theorem 4.1 follows from these results, and also gives a one-sided version of this theorem. Primitive conditions for the misspecified linear IV model, as well as a general construction of a least favorable submodel satisfying the assumptions used in this section, are given in the supplemental appendix.

C.1 Setup

While our focus is on parameter spaces that place restrictions on c , we will also allow for local restrictions on θ in some results. This allows us to bound the scope for “directing power” at particular values of θ . Formally, for some parameter θ^* , we consider the local parameter space that restricts $(\sqrt{n}(\theta - \theta^*)', c)'$ to some set $\mathcal{F} \subseteq \mathbb{R}^{d_\theta + d_g}$. The unrestricted case considered throughout most of the main text corresponds to $\mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$ (in which case θ^* does not affect the definition of the parameter space). We also allow for additional restrictions on θ by placing it in some set Θ_n . Finally, we use \mathcal{P} to denote the set of distributions P over which we require coverage.

With this notation, the set of values of θ that are consistent with the model under P (i.e.

the identified set under P) is

$$\Theta_I(P) = \Theta_I(P; \mathcal{F}, \Theta_n) = \{\theta \in \Theta_n : \sqrt{n}((\theta - \theta^*)', g_P(\theta)')' \in \mathcal{F}\},$$

and the set of pairs (θ, P) over which coverage is required is given by

$$\mathcal{S}_n = \{(\theta, P) \in \Theta_n \times \mathcal{P} : \theta \in \Theta_I(P)\} = \{(\theta, P) \in \Theta_n \times \mathcal{P} : \sqrt{n}((\theta - \theta^*)', g_P(\theta)')' \in \mathcal{F}\},$$

which reduces to the definition in Theorem 4.1 when $\mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$. The coverage requirement for a CI \mathcal{I}_n is then given by (20) with this definition of \mathcal{S}_n . To compare one-sided CIs $[\hat{c}, \infty)$, we will consider the β quantile of excess length. Rather than restricting ourselves to the minimax criterion, we consider worst-case excess length over a potentially smaller parameter space \mathcal{G} , which may place additional restrictions on θ and c . Let

$$q_{\beta,n}(\hat{c}; \mathcal{P}, \mathcal{G}, \Theta_n) = \sup_{P \in \mathcal{P}} \sup_{\theta \in \Theta_I(P; \mathcal{G}, \Theta_n)} q_{P,\beta}(h(\theta) - \hat{c})$$

where $q_{P,\beta}$ denotes the β quantile under P . We will also consider bounds on $q_{P,\beta}(h(\theta) - \hat{c})$ at a single P , which corresponds to the optimistic case of optimizing length at a single distribution. For two-sided CIs, we will consider expected length.

Our efficiency bounds can be thought of as applying the bounds in [Armstrong and Kolesár \(2018\)](#) to a local asymptotic setting, which corresponds to the limiting model (13) with $\Gamma = \Gamma_{\theta^*, P_0}$, $\Sigma = \Sigma_{\theta^*, P_0}$ and $H = H_{\theta^*}$. The between class modulus of continuity for this model is

$$\begin{aligned} \omega(\delta; \mathcal{F}, \mathcal{G}, H, \Gamma, \Sigma) = \sup H(s_1 - s_0) \quad \text{s.t.} \quad & (s'_0, c'_0)' \in \mathcal{F}, (s'_1, c'_1)' \in \mathcal{G}, \\ & [(c_1 - c_0) - \Gamma(s_1 - s_0)]' \Sigma^{-1} [(c_1 - c_0) - \Gamma(s_1 - s_0)] \leq \delta^2. \end{aligned} \quad (31)$$

We use the notation $\omega(\delta)$ and $\omega(\delta; \mathcal{F}, \mathcal{G})$ when the context is clear. In the case where $\mathcal{G} = \mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$ and \mathcal{C} is centrosymmetric, the solution satisfies $s_1 = -s_0$ and $c_1 = -c_0$, which gives the same optimization problem as (15), with the objective multiplied by two (this matches the definition of $\omega(\cdot)$ used to define κ_* in the main text).

For one-sided CIs, we show that, for any CI satisfying the coverage condition (20) for a rich enough class \mathcal{P} , we will have

$$\liminf_{n \rightarrow \infty} \sqrt{n} q_{\beta,n}(\hat{c}; \mathcal{P}, \mathcal{G}, \Theta_n) \geq \omega(\delta_\beta; \mathcal{F}, \mathcal{G}, H, \Gamma, \Sigma) \quad (32)$$

where $\delta_\beta = z_{1-\alpha} + z_\beta$, where z_τ denotes the τ quantile of the $\mathcal{N}(0, 1)$ distribution. For bounds

on excess length at a single P_0 with $E_{P_0}g(w_i, \theta^*) = 0$, we obtain this bound with $\mathcal{G} = \{0\}$:

$$\liminf_{n \rightarrow \infty} \sqrt{n} q_{P_0, \beta}(h(\theta^*) - \hat{c}) \geq \omega(\delta_\beta; \mathcal{F}, \{0\}, H, \Gamma, \Sigma). \quad (33)$$

These results can be thought of as a local asymptotic version of Theorem 3.1 in [Armstrong and Kolesár \(2018\)](#) applied to our setting.

For two-sided CIs, we show that, if a CI $\mathcal{I}_n = \{\hat{h} \pm \hat{\chi}\}$ satisfies the coverage condition (20) for a rich enough class \mathcal{P} , then, for any P_0 with $E_{P_0}g(w_i, \theta^*) = 0$, expected length satisfies

$$\begin{aligned} & \liminf_{T \rightarrow \infty} \liminf_{n \rightarrow \infty} E_{P_0} \min\{\sqrt{n} 2\hat{\chi}, T\} \\ & \geq (1 - \alpha) E[\omega(z_{1-\alpha} - Z; \{0\}, \mathcal{F}, H, \Gamma, \Sigma) + \omega(z_{1-\alpha} - Z; \mathcal{F}, \{0\}, H, \Gamma, \Sigma) | Z \leq z_{1-\alpha}], \end{aligned} \quad (34)$$

where $Z \sim \mathcal{N}(0, 1)$. The above bound uses truncated expected length to avoid technical issues with convergence of moments when achieving the bound (note however that this bound immediately implies the same bound on excess length without truncation). Our results constrain the CI to take the form of an interval. We conjecture that the bound applies to arbitrary confidence sets (with length defined as Lebesgue measure) under additional regularity conditions.

Here, “rich enough” means that \mathcal{P} contains a least favorable submodel. Appendix C.2 begins the derivation of our efficiency results by giving conditions on this submodel. In Supplemental Appendix E.1, we construct a submodel satisfying these conditions under mild conditions.

C.2 Least favorable submodel

Let P_0 be a distribution with $E_{P_0}g(w_i, \theta^*) = 0$ (i.e. the model holds for this data-generating process with $\theta = \theta^*$ and $c = 0$), and consider a parametric submodel P_t indexed by $t \in \mathbb{R}^{d_g}$ (i.e. the dimension of t is the same as the dimension of the values of $g(w_i, \theta)$) with P_t equal to P_0 at $t = 0$. We assume that $\{w_i\}_{i=1}^n$ are i.i.d. under P_t . Let $\pi_t(w_i)$ denote the density of a single observation with respect to its distribution under P_0 , so that $E_{P_t}f(w_i) = E_{P_0}f(w_i)\pi_t(w_i)$ for any function f . We expect that the least favorable submodel for this problem will be the one that makes estimating $E_{P_0}g(W_i, \theta^*)$ most difficult. This corresponds to any subfamily with score function $g(w_i, \theta^*)$. We also place additional conditions on this submodel, given in the following assumption.

Assumption C.1. *The data are i.i.d. under P_t for all t in a neighborhood of zero, and the density $\pi_t(w_i)$ for a single observation is quadratic mean differentiable at $t = 0$ with score*

function $g(w_i, \theta^*)$, where $E_{P_0}g(w_i, \theta^*) = 0$. In addition, the function $(t', \theta')' \mapsto E_{p_t}g(w_i, \theta)$ is continuously differentiable at $(0', \theta^*)'$ with

$$\left[\frac{d}{d(t', \theta')} E_{p_t}g(w_i, \theta) \right]_{t=0, \theta=\theta^*} = (\Sigma, \Gamma) \quad (35)$$

where Σ and Γ are full rank.

To understand Assumption C.1, note that Problem 12.17 in [Lehmann and Romano \(2005\)](#) gives the Jacobian with respect to t as Σ in the case where $g(w_i, \theta^*)$ is bounded, and the Jacobian with respect to θ is equal to Γ by definition. Assumption C.1 requires the slightly stronger condition that $E_{p_t}g(w_i, \theta)$ is continuously differentiable with respect to $(t', \theta')'$ for t close to 0 and θ close to θ^* . This is needed to apply the Implicit Function Theorem in the derivations that follow. In the supplemental materials, we give a construction of a quadratic mean differentiable family satisfying this condition, without requiring boundedness of $g(w_i, \theta^*)$ (Lemma E.1 in Supplemental Appendix E.1).

The bounds in [Armstrong and Kolesár \(2018\)](#) are obtained by bounding the power of a two-point test (simple null and simple alternative) where the null and alternative are given by the points that achieve the modulus. To obtain analogous results in our setting, we use a bound on the power of a two-point test in a least favorable submodel.

Consider sequences of local parameter values $(\theta'_{0,n}, c'_{0,n})'$ and $(\theta'_{1,n}, c'_{1,n})'$ where, for some s_0, c_0, s_1 and c_1 ,

$$\theta_{d,n} = \theta^* + (s_d + o(1))/\sqrt{n}, \quad c_{d,n} = c_d + o(1) \quad d \in \{0, 1\} \quad (36)$$

Consider a sequence of tests of $(\theta'_{0,n}, c'_{0,n})'$ vs $(\theta'_{1,n}, c'_{1,n})'$. Formally, for any $(\theta', c)'$, let

$$\mathcal{P}_n(\theta, c) = \{P \in \mathcal{P} : E_P g(w_i, \theta) = c/\sqrt{n}\} \quad (37)$$

be the set of probability distributions in \mathcal{P} that are consistent with the parameter values $(\theta', c)'$. We derive a bound on the asymptotic minimax power of a level α test of

$$H_{0,n} : P \in \mathcal{P}_n(\theta_{0,n}, c_{0,n}) \quad \text{vs} \quad H_{1,n} : P \in \mathcal{P}_n(\theta_{1,n}, c_{1,n}), \quad (38)$$

as well as a bound on the power of a test of $H_{0,n}$ at P_0 . Let Φ be the standard normal cdf and let

$$\bar{\beta}(s_0, c_0, s_1, c_1) = \Phi \left(\sqrt{[c_1 - c_0 - \Gamma(s_1 - s_0)]' \Sigma^{-1} [c_1 - c_0 - \Gamma(s_1 - s_0)]} - z_{1-\alpha} \right).$$

Lemma C.1. *Let \mathcal{P} be a class of distributions that contains a family P_t that satisfies Assumption C.1. Then, for any sequence of tests ϕ_n satisfying $\limsup_n \sup_{P \in \mathcal{P}_n(\theta_{0,n}, c_{0,n})} E_P \phi_n \leq \alpha$, we have*

$$\limsup_n E_{P_0} \phi_n \leq \bar{\beta}(s_0, c_0, 0, 0) \quad \text{and} \quad \limsup_n \inf_{P \in \mathcal{P}_n(\theta_{1,n}, c_{1,n})} E_P \phi_n \leq \bar{\beta}(s_0, c_0, s_1, c_1).$$

Lemma C.1 says that the asymptotic minimax power of any test of $H_{0,n}$ vs $H_{1,n}$ is bounded by $\bar{\beta}(s_0, c_0, s_1, c_1)$. Furthermore, if we take $s_1 = 0$ and $c_1 = 0$, then this bound is achieved at P_0 . Note that, in keeping with the analogy with the linear model (13), $\bar{\beta}(s_0, c_0, s_1, c_1)$ is the power of the optimal (Neyman-Pearson) test of the simple null (s'_0, c'_0) vs the simple alternative (s'_1, c'_1) in the model (13).

Proof of Lemma C.1. The proof involves two steps. First, we use the Implicit Function Theorem to find sequences $t_{0,n}$ and $t_{1,n}$ such that $P_{t_{0,n}}$ satisfies $H_{0,n}$ and $P_{t_{1,n}}$ satisfies $H_{1,n}$. Next, we apply a standard result on testing in quadratic mean differentiable families to obtain the limiting power of the optimal test of $P_{t_{0,n}}$ vs $P_{t_{1,n}}$, which gives an upper bound on the limiting minimax power of any test of $H_{0,n}$ vs $H_{1,n}$.

Let $f(t, \theta, a) = E_{P_t} g(w_i, \theta) - a$ so that $(\theta', c)'$ is consistent with P_t iff. $f(t, \theta, c/\sqrt{n}) = 0$. Under Assumption C.1, it follows from the Implicit Function Theorem that there exists a function $r(\theta, a)$ such that, for θ in a neighborhood of θ^* and a in a neighborhood of zero,

$$E_{P_{r(\theta, a)}} g(w_i, \theta) - a = f(r(\theta, a), \theta, a) = 0.$$

Thus, letting $t_{0,n} = r(\theta_{0,n}, c_{0,n}/\sqrt{n})$ and $t_{1,n} = r(\theta_{1,n}, c_{1,n}/\sqrt{n})$, $P_{t_{0,n}}$ satisfies $H_{0,n}$ and $P_{t_{1,n}}$ satisfies $H_{1,n}$. Furthermore,

$$\left[\frac{d}{d(\theta', a')} r(\theta, a) \right]_{(\theta', a') = (\theta^*, 0)} = -\Sigma^{-1}(\Gamma, -I_{d_g})$$

so that

$$r(\theta, a) = \Sigma^{-1}a - \Sigma^{-1}\Gamma(\theta - \theta^*) + o(\|\theta - \theta^*\| + \|a\|).$$

Thus, letting $t_{0,\infty} = \Sigma^{-1}c_0 - \Sigma^{-1}\Gamma s_0$, we have

$$\begin{aligned} t_{0,n} &= r(\theta_{0,n}, c_{0,n}/\sqrt{n}) = \Sigma^{-1}c_{0,n}/\sqrt{n} - \Sigma^{-1}\Gamma(\theta_{0,n} - \theta^*) + o(\|\theta_{0,n} - \theta^*\| + \|c_{0,n}\|/\sqrt{n}) \\ &= \Sigma^{-1}c_0/\sqrt{n} - \Sigma^{-1}\Gamma s_0/\sqrt{n} + o(1/\sqrt{n}) = t_{0,\infty}/\sqrt{n} + o(1/\sqrt{n}). \end{aligned}$$

Similarly, $t_{1,n} = t_{1,\infty}/\sqrt{n} + o(1/\sqrt{n})$ where $t_{1,\infty} = \Sigma^{-1}c_1 - \Sigma^{-1}\Gamma s_1$.

Since the information matrix for this submodel evaluated at $t = 0$ is Σ , it follows from

the arguments in Example 12.3.12 in [Lehmann and Romano \(2005\)](#), extended to the case where the null and alternative are both drifting sequences (rather than just the alternative), that the limit of the power of the Neyman-Pearson test of $P_{t_{0,n}}$ vs $P_{t_{1,n}}$ is

$$\Phi \left(\sqrt{[t_{1,\infty} - t_{0,\infty}]' \Sigma [t_{1,\infty} - t_{0,\infty}] - z_{1-\alpha}} \right) = \bar{\beta}(s_0, c_0, s_1, c_1).$$

This gives the required bound on minimax power over $H_{1,n}$. To obtain the bound on power at P_0 , note that, for $\theta_{1,n} = \theta^*$ and $c_{1,n} = 0$, $t_{0,n} = 0$, the bound also corresponds to the power of a test that is optimal for $P_{t_{0,n}}$ vs P_0 . \square

C.3 One-sided CIs

We prove the following efficiency bound for one-sided CIs.

Theorem C.1. *Let \mathcal{P} be a class of distributions that contains a submodel P_t satisfying Assumption C.1. Let $\Theta_n(C) = \{\theta \mid \|\theta - \theta^*\| \leq C/\sqrt{n}\}$ for some constant C , and let \mathcal{F} be given. Let $[\hat{c}, \infty)$ be a sequence of CIs such that, for all C , the coverage condition (20) holds with $\Theta_n = \Theta_n(C)$. Let $\mathcal{G} \subseteq \mathcal{F}$ be a set such that the limiting modulus ω is well-defined and continuous for all δ . Then the asymptotic lower bounds (32) and (33) hold.*

Proof. Consider a sequence of simple null and alternative values of θ and c that satisfy (36) for some s_0, c_0, s_1, c_1 , with $(\sqrt{n}(\theta_{0,n} - \theta^*)', c'_{0,n})' \in \mathcal{F}$ and $(\sqrt{n}(\theta_{1,n} - \theta^*)', c'_{1,n})' \in \mathcal{G}$, for each n . Note that

$$\lim_{n \rightarrow \infty} \sqrt{n}[h(\theta_{1,n}) - h(\theta_{0,n})] = H(s_1 - s_0).$$

Consider the testing problem $H_{0,n} : P \in \mathcal{P}_n(\theta_{0,n}, c_{0,n})$ vs $H_{1,n} : P \in \mathcal{P}_n(\theta_{1,n}, c_{1,n})$ defined in (37) and (38). Suppose that

$$q_{\beta,n}(\hat{c}; \mathcal{P}, \mathcal{G}, \Theta_n) < h(\theta_{1,n}) - h(\theta_{0,n}). \quad (39)$$

Let ϕ_n denote the test that rejects when $h(\theta_{0,n}) \notin [\hat{c}, \infty)$. Since, for any $P \in \mathcal{P}_n(\theta_{1,n}, c_{1,n})$, we have $q_{P,\beta}(h(\theta_{1,n}) - \hat{c}) \leq q_{\beta,n}(\hat{c}; \mathcal{P}, \mathcal{G}, \Theta_n)$ by construction, it follows that, for all $P \in \mathcal{P}_n(\theta_{1,n}, c_{1,n})$,

$$E_P \phi_n = P(h(\theta_{1,n}) - \hat{c} < h(\theta_{1,n}) - h(\theta_{0,n})) \geq P(h(\theta_{1,n}) - \hat{c} \leq q_{P,\beta}(h(\theta_{1,n}) - \hat{c})) \geq \beta,$$

where the last step follows from properties of quantiles (Lemma 21.1 in [van der Vaart, 1998](#)). The coverage requirement (20) implies that the test ϕ_n that rejects when $h(\theta_{0,n}) \notin [\hat{c}, \infty)$ has

asymptotic level α for $H_{0,n}$. Thus, by Lemma C.1, we must have $\beta \leq \bar{\beta}(s_0, c_0, s_1, c_1)$ if (39) holds infinitely often.

It follows that, if $\bar{\beta}(s_0, c_0, s_1, c_1) < \beta$, we must have

$$\liminf_{n \rightarrow \infty} \sqrt{n} q_{\beta,n}(\hat{c}; \mathcal{P}, \mathcal{G}, \Theta_n) \geq H(s_1 - s_0)$$

since otherwise, Eq. (39) would hold infinitely often. Since the sequences and limiting $(s'_0, c'_0) \in \mathcal{F}$ and $(s'_1, c'_1) \in \mathcal{G}$ were arbitrary, the above bound holds for any $(s'_0, c'_0) \in \mathcal{F}$ and $(s'_1, c'_1) \in \mathcal{G}$ with $\bar{\beta}(s_0, c_0, s_1, c_1) \leq \beta - \eta$, where $\eta > 0$ is arbitrary. The maximum of the right-hand side over s_0, c_0, s_1, c_1 in this set is equal to $\omega(\delta_{\beta-\eta}; \mathcal{F}, \mathcal{G}, H, \Gamma, \Sigma)$ by definition, so taking $\eta \rightarrow 0$ gives the result. \square

C.4 Two-sided CIs

We prove the following efficiency bound for two-sided CIs.

Theorem C.2. *Suppose that, for all C , $\{\hat{h} \pm \hat{\chi}\}$ satisfies the local coverage condition (20) with $\Theta_n = \Theta_n(C) = \{\theta \mid \|\theta - \theta^*\| \leq C/\sqrt{n}\}$, where \mathcal{P} contains a submodel P_t satisfying Assumption C.1. Suppose also that $0_{d_\theta+d_g} \in \mathcal{F}$ and a minimizer $(s'_\vartheta, c'_\vartheta)'$ of $(c - \Gamma s)' \Sigma^{-1} (c - \Gamma s)$ subject to $Hs = \vartheta$ and $(s', c')' \in \mathcal{F}$ exists for all $\vartheta \in \mathbb{R}$. Then the asymptotic lower bound (34) holds.*

In the case where $\mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$, which is the focus of the main text, a sufficient condition for the existence of the minimizer $(s'_\vartheta, c'_\vartheta)'$ is that \mathcal{C} is compact, H is not equal to the zero vector and Γ is full rank.

Proof. For each $\vartheta \in \mathbb{R}$, let $\tilde{\theta}_{\vartheta,n} = \theta^* + s_\vartheta/\sqrt{n}$, and let $\phi_{\vartheta,n} = I(h(\tilde{\theta}_{\vartheta,n}) \notin \{\hat{h} \pm \hat{\chi}\})$ be the test that rejects when $h(\tilde{\theta}_{\vartheta,n})$ is not in the CI. When the constant C defining $\Theta_n = \Theta_n(C)$ is large enough, the asymptotic coverage condition (20) implies that $\phi_{\vartheta,n}$ is an asymptotic level α test for $H_{0,n} : P \in \mathcal{P}_n(\tilde{\theta}_{\vartheta,n}, c_\vartheta)$ defined in (37) and (38). Thus, by Lemma C.1,

$$\limsup_{n \rightarrow \infty} E_{P_0} \phi_{\vartheta,n} \leq \Phi(\delta_\vartheta - z_{1-\alpha}) \quad \text{where} \quad \delta_\vartheta = \sqrt{(c_\vartheta - \Gamma s_\vartheta)' \Sigma^{-1} (c_\vartheta - \Gamma s_\vartheta)}. \quad (40)$$

We apply this bound to a grid of values of ϑ . Let $\mathcal{E}_n(m)$ denote the grid centered at zero with length $2m$ and meshwidth $1/m$:

$$\mathcal{E}_n(m) = \{j/m : j \in \mathbb{Z}, |j| \leq m^2\}.$$

Let

$$\tilde{\mathcal{E}}_n(m) = \{\sqrt{n}[h(\tilde{\theta}_{\vartheta,n}) - h(\theta^*)] : \vartheta \in \mathcal{E}_n(m)\}.$$

Note that $h(\tilde{\theta}_{\vartheta,n}) = h(\theta^*) + (1 + o(1))Hs_{\vartheta}/\sqrt{n} = h(\theta^*) + (1 + o(1))\vartheta/\sqrt{n}$. Thus, letting a_1, \dots, a_{2m^2+1} denote the ordered elements in $\mathcal{E}_n(m)$ and $\tilde{a}_1, \dots, \tilde{a}_{2m^2+1}$ the ordered elements in $\tilde{\mathcal{E}}_n$, we have $\tilde{a}_j \rightarrow a_j$ for each j as $n \rightarrow \infty$.

Let $\mathcal{N}(n, m)$ be the number of elements \tilde{a}_j in $\tilde{\mathcal{E}}_n$ such that $h(\theta^*) + \tilde{a}_j/\sqrt{n} = h(\tilde{\theta}_{a_j,n}) \in \{\hat{h} \pm \hat{\chi}\}$. Then

$$E_{P_0}\mathcal{N}(n, m) = \sum_{j=1}^{2m^2+1} E_{P_0}I(h(\tilde{\theta}_{a_j,n}) \in \{\hat{h} \pm \hat{\chi}\}) = \sum_{j=1}^{2m^2+1} [1 - E_{P_0}\phi_{a_j,n}].$$

It follows from (40) that (assuming the constant C that defines $\Theta_n(C)$ is large enough),

$$\liminf_{n \rightarrow \infty} E_{P_0}\mathcal{N}(n, m) \geq \sum_{j=1}^{2m^2+1} [1 - \Phi(\delta_{a_j} - z_{1-\alpha})] = \sum_{j=1}^{2m^2+1} \Phi(z_{1-\alpha} - \delta_{a_j}).$$

Note that $2\hat{\chi} \geq n^{-1/2}[\mathcal{N}(n, m) - 1] \cdot \min_{1 \leq j \leq 2m^2} (\tilde{a}_{j+1} - \tilde{a}_j) = n^{-1/2}[\mathcal{N}(n, m) - 1] \cdot m^{-1} \cdot (1 + \varepsilon_n)$ where $\varepsilon_n = \min_{1 \leq j \leq 2m^2} (\tilde{a}_{j+1} - \tilde{a}_j)/m^{-1} - 1$ is a nonrandom sequence converging to zero. This, combined with the above display, gives

$$\liminf_{n \rightarrow \infty} E_{P_0} \min\{2n^{1/2}\hat{\chi}, T\} \geq \left[m^{-1} \sum_{j=1}^{2m^2+1} \Phi(z_{1-\alpha} - \delta_{a_j}) - m^{-1} \right]$$

for any $T > 2m$. We have

$$m^{-1} \sum_{j=1}^{2m^2+1} \Phi(z_{1-\alpha} - \delta_{a_j}) = m^{-1} \sum_{j=1}^{2m^2+1} \int I(\delta_{a_j} \leq z_{1-\alpha} - z) d\Phi(z). \quad (41)$$

Following the proof of Theorem 3.2 in [Armstrong and Kolesár \(2018\)](#), note that, for $\vartheta \geq 0$, $t \geq 0$, we have $\delta_{\vartheta} \leq t$ iff. $\vartheta \leq \omega(t; \{0\}, \mathcal{F})$. Indeed, note that $\omega(\delta_{\vartheta}; \{0\}, \mathcal{F}) \geq Hs_{\vartheta} = \vartheta$ by feasibility of 0 and $s_{\vartheta}, c_{\vartheta}$ for this modulus problem. Since the modulus is increasing, this means that, if $\delta_{\vartheta} \leq t$, we must have $\vartheta \leq \omega(t; \{0\}, \mathcal{F})$. Now suppose $\vartheta \leq \omega(t; \{0\}, \mathcal{F})$. Then $Hs_{\omega(t; \{0\}, \mathcal{F})} \geq \vartheta$, so, for some $\lambda \in [0, 1]$, $(s'_{\lambda}, c'_{\lambda}) = \lambda(s'_{\omega(t; \{0\}, \mathcal{F})}, c'_{\omega(t; \{0\}, \mathcal{F})})$ satisfies $Hs_{\lambda} = \vartheta$, which means that $\delta_{\vartheta} \leq \sqrt{(c_{\lambda} - \Gamma s_{\lambda})' \Sigma^{-1} (c_{\lambda} - \Gamma s_{\lambda})} \leq t$ as claimed.

Thus, the part of the expression in (41) corresponding to terms in the sum with $a_j \geq 0$

is given by

$$\begin{aligned} m^{-1} \sum_{j=1}^{2m^2+1} \int I(0 \leq a_j \leq \omega(z_{1-\alpha} - z; \{0\}, \mathcal{F})) d\Phi(z) \\ \geq \int_{z \leq z_{1-\alpha}} \min\{\omega(z_{1-\alpha} - z; \{0\}, \mathcal{F}) - 1/m, m\} d\Phi(z). \end{aligned}$$

By the Dominated Convergence Theorem, this converges to $\int_{z \leq z_{1-\alpha}} \omega(z_{1-\alpha} - z; \{0\}, \mathcal{F}) d\Phi(z)$ as $m \rightarrow \infty$. Similarly, for $\vartheta < 0$, $t \geq 0$, we have $\delta_\vartheta \leq t$ iff. $-\vartheta \leq \omega(t; \mathcal{F}, \{0\})$, so that an analogous argument shows that, for arbitrary $\varepsilon > 0$, there exists m such that $\int_{z \leq z_{1-\alpha}} \omega(z_{1-\alpha} - z; \mathcal{F}, \{0\}) d\Phi(z) - \varepsilon$ is an asymptotic lower bound for the part of the expression (41) that corresponds to terms in the sum with $a_j < 0$. Thus, for any $\varepsilon > 0$, there exist constants C and T such that, if the coverage condition (20) holds with $\Theta_n = \Theta_n(C)$,

$$\liminf_{n \rightarrow \infty} E_{P_0} \min\{n^{1/2} 2\hat{\chi}, T\} \geq \int_{z \leq z_{1-\alpha}} [\omega(z_{1-\alpha} - z; \{0\}, \mathcal{F}) + \omega(z_{1-\alpha} - z; \mathcal{F}, \{0\})] d\Phi(z) - 2\varepsilon.$$

This gives the result. □

C.5 Achieving the bound

This section gives formal results showing that the CIs proposed in the main text are asymptotically valid, and that, if the sensitivities are chosen optimally, they achieve the efficiency bound in Theorem C.1 in the one-sided case, and nearly achieve the bound in Theorem C.2 in the two-sided case (where “nearly” means up to the sharp efficiency bound κ_* in the limiting model, given in (18), in the case where \mathcal{C} is centrosymmetric).

We specialize to the case considered in the main text where we require coverage without local restrictions on θ . In the notation of Appendices C.3 and C.4, this corresponds to $\mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$ for a convex (but possibly asymmetric) set \mathcal{C} .

In the main text, we focused on the case where \mathcal{C} is centrosymmetric. To allow for general convex \mathcal{C} , we use estimators that are asymptotically affine, rather than linear. We focus on one-step estimators, which take the form

$$\hat{h} = h(\hat{\theta}_{\text{initial}}) + \hat{k}' g(\hat{\theta}_{\text{initial}}) + \hat{a}/\sqrt{n}.$$

for some vector \hat{k} and some scalar \hat{a} . We continue to require the condition

$$\hat{H} = -\hat{k}' \hat{\Gamma}, \tag{42}$$

where $\hat{\Gamma}$ is an estimator of Γ satisfying conditions to be given below.

To deal with asymmetric \mathcal{C} , and to state results involving worst-case quantiles of excess length over different sets, it will be helpful to separately define worst-case upper and lower bias. For a set $\mathcal{C} \in \mathbb{R}^{d_\theta}$, let

$$\overline{\text{bias}}_{\mathcal{C}}(k, a) = \sup_{c \in \mathcal{C}} k'c + a, \quad \underline{\text{bias}}_{\mathcal{C}}(k, a) = \inf_{c \in \mathcal{C}} k'c + a$$

A one-sided asymptotic $1 - \alpha$ CI is given by $[\hat{c}, \infty)$ where

$$\begin{aligned} \hat{c} &= \hat{h} - \overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a})/\sqrt{n} - z_{1-\alpha} \sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\sqrt{n} \\ &= h(\hat{\theta}_{\text{initial}}) + \hat{k}'g(\hat{\theta}_{\text{initial}}) + \hat{a}/\sqrt{n} - \overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a})/\sqrt{n} - z_{1-\alpha} \sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\sqrt{n} \\ &= h(\hat{\theta}_{\text{initial}}) + \hat{k}'g(\hat{\theta}_{\text{initial}}) - \overline{\text{bias}}_{\mathcal{C}}(\hat{k}, 0)/\sqrt{n} - z_{1-\alpha} \sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\sqrt{n}, \end{aligned}$$

and $\hat{\Sigma}$ is an estimate of Σ . Thus, the intercept term \hat{a} does not matter for the one-sided CI and can be taken to be zero in this case. For two-sided CIs, however, the choice of \hat{a} matters, and we assume that \hat{a} is chosen so that the estimator is centered:

$$\overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a}) = \sup_{c \in \mathcal{C}} \hat{k}'c + \hat{a} = - \left(\inf_{c \in \mathcal{C}} \hat{k}'c + \hat{a} \right) = - \underline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a}). \quad (43)$$

A two-sided asymptotic $1 - \alpha$ CI is then given by $\hat{h} \pm \hat{\chi}$ where

$$\hat{\chi} = \text{cv}_\alpha \left(\overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a})/\sqrt{\hat{k}'\hat{\Sigma}\hat{k}} \right) \sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\sqrt{n}, \quad \text{where } \text{cv}_\alpha(t) \text{ is the } 1 - \alpha \text{ quantile of } |\mathcal{N}(t, 1)|.$$

For both forms of CIs, we first state a result for general \hat{k} , \hat{a} , and then specialize to optimal weights. For the one-sided case, we consider CIs that optimize worst-case length over $(\sqrt{n}(\theta - \theta^*)', c)'$ in some set \mathcal{G} , subject to coverage over $\mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$. In principle, this allows for confidence sets that “direct power” not only at particular values of c but also at particular values of θ . However, Lemma E.2 in Supplemental Appendix E.3.1 shows that the optimal weights for this problem are the same as the optimal weights when \mathcal{G} is replaced by $\mathbb{R}^{d_\theta} \times \mathcal{D}(\mathcal{G})$, where $\mathcal{D}(\mathcal{G}) = \{c: \text{there exists } s \text{ s.t. } (s', c)' \in \mathcal{G}\}$. Thus, it is without loss of generality to consider weights that optimize worst-case excess length over $c \in \mathcal{D}$ subject to coverage over $c \in \mathcal{C}$ where $\mathcal{D} \subseteq \mathcal{C}$ is a compact convex set.

The optimal weights take the form $\hat{k} = k(\delta_\beta, \hat{H}, \hat{\Gamma}, \hat{\Sigma})$ where

$$k(\delta, H, \Gamma, \Sigma)' = \frac{((c_{1,\delta}^* - c_{0,\delta}^*) - \Gamma(s_{1,\delta}^* - s_{0,\delta}^*))'\Sigma^{-1}}{((c_{1,\delta}^* - c_{0,\delta}^*) - \Gamma(s_{1,\delta}^* - s_{0,\delta}^*))'\Sigma^{-1}\Gamma H'/HH'} \quad (44)$$

and $c_{0,\delta}$, $s_{0,\delta}$, $c_{1,\delta}$, $s_{1,\delta}$ solve the between class modulus problem (31) with $\mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$ and $\mathcal{G} = \mathbb{R}^{d_\theta} \times \mathcal{D}$. For a two-sided CI of the form given above, the optimal weights take this form with $\mathcal{D} = \mathcal{C}$, δ minimizing $\hat{\chi}$, and with \hat{a} chosen to center the CI so that (43) holds. We note that, in the case where $\mathcal{D} = \mathcal{C}$ and \mathcal{C} is centrosymmetric, $s_{1,\delta}^* = s_{0,\delta}^*$ and $c_{1,\delta}^* = c_{0,\delta}^*$, and (31) reduces to two times the optimization problem (15). The weights \hat{k} then take the form given in (14) in the main text, and, since \mathcal{C} is centrosymmetric, $\hat{a} = 0$, which gives the two-sided CI proposed in the main text.

For our general result showing coverage for possibly suboptimal weights \hat{k} , \hat{a} , we make the following assumptions. In the following, for a set \mathcal{A}_n , random variables $A_{n,\theta,P}$ and $B_{n,\theta,P}$ and a sequence a_n , we say $A_{n,\theta,P} = B_{n,\theta,P} + o_P(a_n)$ uniformly over (θ, P) in \mathcal{A}_n if, for all $\varepsilon > 0$, $\sup_{(\theta,P) \in \mathcal{A}_n} P(a_n^{-1} \|A_{n,\theta,P} - B_{n,\theta,P}\| > \varepsilon) \rightarrow 0$. We say $A_{n,\theta,P} = B_{n,\theta,P} + \mathcal{O}_P(a_n)$ uniformly over (θ, P) in a set \mathcal{A}_n if $\lim_{C \rightarrow \infty} \limsup_{n \rightarrow \infty} \sup_{(\theta,P) \in \mathcal{A}_n} P(a_n^{-1} \|A_{n,\theta,P} - B_{n,\theta,P}\| > C) = 0$. In the following, the set \mathcal{S}_n defined in Appendix C.1 over which coverage is required is defined with $\mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$.

Assumption C.2. *The set \mathcal{C} is compact or takes the form $\tilde{\mathcal{C}} \times \mathbb{R}^{d_{g_2}}$ where $d_{g_1} + d_{g_2} = d_g$ and $\tilde{\mathcal{C}}$ is a compact subset of $\mathbb{R}^{d_{g_1}}$. In addition, $\hat{\theta}_{initial} - \theta = \mathcal{O}_P(1/\sqrt{n})$, $\hat{g}(\hat{\theta}_{initial}) - \hat{g}(\theta) = \Gamma_{\theta,P}(\hat{\theta}_{initial} - \theta) + o_P(1/\sqrt{n})$ and $h(\hat{\theta}_{initial}) - h(\theta) = H_\theta(\hat{\theta}_{initial} - \theta) + o_P(1/\sqrt{n})$ uniformly over $(\theta, P) \in \mathcal{S}_n$.*

Assumption C.3. *$\hat{g}(\theta) - g_P(\theta) = \mathcal{O}(1/\sqrt{n})$ uniformly over $(\theta, P) \in \mathcal{S}_n$. Furthermore, for a collection of matrices $\Sigma_{\theta,P}$ such that $k'_{\theta,P} \Sigma_{\theta,P} k_{\theta,P}$ is bounded away from zero and infinity,*

$$\sup_{t \in \mathbb{R}} \sup_{(\theta,P) \in \mathcal{S}_n} \left| P \left(\frac{\sqrt{n} k'_{\theta,P} (\hat{g}(\theta) - g_P(\theta))}{\sqrt{k'_{\theta,P} \Sigma_{\theta,P} k_{\theta,P}}} \leq t \right) - \Phi(t) \right| \rightarrow 0.$$

Assumption C.4. *$\hat{k} - k_{\theta,P} = o_P(1)$ uniformly over $(\theta, P) \in \mathcal{S}_n$, and similarly for \hat{a} , $\hat{\Gamma}$, \hat{H} and $\hat{\Sigma}$. Furthermore, $k_{\theta,P}$, $a_{\theta,P}$, $\Gamma_{\theta,P}$, H_θ and $\Sigma_{\theta,P}$ are bounded uniformly over $(\theta, P) \in \mathcal{S}_n$. In the case where $\mathcal{C} = \tilde{\mathcal{C}} \times \mathbb{R}^{d_{g_2}}$, assume that the last d_{g_2} elements of \hat{k} are zero with probability one for all $P \in \mathcal{P}$.*

Theorem C.3. *Suppose that Assumptions C.2, C.3 and C.4 hold and let \hat{c} be defined above with \hat{k} , $\hat{\Gamma}$ and \hat{H} satisfying (42). Then*

$$\liminf_{n \rightarrow \infty} \inf_{(\theta,P) \in \mathcal{S}_n} P(h(\theta) \in [\hat{c}, \infty)) \geq 1 - \alpha,$$

and

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \sup_{\theta \in \Theta_I(P; \mathbb{R}^{d_\theta} \times \mathcal{D}, \Theta_n)} \left\{ \sqrt{n} q_{\beta, P}(h(\theta) - \hat{c}) - \left[\overline{\text{bias}}_{\mathcal{C}}(k_{\theta, P}, 0) - \underline{\text{bias}}_{\mathcal{D}}(k_{\theta, P}, 0) + (z_{1-\alpha} + z_\beta) \sqrt{k'_{\theta, P} \Sigma_{\theta, P} k_{\theta, P}} \right] \right\} \leq 0.$$

Proof. If $\mathcal{C} = \tilde{\mathcal{C}} \times \mathbb{R}^{d_{g_2}}$ with $\tilde{\mathcal{C}}$ compact, the theorem can equivalently be stated as holding with \hat{k} redefined to be the vector in $\mathbb{R}^{d_{g_1}}$ that contains the first d_{g_1} elements of the original sensitivity \hat{k} , and with other objects redefined similarly. Therefore, it suffices to consider the case where \mathcal{C} is compact.

Note that

$$\begin{aligned} \sqrt{n}(\hat{h} - h(\theta)) &= H_\theta \sqrt{n}(\hat{\theta}_{\text{initial}} - \theta) + \hat{k} \sqrt{n} \hat{g}(\theta) + \hat{k} \sqrt{n}(\hat{g}(\hat{\theta}_{\text{initial}}) - \hat{g}(\theta)) + \hat{a} + o_P(1) \\ &= H_\theta \sqrt{n}(\hat{\theta}_{\text{initial}} - \theta) + \hat{k} \sqrt{n}(\hat{g}(\theta) - g_P(\theta)) + \hat{k}' c + \hat{k} \sqrt{n} \Gamma_{\theta, P}(\hat{\theta}_{\text{initial}} - \theta) + \hat{a} + o_P(1) \\ &= (H_\theta + k'_{\theta, P} \Gamma_{\theta, P}) \sqrt{n}(\hat{\theta}_{\text{initial}} - \theta) + k'_{\theta, P} c + a_{\theta, P} + k'_{\theta, P} \sqrt{n}(\hat{g}(\theta) - g_P(\theta)) + o_P(1), \end{aligned}$$

where $c = \sqrt{n} g_P(\theta)$ and the $o_P(1)$ terms are uniform over $(\theta, P) \in \mathcal{S}_n$ (the last equality uses the fact that \mathcal{C} is compact). By Assumption C.4 and (42), $H_\theta + k'_{\theta, P} \Gamma_{\theta, P} = 0$ so this implies

$$\sqrt{n}(\hat{h} - h(\theta)) = k'_{\theta, P} c + a_{\theta, P} + k'_{\theta, P} \sqrt{n}(\hat{g}(\theta) - g_P(\theta)) + o_P(1) \quad (45)$$

uniformly over $(\theta, P) \in \mathcal{S}_n$. By compactness of \mathcal{C} and Assumption C.4, we also have

$$\overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a}) = \overline{\text{bias}}_{\mathcal{C}}(k_{\theta, P}, a_{\theta, P}) + o_P(1), \quad \hat{k}' \hat{\Sigma} \hat{k} = k'_{\theta, P} \Sigma_{\theta, P} k_{\theta, P} + o_P(1)$$

uniformly over $(\theta, P) \in \mathcal{S}_n$. Thus,

$$\begin{aligned} \sqrt{n}(\hat{c} - h(\theta)) &= \sqrt{n}(\hat{h} - h(\theta)) - \overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a}) - z_{1-\alpha} \sqrt{\hat{k}' \hat{\Sigma} \hat{k}} = k'_{\theta, P} c + \\ & a_{\theta, P} + k'_{\theta, P} \sqrt{n}(\hat{g}(\theta) - g_P(\theta)) - \overline{\text{bias}}_{\mathcal{C}}(k_{\theta, P}, a_{\theta, P}) - z_{1-\alpha} \sqrt{k'_{\theta, P} \Sigma_{\theta, P} k_{\theta, P}} + o_P(1) \quad (46) \end{aligned}$$

uniformly over $(\theta, P) \in \mathcal{S}_n$. Since $k'_{\theta, P} c + a_{\theta, P} - \overline{\text{bias}}_{\mathcal{C}}(k_{\theta, P}, a_{\theta, P}) \leq 0$ by definition, the first part of the theorem (coverage) now follows from Assumption C.3. For the last part of the theorem, note that, using the above display and the fact that $k'_{\theta, P} c + a_{\theta, P} \geq \underline{\text{bias}}_{\mathcal{D}}(k_{\theta, P}, a_{\theta, P})$ for any (θ, P) with $c = \sqrt{n} E_P g(w_i, \theta) \in \mathcal{D}$, it follows that $\sqrt{n}(h(\theta) - c)$ is less than or equal to

$$\overline{\text{bias}}_{\mathcal{C}}(k_{\theta, P}, a_{\theta, P}) - \underline{\text{bias}}_{\mathcal{D}}(k_{\theta, P}, a_{\theta, P}) + z_{1-\alpha} \sqrt{k'_{\theta, P} \Sigma_{\theta, P} k_{\theta, P}} + k'_{\theta, P} \sqrt{n}(\hat{g}(\theta) - g_P(\theta)) + o_P(1)$$

uniformly over (θ, P) with $\sqrt{n}E_{Pg}(w_i, \theta) \in \mathcal{D}$. This, along with Assumption C.3, gives the last part of the theorem. \square

Theorem C.4. *Suppose that Assumptions C.2, C.3 and C.4 hold and let \hat{h} and $\hat{\chi}$ be defined above with \hat{k} , \hat{a} , $\hat{\Gamma}$ and \hat{H} satisfying (42) and (43). Then*

$$\liminf_{n \rightarrow \infty} \inf_{(\theta, P) \in \mathcal{S}_n} P \left(h(\theta) \in \{\hat{h} \pm \hat{\chi}\} \right) \geq 1 - \alpha.$$

In addition, we have

$$\sqrt{n}\hat{\chi} - \text{cv}_\alpha \left(\frac{\overline{\text{bias}}_{\mathcal{C}}(k_{\theta, P}, a_{\theta, P})}{\sqrt{k'_{\theta, P} \Sigma_{\theta, P} k_{\theta, P}}} \right) \sqrt{k'_{\theta, P} \Sigma_{\theta, P} k_{\theta, P}} \xrightarrow{p} 0$$

uniformly over $(\theta, P) \in \mathcal{S}_n$.

Proof. As with Theorem C.3, it suffices to consider the case where \mathcal{C} is compact. Let (θ_n, P_n) be a sequence in \mathcal{S}_n and let $c_n = \sqrt{n}g_{P_n}(\theta_n)$. Let $b_n = k'_{\theta_n, P_n}c_n + a_{\theta_n, P_n}$, $\text{sd}_n = \sqrt{k'_{\theta_n, P_n} \Sigma_{\theta_n, P_n} k_{\theta_n, P_n}}$ and $\bar{b}_n = \overline{\text{bias}}_{\mathcal{C}}(k_{\theta_n, P_n}, a_{\theta_n, P_n})$. Note that, by (43), $\overline{\text{bias}}_{\mathcal{C}}(k_{\theta_n, P_n}, a_{\theta_n, P_n}) = -\underline{\text{bias}}_{\mathcal{C}}(k_{\theta_n, P_n}, a_{\theta_n, P_n})$ when Assumption C.4 holds. It therefore follows that $-\bar{b}_n \leq b_n \leq \bar{b}_n$.

Let $Z_n = \sqrt{n}k'_{\theta_n, P_n}(\hat{g}(\theta_n) - g_{P_n}(\theta_n))/\text{sd}_n$. Note that Z_n converges in distribution (under P_n) to a $\mathcal{N}(0, 1)$ random variable by Assumption C.3. By (45),

$$\sqrt{n}(\hat{h} - h(\theta_n)) = b_n + \text{sd}_n Z_n + o_{P_n}(1).$$

Using the fact that sd_n is bounded away from zero and $\sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\text{sd}_n$ converges in probability to one under P_n , it also follows that

$$\sqrt{n}(\hat{h} - h(\theta_n))/\sqrt{\hat{k}'\hat{\Sigma}\hat{k}} = b_n/\text{sd}_n + Z_n + o_{P_n}(1).$$

Also, by Assumption C.4, we have, for a large enough constant K ,

$$\left| \text{cv}_\alpha \left(\frac{\overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a})}{\sqrt{\hat{k}'\hat{\Sigma}\hat{k}}} \right) - \text{cv}_\alpha \left(\frac{\bar{b}_n}{\text{sd}_n} \right) \right| \leq K \left\{ \left[\overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a}) - \bar{b}_n \right] + \left[\sqrt{\hat{k}'\hat{\Sigma}\hat{k}} - \text{sd}_n \right] \right\} \xrightarrow{p} 0.$$

This, along with the fact that $\sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\text{sd}_n$ converges in probability to one under P_n , gives

the second part of the theorem. Furthermore, it follows from the above display that

$$\begin{aligned} P_n \left(h(\theta_n) > \hat{h} + \hat{\chi} \right) &= P_n \left(\frac{\sqrt{n} \left(\hat{h} - h(\theta_n) \right)}{\sqrt{\hat{k}' \hat{\Sigma} \hat{k}}} < -\text{cv}_\alpha \left(\overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a}) / \sqrt{\hat{k}' \hat{\Sigma} \hat{k}} \right) \right) \\ &= P_n \left(b_n / \text{sd}_n + Z_n < -\text{cv}_\alpha \left(\bar{b}_n / \text{sd}_n \right) + o_{P_n}(1) \right) = \Phi \left(-b_n / \text{sd}_n - \text{cv}_\alpha \left(\bar{b}_n / \text{sd}_n \right) \right) + o(1). \end{aligned}$$

Similarly,

$$\begin{aligned} P_n \left(h(\theta_n) < \hat{h} - \hat{\chi} \right) &= P_n \left(\frac{\sqrt{n} \left(\hat{h} - h(\theta_n) \right)}{\sqrt{\hat{k}' \hat{\Sigma} \hat{k}}} > \text{cv}_\alpha \left(\overline{\text{bias}}_{\mathcal{C}}(\hat{k}, \hat{a}) / \sqrt{\hat{k}' \hat{\Sigma} \hat{k}} \right) \right) \\ &= P_n \left(b_n / \text{sd}_n + Z_n > \text{cv}_\alpha \left(\bar{b}_n / \text{sd}_n \right) + o_{P_n}(1) \right) = 1 - \Phi \left(-b_n / \text{sd}_n + \text{cv}_\alpha \left(\bar{b}_n / \text{sd}_n \right) \right) + o(1). \end{aligned}$$

Thus, the probability of the CI not covering is given, up to $o(1)$, by

$$1 - \Phi \left(-b_n / \text{sd}_n + \text{cv}_\alpha \left(\bar{b}_n / \text{sd}_n \right) \right) + \Phi \left(-b_n / \text{sd}_n - \text{cv}_\alpha \left(\bar{b}_n / \text{sd}_n \right) \right).$$

This is the probability that the absolute value of a $\mathcal{N}(b_n / \text{sd}_n, 1)$ variable is greater than $\text{cv}_\alpha \left(\bar{b}_n / \text{sd}_n \right)$, which is less than $1 - \alpha$ since $|b_n| \leq \bar{b}_n$. \square

We now specialize to the case where the optimal weights are used. We make a uniform consistency assumption on $\hat{\Gamma}$, \hat{H} and $\hat{\Sigma}$, as well as assumptions on the rank of H , Γ and Σ . The latter are standard regularity conditions for the correctly specified ($\mathcal{C} = \{0\}$) case.

Assumption C.5. *The estimators $\hat{\Gamma}$, \hat{H} and $\hat{\Sigma}$ are full rank with probability one and satisfy $\hat{\Gamma} - \Gamma_{\theta, P} = o_P(1)$, $\hat{H} - H_\theta = o_P(1)$ and $\hat{\Sigma} - \Sigma_{\theta, P} = o_P(1)$ uniformly over $(\theta, P) \in \mathcal{S}_n$.*

Assumption C.6. *There exists a compact set \mathcal{B} that contains the set $\{(H_\theta, \Gamma_{\theta, P}, \Sigma_{\theta, P}) : \theta \in \Theta_n, P \in \mathcal{P}\}$ for all n , such that (i) in the case where \mathcal{C} is compact, $H \neq 0$ and Γ and Σ are full rank for any $(H, \Gamma, \Sigma) \in \mathcal{B}$ or (ii) in the case where $\mathcal{C} = \tilde{\mathcal{C}} \times \mathbb{R}^{d_{g_2}}$ with $\tilde{\mathcal{C}}$ compact, the same holds for the sub-matrices corresponding to the first d_{g_1} moments.*

Using these assumptions, we can verify that Assumption C.4 holds with weights $k_{\theta, P}$ that achieve the efficiency bound in Theorem C.1 and nearly achieve the efficiency bound in Theorem C.2. This gives the following results.

Theorem C.5. *Suppose that Assumptions C.2, C.3, C.5 and C.6 hold and let \hat{c} be defined above with $\hat{k} = k(\delta_\beta, \hat{H}, \hat{\Gamma}, \hat{\Sigma})$. Then*

$$\liminf_{n \rightarrow \infty} \inf_{(\theta, P) \in \mathcal{S}_n} P(h(\theta) \in [\hat{c}, \infty)) \geq 1 - \alpha$$

and

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \sup_{\theta \in \Theta_I(P; \mathbb{R}^{d_\theta} \times \mathcal{D}, \Theta_n)} \left[\sqrt{n} q_{\beta, P}(h(\theta) - \hat{c}) - \omega(\delta_\beta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \mathcal{D}, H_\theta, \Gamma_{\theta, P}, \Sigma_{\theta, P}) \right] \leq 0.$$

Proof. In the case where \mathcal{C} is compact, it follows from Lemma E.6 in Supplemental Appendix E.3.2, $k(\delta, H, \Gamma, \Sigma)$ is continuous on $\{\delta\} \times \mathcal{B}$. Since \mathcal{B} is compact, this means that $k(\delta, H, \Gamma, \Sigma)$ is uniformly continuous. Thus, Assumption C.5 implies that \hat{k} satisfies Assumption C.4 with $k_{\theta, P} = k(\delta, H_\theta, \Gamma_{\theta, P}, \Sigma_{\theta, P})$. Furthermore, \hat{k} satisfies (42) by assumption. By properties of the modulus (Equation (24) in [Armstrong and Kolesár, 2018](#)),

$$\begin{aligned} \overline{\text{bias}}_{\mathcal{C}}(k_{\theta, P}, 0) - \underline{\text{bias}}_{\mathcal{D}}(k_{\theta, P}, 0) + (z_{1-\alpha} + z_\beta) \sqrt{k'_{\theta, P} \Sigma_{\theta, P} k_{\theta, P}} \\ = \omega(\delta_\beta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \mathcal{D}, H_\theta, \Gamma_{\theta, P}, \Sigma_{\theta, P}) \end{aligned}$$

for this $k_{\theta, P}$. Applying Theorem C.3 gives the result.

In the case where $\mathcal{C} = \tilde{\mathcal{C}} \times \mathbb{R}^{d_{g_2}}$ with $\tilde{\mathcal{C}}$ compact, the last d_{g_2} elements of \hat{k} are equal to zero as required by Assumption C.4, and the first d_{g_1} elements are the same as the weights computed from the modulus problem with the last d_{g_2} components thrown away and H , Γ and Σ redefined to be the sub-matrices corresponding to the first d_{g_1} elements of the moments. Thus, the same arguments apply in this case. \square

For two-sided CIs, we consider weights $\hat{k} = k(\delta^*(\hat{H}, \hat{\Gamma}, \hat{\Sigma}), \hat{H}, \hat{\Gamma}, \hat{\Sigma})$ given by (44) with $\mathcal{G} = \mathcal{F} = \mathbb{R}^{d_\theta} \times \mathcal{C}$, where δ^* may depend on the data through \hat{H} , $\hat{\Gamma}$ and $\hat{\Sigma}$. If δ^* is chosen to optimize CI length, it will be given by $\delta_\chi(\hat{H}, \hat{\Gamma}, \hat{\Sigma})$ where

$$\delta_\chi(H, \Gamma, \Sigma) = \underset{\delta}{\operatorname{argmin}} \operatorname{cv}_\alpha \left(\frac{\omega(\delta)}{2\omega'(\delta)} - \frac{\delta}{2} \right) \omega'(\delta) \quad (47)$$

where $\omega(\delta) = \omega(\delta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \mathcal{C}, H, \Gamma, \Sigma)$ is the single class modulus (see Section 3.4 in [Armstrong and Kolesár, 2018](#)).

We make a continuity assumption on δ^* .

Assumption C.7. δ^* is a continuous function of its arguments on the set \mathcal{B} given in Assumption C.6.

Theorem C.6. Suppose that Assumptions C.2, C.3, C.5, C.6 and C.7 hold and let \hat{h} be defined above with $\hat{k} = k(\delta^*(\hat{H}, \hat{\Gamma}, \hat{\Sigma}), \hat{H}, \hat{\Gamma}, \hat{\Sigma})$. Then the conclusion of Theorem C.4 holds. If, in addition, $\delta^* = \delta_\chi(\hat{H}, \hat{\Gamma}, \hat{\Sigma})$ for δ_χ the CI length optimizing choice of δ given in (47),

then the half-length $\hat{\chi}$ satisfies $\sqrt{n}\hat{\chi} = \chi(\theta, P) + o_P(1)$ uniformly over $(\theta, P) \in \mathcal{S}_n$, where

$$\chi(\theta, P) = \min_{\delta} \text{cv}_{\alpha} \left(\frac{\omega(\delta)}{2\omega'(\delta)} - \frac{\delta}{2} \right) \omega'(\delta), \quad \omega(\delta) = \omega(\delta; \mathbb{R}^{d_{\theta}} \times \mathcal{C}, \mathbb{R}^{d_{\theta}} \times \mathcal{C}, H_{\theta}, \Gamma_{\theta, P}, \Sigma_{\theta, P}).$$

Proof. The result follows from using the same arguments as in the proof of Theorem C.5, along with continuity of δ^* , to verify Assumption C.4. The form of the limiting half-length for the optimal weights follows from properties of the modulus (see Section 3.4 in [Armstrong and Kolesár, 2018](#)). \square

C.6 Centrosymmetric case

Theorem 4.1 in Section 4 gives a bound for two-sided CIs in the case where \mathcal{C} is centrosymmetric. This follows from applying Theorems C.2 and C.6 in the centrosymmetric case. In particular, comparing the asymptotic length in Theorem C.6 to the bound in Theorem C.2 and using the fact that $\omega(\delta; \mathbb{R}^{d_{\theta}} \times \mathcal{C}, \{0\}, H_{\theta}, \Gamma_{\theta, P}, \Sigma_{\theta, P}) = \omega(\delta; \{0\}, \mathbb{R}^{d_{\theta}} \times \mathcal{C}, H_{\theta}, \Gamma_{\theta, P}, \Sigma_{\theta, P}) = \frac{1}{2}\omega(2\delta; \mathbb{R}^{d_{\theta}} \times \mathcal{C}, \mathbb{R}^{d_{\theta}} \times \mathcal{C}, H_{\theta}, \Gamma_{\theta, P}, \Sigma_{\theta, P})$ when \mathcal{C} is centrosymmetric gives the bound $\kappa_*(H_{\theta}, \Gamma_{\theta, P_0}, \Sigma_{\theta, P_0}, \mathcal{C})$ from the statement of Theorem 4.1. This corresponds to the bound in Corollary 3.3 of [Armstrong and Kolesár \(2018\)](#). The universal lower bound for κ_* follows from the following result:

Theorem C.7. *For any H, Γ, Σ and \mathcal{C} , the efficiency κ_* given in (18) is lower bounded by*

$$(z_{1-\alpha}(1-\alpha) - \tilde{z}_{\alpha}\Phi(\tilde{z}_{\alpha}) + \phi(z_{1-\alpha}) - \phi(\tilde{z}_{\alpha}))/z_{1-\alpha/2}$$

where $\tilde{z}_{\alpha} = z_{1-\alpha} - z_{1-\alpha/2}$ and Φ and ϕ denote the standard normal cdf, and pdf respectively. The lower bound is sharp in the sense that it holds with equality if $\omega(\delta) = K_0 \min\{\delta, 2z_{1-\alpha/2}\}$, for some constant K_0 .

Proof. Since $\text{cv}_{\alpha}(b) \leq b + z_{1-\alpha/2}$, the denominator in (18) is upper-bounded by

$$\begin{aligned} \min_{\delta} 2 \text{cv}_{\alpha} \left(\frac{\omega(\delta)}{2\omega'(\delta)} - \frac{\delta}{2} \right) \omega'(\delta) &\leq \\ &2 \text{cv}_{\alpha} \left(\frac{\omega(2z_{1-\alpha/2})}{2\omega'(2z_{1-\alpha/2})} - z_{1-\alpha/2} \right) \omega'(2z_{1-\alpha/2}) \leq \omega(2z_{1-\alpha/2}). \end{aligned} \quad (48)$$

On the other hand, the numerator in (18) can be decomposed as

$$\begin{aligned} (1-\alpha)E[\omega(2(z_{1-\alpha} - Z)) \mid Z \leq z_{1-\alpha}] &= E[\omega(2(z_{1-\alpha} - Z)) \mathbb{I}\{Z \leq z_{1-\alpha} - z_{1-\alpha/2}\}] \\ &\quad + E[\omega(2(z_{1-\alpha} - Z)) \mathbb{I}\{z_{1-\alpha} - z_{1-\alpha/2} \leq Z \leq z_{1-\alpha}\}]. \end{aligned}$$

Since the modulus $\omega(\delta)$ is non-decreasing, the first summand is lower-bounded by

$$E \left[\omega(2z_{1-\alpha/2}) \mathbb{I}\{Z \leq z_{1-\alpha} - z_{1-\alpha/2}\} \right] = \omega(2z_{1-\alpha/2})\Phi(z_{1-\alpha} - z_{1-\alpha/2}).$$

Since the modulus $\omega(\delta)$ is concave, $\omega(2(z_{1-\alpha} - Z)) \geq (z_{1-\alpha} - Z)/z_{1-\alpha/2} \cdot \omega(2z_{1-\alpha/2})$, so that the second summand is lower-bounded by

$$\begin{aligned} & \frac{\omega(2z_{1-\alpha/2})}{z_{1-\alpha/2}} E \left[(z_{1-\alpha} - Z) \mathbb{I}\{z_{1-\alpha} - z_{1-\alpha/2} \leq Z \leq z_{1-\alpha}\} \right] \\ &= \frac{\omega(2z_{1-\alpha/2})}{z_{1-\alpha/2}} \left(z_{1-\alpha}(1 - \alpha - \Phi(z_{1-\alpha} - z_{1-\alpha/2})) + \phi(z_{1-\alpha}) - \phi(z_{1-\alpha} - z_{1-\alpha/2}) \right), \end{aligned}$$

where the equality follows by the formula for the expectation of a truncated normal random variable. Combining the two preceding displays then yields

$$\begin{aligned} (1 - \alpha)E \left[\omega(2(z_{1-\alpha} - Z)) \mid Z \leq z_{1-\alpha} \right] \\ \geq \omega(2z_{1-\alpha/2}) \frac{z_{1-\alpha}(1 - \alpha) - \tilde{z}_\alpha \Phi(\tilde{z}_\alpha) + \phi(z_{1-\alpha}) - \phi(\tilde{z}_\alpha)}{z_{1-\alpha/2}}, \end{aligned} \quad (49)$$

where $\tilde{z}_\alpha = z_{1-\alpha} - z_{1-\alpha/2}$. Combining this with the bound in (48) then yields the result. The sharpness of the bound for the case $\omega(\delta) = K_0 \min\{\delta, 2z_{1-\alpha/2}\}$ follows from by noting that in this case, both (48) and (49) hold as equalities. \square

For the one-sided case, we obtain the following bound.

Theorem C.8. *Consider the setting of Theorem C.5, with \mathcal{C} centrosymmetric. Then the weights $\hat{k} = \hat{k}(\delta_\beta, \hat{H}, \hat{\Gamma}, \hat{\Sigma})$ with $\mathcal{D} = \mathcal{C}$ are identical to the weights $\hat{k}(\delta_\beta, \hat{H}, \hat{\Gamma}, \hat{\Sigma})$ computed with $\mathcal{D} = \{0\}$, but with $\tilde{\beta} = \Phi((z_\beta - z_{1-\alpha})/2)$. Furthermore, letting \hat{c}_{minimax} denote the lower endpoint of the CI computed with these weights ($\hat{k}(\delta_\beta, \hat{H}, \hat{\Gamma}, \hat{\Sigma})$ with $\mathcal{D} = \mathcal{C}$), we have*

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \sup_{\theta \in \Theta_I(P; \mathbb{R}^{d_\theta} \times \{0\}, \Theta_n)} \left\{ \sqrt{n} q_{\beta, P}(h(\theta) - \hat{c}_{\text{minimax}}) - \frac{1}{2} [\omega_{\theta, P}(\delta_\beta) + \delta_\beta \omega'_{\theta, P}(\delta_\beta)] \right\} \leq 0$$

where $\omega_{\theta, P}(\delta) = \omega(\delta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \mathcal{C}, H_\theta, \Gamma_{\theta, P}, \Sigma_{\theta, P})$. For \hat{c} computed instead with $\mathcal{D} = \{0\}$, we obtain

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \sup_{\theta \in \Theta_I(P; \mathbb{R}^{d_\theta} \times \{0\}, \Theta_n)} \left\{ \sqrt{n} q_{\beta, P}(h(\theta) - \hat{c}) - \frac{1}{2} \omega_{\theta, P}(2\delta_\beta) \right\} \leq 0.$$

Proof. The first statement follows from Corollary 3.2 in [Armstrong and Kolesár \(2018\)](#). The second statement follows from applying Theorem C.3 as in the proof of Theorem C.5,

noting that $\text{bias}_{\{0\}}(k_{\theta,P}, 0) = 0$, and using arguments from the proof of Corollary 3.2 in [Armstrong and Kolesár \(2018\)](#). The last statement follows from Theorem C.5 and the fact that $\omega(\delta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \{0\}, H_\theta, \Gamma_{\theta,P}, \Sigma_{\theta,P}) = \frac{1}{2}\omega(2\delta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \mathcal{C}, H_\theta, \Gamma_{\theta,P}, \Sigma_{\theta,P})$. \square

Thus, directing power toward the correctly specified case yields the same one-sided CI once one changes the quantile over which one optimizes excess length. If one does attempt to direct power, the scope for doing so is bounded by a factor of

$$\kappa_*^{\text{OCI},\beta}(H_\theta, \Gamma_{\theta,P_0}, \Sigma_{\theta,P_0}, \mathcal{C}) = \frac{\omega_{\theta,P}(2\delta_\beta)}{\omega_{\theta,P}(\delta_\beta) + \delta_\beta \omega'_{\theta,P}(\delta_\beta)}. \quad (50)$$

This gives a bound for the one-sided case analogous to the bound κ_* in (18) for two-sided CIs.

A consistent estimate of these bounds can be obtained by plugging in $\omega(\delta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \mathcal{C}, \hat{H}, \hat{\Gamma}, \hat{\Sigma})$ for $\omega_{\theta,P}(\delta) = \omega(\delta; \mathbb{R}^{d_\theta} \times \mathcal{C}, \mathbb{R}^{d_\theta} \times \mathcal{C}, H_\theta, \Gamma_{\theta,P}, \Sigma_{\theta,P})$. Table 2 reports estimates of this bound under different forms of misspecification in the empirical application in Section 6.

D Global Misspecification

We now describe two approaches to the construction of CIs that are robust to global misspecification. The first approach is generally applicable, and, for a one-sided CIs, yields CIs that are asymptotically equivalent under local misspecification to the CIs proposed in the main text. The second approach also exhibits this equivalence property for two-sided CIs, but the regularity conditions it imposes may not be satisfied in all applications.

Before describing the procedures in Appendices D.1 and D.2 below, let us briefly describe the setup. Under global misspecification, the true parameter θ_0 satisfies

$$g_P(\theta_0) = E_P g(w_i, \theta_0) = \tilde{c}, \quad \tilde{c} \in \tilde{\mathcal{C}}, \quad (51)$$

where $\tilde{\mathcal{C}}$ is fixed with the sample size n . To accommodate both local and global misspecification with the same notation, we consider a sequence $\tilde{\mathcal{C}} = \tilde{\mathcal{C}}_n$ of sets. Under global misspecification $\tilde{\mathcal{C}}_n$ is fixed with n , whereas, under local misspecification, $\tilde{\mathcal{C}}_n = \mathcal{C}/\sqrt{n} = \{c/\sqrt{n} : c \in \mathcal{C}\}$ where \mathcal{C} is fixed with n . The rest of the setup is the same as the formal setup in Section 4.2 and appendix C: we are interested in a CI for $h(\theta)$ that satisfies the coverage condition (20), where $\mathcal{S}_n = \{(\theta, P) \in \Theta_n \times \mathcal{P} : g_P(\theta) \in \tilde{\mathcal{C}}\}$ denotes the set of pairs (θ, P) such that θ is in the identified set under P .

For concreteness, we focus on GMM estimators. We treat the weighting matrix W as given, and construct CIs that are asymptotically equivalent to the CIs given in Section 2.1

with sensitivity $k' = -H(\Gamma'W\Gamma)^{-1}\Gamma'W$. To make these CIs optimal under local misspecification, we need to choose the weighting matrix W so that this sensitivity is optimal under local misspecification. This can be done by computing the optimal sensitivity \hat{k} under local misspecification using first stage estimates, following our implementation in Section 3, and then computing an equivalent GMM weighting matrix as described in Remark 3.2.

D.1 CIs based on recentering the moments

Let $\mathcal{I}_{\tilde{c}}$ be a family of CIs, indexed by $\tilde{c} \in \tilde{\mathcal{C}}$, such that, for each \tilde{c} , $\mathcal{I}_{\tilde{c}}$ is asymptotically valid for the GMM model defined by the moment function $\theta \mapsto g(w_i, \theta) - \tilde{c}$. We consider the CI $\mathcal{I} = \cup_{\tilde{c} \in \tilde{\mathcal{C}}} \mathcal{I}_{\tilde{c}}$. Since this CI contains a CI based on the moment conditions $\theta \mapsto \hat{g}(\theta) - \tilde{c}_0$, where $\tilde{c}_0 = E_P g(w_i, \theta_0)$ is the true value of \tilde{c} , it will have correct asymptotic coverage under standard conditions. (While we omit a formal statement, we note that this follows by showing that \mathcal{I}_{c_n} has correct coverage under a drifting sequence of moment functions $\tilde{g}_n(w_i, \theta) = g(w_i, \theta) - \tilde{c}_n$ for sequences $\tilde{c}_n \in \tilde{\mathcal{C}}$. This follows from the usual arguments for asymptotic coverage of GMM estimators under correct specification.) As we discuss in Appendix D.1.1 below, this CI can be computed using constrained optimization, where the objective function will typically be smooth so long as $\theta \mapsto g(w_i, \theta)$ is smooth. However, since the problem involves minimization of a GMM objective function, it will typically not be convex.

We now show that this approach can be used to construct a one-sided CI that is asymptotically equivalent under local misspecification to the CI proposed in the main text. We consider CIs based on the GMM estimator $\hat{\theta}_{W, \tilde{c}} = \arg \min_{\theta} [\hat{g}(\theta) - \tilde{c}]' W [\hat{g}(\theta) - \tilde{c}]$ based on the moment function $g(w_i, \theta) - \tilde{c}$. Let $\hat{\Gamma}_{\theta}$ denote an estimate of the Jacobian matrix $\Gamma_{\theta, P} = \frac{d}{dt} E_P g(w_i, t) |_{t=\theta}$. Under local misspecification, the identified set for θ shrinks toward a point, so that the same estimate $\hat{\Sigma}$ is consistent for $\Sigma_{\theta, P}$ uniformly over the identified set. This is no longer the case under global misspecification, and we instead need to use a class of estimates $\hat{\Sigma}_{\theta}$ for $\Sigma_{\theta, P}$ indexed by θ . In the i.i.d. case, we can take $\hat{\Sigma}_{\theta} = \frac{1}{n} \sum_{i=1}^n g(w_i, \theta) g(w_i, \theta)'$, in which case the estimate of $\hat{\Sigma}$ in Section 3 corresponds to $\hat{\Sigma} = \hat{\Sigma}_{\hat{\theta}_{\text{initial}}}$. We allow for the possibility that W is data dependent, in which case we will assume that it converges to some W_P , which may depend on the underlying distribution P . Let $\hat{k}'_{\theta} = -H_{\theta} \left(\hat{\Gamma}'_{\theta} W \hat{\Gamma}_{\theta} \right)^{-1} \hat{\Gamma}'_{\theta} W$ denote an estimate of the sensitivity $k'_{\theta, P} = -H_{\theta} \left(\Gamma'_{\theta, P} W_P \Gamma_{\theta, P} \right)^{-1} \Gamma'_{\theta, P} W$. This gives the standard error $se_{\hat{\theta}_{W, \tilde{c}}}$ where $se_{\theta} = \sqrt{\hat{k}'_{\theta} \hat{\Sigma}_{\theta} \hat{k}_{\theta}} / n$. The nominal $1 - \alpha$ one-sided CI based on the given \tilde{c} is then $[h(\hat{\theta}_{W, \tilde{c}}) - z_{1-\alpha} se_{W, \tilde{c}}, \infty)$. The lower endpoint of a one-sided CI that is asymptotically valid for the set $\tilde{\mathcal{C}}$ is then given by $[\hat{c}_{\text{glob}}, \infty)$ where

$$\hat{c}_{\text{glob}} = \inf_{\tilde{c} \in \tilde{\mathcal{C}}} \left[h(\hat{\theta}_{W, \tilde{c}}) - z_{1-\alpha} se_{\hat{\theta}_{W, \tilde{c}}} \right].$$

The CI proposed in the main text takes the form $[\hat{c}_{\text{loc}}, \infty)$ where

$$\hat{c}_{\text{loc}} = h(\hat{\theta}_{\text{initial}}) + \hat{k}'\hat{g}(\hat{\theta}_{\text{initial}}) - \sup_{c \in \mathcal{C}} \hat{k}'c/\sqrt{n} - z_{1-\alpha} \sqrt{\hat{k}'\hat{\Sigma}\hat{k}}/\sqrt{n}.$$

We assume that the sensitivity \hat{k} corresponds to the same GMM weighting matrix, so that Assumptions C.2, C.3 and C.4 hold with $k_{\theta,P}$, $\Gamma_{\theta,P}$, H_{θ} and $\Sigma_{\theta,P}$ given above. In addition, we use some further regularity conditions for the estimator $\hat{\theta}_{W,\tilde{c}}$.

Assumption D.1. *The estimator $\hat{\theta}_{W,\tilde{c}}$ satisfies*

$$\sup_{c \in \mathcal{C}} \left| h(\hat{\theta}_{W,c/\sqrt{n}}) - h(\theta) - k'_{\theta,P}[\hat{g}(\theta) - c/\sqrt{n}] \right| = o_P(1/\sqrt{n})$$

and

$$\sup_{c \in \mathcal{C}} \left| \hat{k}'_{\hat{\theta}_{W,c/\sqrt{n}}} \hat{\Sigma}_{\hat{\theta}_{W,c/\sqrt{n}}} \hat{k}_{\hat{\theta}_{W,c/\sqrt{n}}} - k'_{\theta,P} \Sigma_{\theta,P} k_{\theta,P} \right| = o_P(1)$$

uniformly over $(\theta, P) \in \mathcal{S}_n$, where $\tilde{\mathcal{C}} = \tilde{\mathcal{C}}_n = \mathcal{C}/\sqrt{n}$.

Assumption D.1 imposes an influence function representation for the GMM estimator $\hat{\theta}_{W,c}$, and a uniform consistency condition for the asymptotic variance estimator. We verify this assumption under primitive conditions in the linear IV setting in Supplemental Appendix E.2.3. Note that, while the CI is robust to global misspecification, Assumption D.1 imposes conditions that are required only under local misspecification.

Theorem D.1. *Let \hat{c}_{loc} and \hat{c}_{glob} be given above with $\tilde{\mathcal{C}} = \mathcal{C}/\sqrt{n}$, where \mathcal{C} is a compact set. Suppose that Assumptions C.2, C.3 and C.4 hold for \hat{c}_{loc} , and Assumption D.1 holds for \hat{c}_{glob} , with the same $k_{\theta,P}$ and $\Sigma_{\theta,P}$, and assume that, for each $P \in \mathcal{P}$, there exists $\theta \in \Theta_n$ such that $(\theta, P) \in \mathcal{S}_n$ (i.e. the identified set for θ under each $P \in \mathcal{P}$ is nonempty). Then $\sqrt{n}(\hat{c}_{\text{loc}} - \hat{c}_{\text{glob}})$ converges in probability to zero uniformly over $P \in \mathcal{P}$.*

Proof. Under these assumptions,

$$\begin{aligned} \hat{c}_{\text{glob}} &= \inf_{c \in \mathcal{C}} h(\theta) + k'_{\theta,P} \hat{g}(\theta) - k'_{\theta,P} c/\sqrt{n} - z_{1-\alpha} \sqrt{k'_{\theta,P} \Sigma_{\theta,P} k_{\theta,P}}/\sqrt{n} + o_P(1/\sqrt{n}) \\ &= h(\theta) + k'_{\theta,P} \hat{g}(\theta) - \overline{\text{bias}}_{\mathcal{C}}(k_{\theta,P}, 0)/\sqrt{n} - z_{1-\alpha} \sqrt{k'_{\theta,P} \Sigma_{\theta,P} k_{\theta,P}}/\sqrt{n} + o_P(1/\sqrt{n}) \end{aligned}$$

uniformly over $(\theta, P) \in \mathcal{S}_n$. The result follows by noting that this matches the asymptotic expression for \hat{c}_{loc} given in Eq. (46) in the proof of Theorem C.3. \square

The CI $[\hat{c}_{\text{glob}}, \infty)$ is valid under global misspecification under standard conditions, and Theorem D.1 shows that this CI is asymptotically equivalent to the CI $[\hat{c}_{\text{loc}}, \infty)$ under local misspecification.

D.1.1 Computation

The problem of computing this CI can be written as a nested optimization problem, in which one minimizes the GMM objective function for a given \tilde{c} , and then minimizes the lower endpoint of the CI over \tilde{c} . Alternatively, in the spirit of recent papers on MPEC (e.g. [Dubé et al., 2012](#)), one can write this as an optimization problem over θ, \tilde{c} subject to the constraint that θ minimizes the GMM objective under \tilde{c} :

$$\min_{\theta, \tilde{c}} h(\theta) - z_{1-\alpha} \text{se}_\theta \quad \text{s.t.} \quad \theta = \arg \min_t [\hat{g}(t) - \tilde{c}]' W [\hat{g}(t) - \tilde{c}], \quad \tilde{c} \in \tilde{\mathcal{C}}.$$

In the case where $\hat{g}(\theta)$ is smooth, one may also relax the constraint that θ minimizes the GMM objective by instead imposing only the first order conditions:

$$\min_{\theta, \tilde{c}} h(\theta) - z_{1-\alpha} \text{se}_\theta \quad \text{s.t.} \quad \hat{\Gamma}'_\theta W [\hat{g}(\theta) - \tilde{c}] = 0, \quad \tilde{c} \in \tilde{\mathcal{C}}.$$

Since a constraint is relaxed, this can only make the resulting CI more conservative.

D.2 CIs based on misspecification-robust standard errors

In some cases, it may be possible to construct estimates \hat{B} of the worst-case asymptotic bias of the estimator \hat{h} that are asymptotically normal. In such cases, one can construct CIs that are valid under global misspecification by using misspecification-robust standard errors.

Let $\theta_c^* = \arg \min_\theta (g_P(\theta) - c)' W (g_P(\theta) - c)$, so that under the model (51), the true parameter is given by $\theta_0 = \theta_c^*$, and the pseudo-true parameter, the estimand of the GMM estimator $\hat{\theta}$ with weighting matrix W , is given by θ_0^* . Note that θ_c^* depends on P , but we leave this dependence implicit for clarity of notation. [Hall and Inoue \(2003\)](#) show that under regularity conditions, if $\tilde{\mathcal{C}}$ is fixed with n ,

$$\sqrt{n}(\hat{h} - h(\theta_0^*)) \xrightarrow{d} \mathcal{N}(0, \Omega_h),$$

where, unlike in the correctly specified case, the asymptotic variance Ω_h generally depends on the weighting matrix W . [Hall and Inoue \(2003\)](#) also show how to construct consistent estimates of Ω_h . Let $B_{\tilde{\mathcal{C}}} = \max_{\tilde{c} \in \tilde{\mathcal{C}}} |h(\theta_0^*) - h(\theta_{\tilde{c}}^*)|$ denote the worst-case bias. Suppose that we have available an estimator \hat{B} of $B_{\tilde{\mathcal{C}}}$, such that (\hat{h}, \hat{B}) are jointly asymptotically normal. Then, by the delta method, for some asymptotic variances Ω_+ and Ω_- , it holds that

$$\frac{\sqrt{n}(\hat{h} - \hat{B} - (h(\theta_0^*) - B_{\tilde{\mathcal{C}}}))}{\Omega_-^{1/2}} \xrightarrow{d} \mathcal{N}(0, 1), \quad \frac{\sqrt{n}(\hat{h} + \hat{B} - (h(\theta_0^*) + B_{\tilde{\mathcal{C}}}))}{\Omega_+^{1/2}} \xrightarrow{d} \mathcal{N}(0, 1). \quad (52)$$

Note that $B_{\tilde{C}}$, Ω_+ and Ω_- may depend on P , although we leave this implicit in the notation.

If one has available estimators $\hat{\Omega}_+$ and $\hat{\Omega}_-$ that satisfy

$$\hat{\Omega}_+/\Omega_+ \xrightarrow{p} 1, \quad \hat{\Omega}_-/\Omega_- \xrightarrow{p} 1, \quad (53)$$

one can construct one-sided CIs as $[\hat{h} - \hat{B} - z_{1-\alpha}\hat{\Omega}_-^{1/2}/\sqrt{n}, \infty)$, and $(-\infty, \hat{h} + \hat{B} + z_{1-\alpha}\hat{\Omega}_+^{1/2}/\sqrt{n}]$.

Two-sided CIs can be constructed as

$$\widetilde{CI} = \left[\hat{h} - cv_\alpha(\sqrt{n}\hat{B}/\hat{\Omega}_-^{1/2}) \cdot \hat{\Omega}_-^{1/2}/\sqrt{n}, \hat{h} + cv_\alpha(\sqrt{n}\hat{B}/\hat{\Omega}_+^{1/2}) \cdot \hat{\Omega}_+^{1/2}/\sqrt{n} \right].$$

The next result shows that these CIs are valid under global misspecification.

Theorem D.2. *Suppose that the convergence in (52) and (53) holds uniformly over $(\theta, P) \in \mathcal{S}_n = \{(\theta, P) \in \Theta_n \times \mathcal{P} : g_P(\theta) \in \tilde{\mathcal{C}}\}$, with $\tilde{\mathcal{C}}$ is fixed. Then*

$$\liminf_{n \rightarrow \infty} \inf_{(\theta, P) \in \mathcal{S}_n} P(h(\theta) \in [\hat{h} - \hat{B} - z_{1-\alpha}\hat{\Omega}_-^{1/2}/\sqrt{n}, \infty)) \geq 1 - \alpha.$$

Suppose, in addition, that $\hat{B}/B_{\tilde{C}} \xrightarrow{p} 1$ and $|\Omega_+ - \Omega_-|/\sqrt{n}B_{\tilde{C}} \rightarrow 0$ uniformly over $(\theta, P) \in \mathcal{S}_n$. Then

$$\liminf_{n \rightarrow \infty} \inf_{(\theta, P) \in \mathcal{S}_n} P(h(\theta) \in \widetilde{CI}) \geq 1 - \alpha.$$

The proof of this theorem is deferred to Appendix D.2.1. Under global misspecification, when $\sqrt{n}B_{\tilde{C}} \rightarrow \infty$, the condition $|\Omega_+ - \Omega_-|/\sqrt{n}B_{\tilde{C}} \rightarrow 0$ holds if Ω_+ and Ω_- are of the same order, which is typically the case. In this case, \widetilde{CI} is asymptotically equivalent to the CI $[\hat{h} - \hat{B} - z_{1-\alpha}\hat{\Omega}_-/\sqrt{n}, \hat{h} + \hat{B} + z_{1-\alpha}\hat{\Omega}_+/\sqrt{n}]$. Since in large samples, the uncertainty about the endpoints of the identified set $[h(\theta_0^*) - B_{\tilde{C}}, h(\theta_0^*) + B_{\tilde{C}}]$ dominates by the uncertainty about the location of the endpoints, it suffices to use a one-sided critical value $z_{1-\alpha}$ (see [Imbens and Manski, 2004](#), for a discussion).

Under local misspecification, if the estimator $\hat{\theta}$ is asymptotically linear with sensitivity k , $\sqrt{n}B_{\tilde{C}} = \overline{\text{bias}}_{\sqrt{n}\tilde{\mathcal{C}}}(k)$ is bounded, so that the condition $|\Omega_+ - \Omega_-|/\sqrt{n}B_{\tilde{C}} \rightarrow 0$ holds if $\Omega_+ - \Omega_- \rightarrow 0$. This is indeed the case if $\sqrt{n}\hat{B} = \overline{\text{bias}}_{\sqrt{n}\tilde{\mathcal{C}}}(k) + o_p(1)$, since then Ω_+ and Ω_- both equal $k'\Sigma k$. In this case \widetilde{CI} is asymptotically equivalent to the CI in Eq. (7). The CI thus automatically adapts to the misspecification magnitude.

Example D.1 (linear IV model). To give an example of a setting in which the condition (52) holds, consider the linear instrumental variables (IV) model from Section 5.1. In particular, suppose $h = H\theta$, and suppose that $g(\theta) = E[z_i(y_i - x_i'\theta)] \in \tilde{\mathcal{C}}$, with $\tilde{\mathcal{C}} = \{\sqrt{n}B\gamma : \|\gamma\| \leq M_n\}$, $B = E[z_i z_i']$, and for concreteness, suppose $\|\cdot\|$ corresponds to an ℓ_2 norm. If $M_n = M$ is

fixed, this reduces to the local misspecification setup in the main text, but if $M_n = \sqrt{n}M$, the misspecification is global. Consider the 2SLS estimator $\hat{h} = \hat{k}' \sum_{i=1}^n z_i y_i$, with $\hat{k} = -H(\hat{\Gamma}'\hat{W}\hat{\Gamma})^{-1}\hat{\Gamma}'\hat{W}$, $\hat{\Gamma} = -n^{-1} \sum_{i=1}^n z_i x_i'$ and $\hat{W}^{-1} = \sum_{i=1}^n z_i z_i'$; let $\Gamma = -E[z_i x_i']$, $W = E[z_i z_i']^{-1}$, and $k = -H(\Gamma'W\Gamma)^{-1}\Gamma'W$.

Then $h(\theta_0^*) = H\theta + k'E[z_i z_i']\gamma/\sqrt{n}$, and $B_{\tilde{c}} = \|k'E[z_i z_i']\|M_n/\sqrt{n}$. Consider the estimator $\hat{B} = \|\hat{k}' \sum_{i=1}^n z_i z_i'\|M_n/\sqrt{n}$ of the worst-case bias, which is the same as the estimator $\overline{\text{bias}}_{\sqrt{n}\tilde{c}}(\hat{k})/\sqrt{n}$ under local misspecification used in the main text. Since \hat{B} and \hat{h} depend on the data only through the sample means $S_n = n^{-1}(\sum_i z_i z_i', \sum_i x_i z_i', \sum_i x_i y_i')$, Eq. (52) holds by the delta method, and consistent estimates $\hat{\Omega}_+$ and $\hat{\Omega}_-$ of Ω_+ and Ω_- can be constructed using a consistent estimator of the asymptotic variance of S_n , which yields the CI

$$\widetilde{CI} = \left[\hat{h} - \text{cv}_\alpha(\overline{\text{bias}}_{\sqrt{n}\tilde{c}}(\hat{k})/\hat{\Omega}_-^{1/2}) \cdot \hat{\Omega}_-^{1/2}/\sqrt{n}, \hat{h} + \text{cv}_\alpha(\overline{\text{bias}}_{\sqrt{n}\tilde{c}}(\hat{k})/\hat{\Omega}_+^{1/2}) \cdot \hat{\Omega}_+^{1/2}/\sqrt{n} \right].$$

Thus, relative to the CI described in the main text, \widetilde{CI} differs only in that it uses variance estimates $\hat{\Omega}_+$ and $\hat{\Omega}_-$ that are valid under global misspecification. If $M_n = M$, so that misspecification is local, $\hat{\Omega}_+ = k'\Sigma k + o_p(1)$ and $\hat{\Omega}_- = k'\Sigma k + o_p(1)$, and the CI is asymptotically equivalent to the CI described in the main text.

D.2.1 Proof of Theorem D.2

Proof. Let $Z_- = \sqrt{n}(\hat{h} - \hat{B} - (h(\theta_0^*) - B_{\tilde{c}}))/\Omega_-^{1/2}$, and let $Z_+ = \sqrt{n}(\hat{h} + \hat{B} - (h(\theta_0^*) + B_{\tilde{c}}))/\Omega_+^{1/2}$. Then

$$\begin{aligned} P(h(\theta) \geq \hat{h} - \hat{B} - z_{1-\alpha}\hat{\Omega}_-^{1/2}/\sqrt{n}) \\ &= P(Z_- \leq \sqrt{n}(B_{\tilde{c}} - (h(\theta_0^*) - h(\theta)))/\Omega_-^{1/2} + z_{1-\alpha}\hat{\Omega}_-^{1/2}/\Omega_-^{1/2}) \\ &\geq P(Z_- \leq z_{1-\alpha}\hat{\Omega}_-^{1/2}/\Omega_-^{1/2}) \geq 1 - \alpha + o(1), \end{aligned}$$

where the equality follows by definition of Z_- , the first inequality follows since $B_{\tilde{c}} \geq h(\theta_0^*) - h(\theta)$ by definition of $B_{\tilde{c}}$, and the second inequality follows since $\hat{\Omega}_-^{1/2}/\Omega_-^{1/2} \xrightarrow{p} 1$, $Z_- \xrightarrow{d} \mathcal{N}(0, 1)$, and since convergence in distribution to a continuous distribution implies uniform convergence of the cdfs (van der Vaart, 1998, Lemma 2.11). To show the result for the two-sided CI, let $b = h(\theta^*) - h(\theta)$ denote the asymptotic bias. Then

$$\begin{aligned} P(h(\theta) \in \widetilde{CI}) &= P(\text{cv}_\alpha(\sqrt{n}\hat{B}/\hat{\Omega}_-^{1/2}) \cdot \hat{\Omega}_-^{1/2}/\sqrt{n} \geq \hat{h} - h(\theta) \geq -\text{cv}_\alpha(\sqrt{n}\hat{B}/\hat{\Omega}_+^{1/2}) \cdot \hat{\Omega}_+^{1/2}/\sqrt{n}) \\ &= P(Z_- + \sqrt{nb}/\Omega_-^{1/2} \leq A_-) + P(-Z_+ - \sqrt{nb}/\Omega_+^{1/2} \leq A_+) - 1, \end{aligned}$$

where $A_- = \sqrt{n}(B_{\bar{c}} - \hat{B})/\Omega_-^{1/2} + \text{cv}_\alpha(\sqrt{n}\hat{B}/\hat{\Omega}_-^{1/2}) \cdot \hat{\Omega}_-^{1/2}/\Omega_-^{1/2}$ and $A_+ = \text{cv}_\alpha(\sqrt{n}\hat{B}/\hat{\Omega}_+^{1/2}) \cdot \hat{\Omega}_+^{1/2}/\Omega_+^{1/2} + \sqrt{n}(B_{\bar{c}} - \hat{B})/\Omega_+^{1/2}$. Now, since $\hat{\Omega}_-/\Omega_- \xrightarrow{p} 1$, applying first Eq. (55) and next Eq. (56) in Lemma D.1 below yields

$$A_- = \sqrt{n}(B_{\bar{c}} - \hat{B})/\Omega_-^{1/2} + \text{cv}_\alpha(\sqrt{n}\hat{B}/\Omega_-^{1/2}) + o_p(1) = \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_-^{1/2}) + o_p(1),$$

where the $o_p(1)$ term is asymptotically negligible uniformly over \mathcal{S}_n . By similar argument, $A_+ = \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2}) + o_p(1)$. By Eq. (52), it therefore follows that, up to a term that's asymptotically negligible uniformly over \mathcal{S}_n , $P(h(\theta) \in \widetilde{CI})$ equals

$$P(Z + \sqrt{nb}/\Omega_-^{1/2} \leq \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_-^{1/2})) + P(Z - \sqrt{nb}/\Omega_+^{1/2} \leq \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2})) - 1, \quad (54)$$

where Z denotes as standard normal random variable. Fix $\epsilon > 0$. To conclude the proof, we will show that for n large enough, this expression is bounded below by $1 - \alpha - \epsilon$.

Since Eq. (54) is symmetric in Ω_+ and Ω_- , suppose without loss of generality that $\Omega_+ > \Omega_-$. By the assumption of the theorem, for n large enough and $\eta > 0$ specified below, $|\Omega_-^{1/2}/\Omega_+^{1/2} - 1| \leq \eta\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2}$.

We'll consider two cases, $\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2} > z_{1-\epsilon}$, and $\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2} \leq z_{1-\epsilon}$. Suppose first $\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2} > z_{1-\epsilon}$. Then, if $b < 0$, Eq. (54) is bounded below by $\Phi(\text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_-^{1/2})) + 1 - \alpha - 1 \geq \Phi(\text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2})) - \alpha \geq 1 - \epsilon - \alpha$, where the last inequality follows since $\text{cv}_\alpha(t) \geq t$. If b is positive, it is bounded below by $1 - \alpha + \Phi(\text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2})) - 1 \geq 1 - \alpha - \epsilon$.

Next, suppose, $\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2} \leq z_{1-\epsilon}$. Then $|\Omega_-^{1/2}/\Omega_+^{1/2} - 1| \leq \eta z_{1-\epsilon}$, so that by Eq. (56),

$$|\text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_-^{1/2}) - \sqrt{n}B_{\bar{c}}/\Omega_-^{1/2} - \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2}) + \sqrt{n}B_{\bar{c}}/\Omega_+^{1/2}| \leq \nu,$$

where $\nu = \eta z_{1-\epsilon}(z_{1-\alpha/2} - z_{1-\alpha})$. Therefore,

$$\begin{aligned} & P(Z + \sqrt{nb}/\Omega_-^{1/2} \leq \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_-^{1/2})) \\ & \geq P(Z + \sqrt{nb}/\Omega_-^{1/2} \leq \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2}) - \sqrt{n}B_{\bar{c}}/\Omega_+^{1/2} + \sqrt{n}B_{\bar{c}}/\Omega_-^{1/2} - \nu) \\ & = P\left(Z + \sqrt{nb}/\Omega_+^{1/2} \leq \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2}) + \sqrt{n}(B_{\bar{c}} - b)(\Omega_-^{-1/2} - \Omega_+^{-1/2}) - \nu\right) \\ & \geq P\left(Z + \sqrt{nb}/\Omega_+^{1/2} \leq \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2}) - \nu\right) \\ & \geq P\left(Z + \sqrt{nb}/\Omega_+^{1/2} \leq \text{cv}_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2})\right) + 1 - 2\Phi(\nu/2), \end{aligned}$$

where the last equality follows since $\inf_x \{\Phi(x - \nu) - \Phi(x)\} = 1 - 2\Phi(\nu/2)$. It therefore

follows that the coverage probability in Eq. (54) is bounded below by

$$P(Z + \sqrt{nb}/\Omega_+^{1/2} \leq cv_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2})) + P(Z - \sqrt{nb}/\Omega_+^{1/2} \leq cv_\alpha(\sqrt{n}B_{\bar{c}}/\Omega_+^{1/2})) - 1 + (1 - 2\Phi(\nu/2)) \geq 1 - \alpha + (1 - 2\Phi(\nu/2)).$$

where the inequality follows by definition of cv_α . Setting $\eta = 2z_{1/2+\epsilon/2}/(z_{1-\epsilon}(z_{1-\alpha/2} - z_{1-\alpha}))$ then implies that the right-hand side evaluates to $1 - \alpha - \epsilon$. Thus, Eq. (54) is bounded below by $1 - \alpha - \epsilon$, concluding the proof. \square

Lemma D.1. *The critical value $cv_\alpha(t)$ satisfies, for any $a > 0$,*

$$\sup_{b \geq 0} |cv_\alpha(ab)/a - cv_\alpha(b)| \leq z_{1-\alpha/2} \frac{|1-a|}{\max\{a, 1\}}, \quad (55)$$

and

$$\sup_{b \geq 0} |cv_\alpha(ab) - ab - cv_\alpha(b) + b| \leq (z_{1-\alpha/2} - z_{1-\alpha}) \frac{|1-a|}{\max\{a, 1\}}. \quad (56)$$

Proof. Since the function cv_α is increasing and convex, with slope bounded by 1, for $b_1, b_2 \geq 0$

$$cv_\alpha(b_1 + b_2) \leq b_1 + cv_\alpha(b_2), \quad (57)$$

and for $a \geq 1$ and $b \geq 0$,

$$cv_\alpha(ba)/a + cv_\alpha(0)(1 - 1/a) \geq cv_\alpha(b). \quad (58)$$

Suppose $a \geq 1$. Then by Eq. (57)

$$cv_\alpha(ab)/a - cv_\alpha(b) \leq (cv_\alpha(b) - b)(1/a - 1) \leq 0,$$

since $cv_\alpha(b) - b$ is bounded below by $z_{1-\alpha}$. On the other hand, by Eq. (58), the left-hand side is greater than $-cv_\alpha(0)(1 - 1/a)$. If $a \leq 1$, the same argument with ab and b reversed then yields Eq. (55).

To show Eq. (56), suppose first that $a \geq 1$. By Eq. (57), $cv_\alpha(ab) - ab - cv_\alpha(b) + b \leq 0$. On the other hand, by Eq. (58),

$$cv_\alpha(ab) - ab - cv_\alpha(b) + b \geq (cv_\alpha(ab) - ab - cv_\alpha(0))(1 - 1/a) \geq (1 - 1/a)(z_{1-\alpha} - z_{1-\alpha/2})$$

where the second inequality follows since $cv_\alpha(ab) - ab \geq z_{1-\alpha}$ and $cv_\alpha(0) = z_{1-\alpha/2}$. If $a < 1$, the same argument with ab and b reversed then yields Eq. (56). \square

References

- Andrews, D. W. K. and Guggenberger, P. (2009). Hybrid and size-corrected subsampling methods. *Econometrica*, 77(3):721–762.
- Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to sample statistics. *Quarterly Journal of Economics*, 132(4):1553–1592.
- Andrews, I., Gentzkow, M., and Shapiro, J. M. (2018). On the informativeness of descriptive statistics for structural estimates. Working Paper 25217, National Bureau of Economic Research.
- Armstrong, T. B. and Kolesár, M. (2018). Optimal inference in a class of regression models. *Econometrica*, 86(2):655–683.
- Armstrong, T. B. and Kolesár, M. (2020). Sensitivity analysis using approximate moment condition models. ArXiv: 1808.07387v4.
- Berkowitz, D., Caner, M., and Fang, Y. (2012). The validity of instruments revisited. *Journal of Econometrics*, 166(2):255–266.
- Berry, S. T., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.
- Bonhomme, S. and Weidner, M. (2020). Minimizing sensitivity to model misspecification. ArXiv: 1807.02161.
- Brinch, C. N., Mogstad, M., and Wiswall, M. (2017). Beyond LATE with a Discrete Instrument. *Journal of Political Economy*, 125(4):985–1039.
- Bugni, F. A. and Ura, T. (2019). Inference in dynamic discrete choice problems under local misspecification. *Quantitative Economics*, 10(1):67–103.
- Cai, T. T. and Low, M. G. (2004). An adaptation theory for nonparametric confidence intervals. *The Annals of Statistics*, 32(5):1805–1840.
- Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics*, 34(3):305–334.
- Conley, T. G., Hansen, C. B., and Rossi, P. E. (2012). Plausibly exogenous. *The Review of Economics and Statistics*, 94(1):260–272.

- Cressie, N. and Read, T. R. C. (1984). Multinomial goodness-of-fit tests. *Journal of the Royal Statistical Society. Series B (Methodological)*, 46(3):440–464.
- DiTraglia, F. J. (2016). Using invalid instruments on purpose: Focused moment selection and averaging for GMM. *Journal of Econometrics*, 195(2):187–208.
- Donoho, D. L. (1994). Statistical estimation and optimal recovery. *The Annals of Statistics*, 22(1):238–270.
- Dubé, J.-P., Fox, J. T., and Su, C.-L. (2012). Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica*, 80(5):2231–2267.
- Dufo, E., Greenstone, M., Pande, R., and Ryan, N. (2018). The value of regulatory discretion: Estimates from environmental inspections in India. *Econometrica*, 86(6):2123–2160.
- Efron, B., Hastie, T., Johnstone, I. M., and Tibshirani, R. J. (2004). Least angle regression. *The Annals of Statistics*, 32(2):407–451.
- Freyberger, J. and Horowitz, J. L. (2015). Identification and shape restrictions in nonparametric instrumental variables estimation. *Journal of Econometrics*, 189(1):41–53.
- Gayle, G.-L. and Shephard, A. (2019). Optimal taxation, marriage, home production, and family labor supply. *Econometrica*, 87(1):291–326.
- Guggenberger, P. (2012). On the asymptotic size distortion of tests when instruments locally violate the exogeneity assumption. *Econometric Theory*, 28(2):387–421.
- Hahn, J. and Hausman, J. A. (2005). IV estimation with valid and invalid instruments. *Annales d'Économie et de Statistique*, (79/80):25–57.
- Hall, A. R. and Inoue, A. (2003). The large sample behaviour of the generalized method of moments estimator in misspecified models. *Journal of Econometrics*, 114(2):361–394.
- Hansen, L. P. (1985). A method for calculating bounds on the asymptotic covariance matrices of generalized method of moments estimators. *Journal of Econometrics*, 30(1-2):203–238.
- Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press, Princeton, NJ.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.

- Heckman, J. J. and Vytlacil, E. (2005). Structural Equations, Treatment Effects, and Econometric Policy Evaluation1. *Econometrica*, 73(3):669–738.
- Hong, H., Mahajan, A., and Nekipelov, D. (2015). Extremum estimation and numerical derivatives. *Journal of Econometrics*, 188(1):250–263.
- Huber, P. J. and Ronchetti, E. M. (2009). *Robust Statistics*. John Wiley & Sons, Hoboken, NJ, second edition.
- Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.
- Imbens, G. W. and Manski, C. F. (2004). Confidence intervals for partially identified parameters. *Econometrica*, 72(6):1845–1857.
- Joshi, V. M. (1969). Admissibility of the usual confidence sets for the mean of a univariate or bivariate normal population. *The Annals of Mathematical Statistics*, 40(3):1042–1067.
- Kitamura, Y., Otsu, T., and Evdokimov, K. (2013). Robustness, infinitesimal neighborhoods, and moment restrictions. *Econometrica*, 81(3):1185–1201.
- Kowalski, A. E. (2016). Doing more when you’re running LATE: Applying marginal treatment effect methods to examine treatment effect heterogeneity in experiments. Working Paper 22363, National Bureau of Economic Research.
- Lehmann, E. L. and Romano, J. P. (2005). *Testing Statistical Hypotheses*. Springer, New York, NY, third edition.
- Low, M. G. (1995). Bias-variance tradeoffs in functional estimation problems. *The Annals of Statistics*, 23(3):824–835.
- McCloskey, A. (2019). Asymptotically uniform tests after consistent model selection in the linear regression model. *Journal of Business & Economic Statistics*, forthcoming.
- Mogstad, M., Santos, A., and Torgovitsky, A. (2018). Using instrumental variables for inference about policy relevant treatment effects. *Econometrica*, 86(5):1589–1619.
- Moon, H. R. and Schorfheide, F. (2009). Estimation with overidentifying inequality moment conditions. *Journal of Econometrics*, 153(2):136–154.
- Mukhin, Y. (2018). Sensitivity of regular estimators. ArXiv: 1805.08883.

- Newey, W. K. (1985). Generalized method of moments specification testing. *Journal of Econometrics*, 29(3):229–256.
- Newey, W. K. (1990). Semiparametric efficiency bounds. *Journal of Applied Econometrics*, 5(2):99–135.
- Newey, W. K. and McFadden, D. L. (1994). Large sample estimation and hypothesis testing. In McFadden, D. L. and Engle, R. F., editors, *Handbook of Econometrics*, volume 4, chapter 36, pages 2111–2245. Elsevier.
- Pratt, J. W. (1961). Length of confidence intervals. *Journal of the American Statistical Association*, 56(295):549–567.
- Rambachan, A. and Roth, J. (2019). An honest approach to parallel trends. Unpublished manuscript, Harvard University.
- Rosset, S. and Zhu, J. (2007). Piecewise linear regularized solution paths. *The Annals of Statistics*, 35(3):1012–1030.
- Sacks, J. and Ylvisaker, D. (1978). Linear estimation for approximately linear models. *The Annals of Statistics*, 6(5):1122–1137.
- Sun, Y., Baricz, Á., and Zhou, S. (2010). On the monotonicity, log-concavity, and tight bounds of the generalized Marcum and Nuttall Q -functions. *IEEE Transactions on Information Theory*, 56(3):1166–1186.
- Tsybakov, A. B. (2009). *Introduction to nonparametric estimation*. Springer, New York, NY.
- van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press, New York, NY.

Table 1: J -test of overidentifying restrictions in the application to [Berry et al. \(1995\)](#) under different forms of ℓ_p misspecification.

Instrument set	$p = 2$	$p = \infty$
D/F: # cars	10.21	10.21
S/F: # cars	15.00	15.00
Supply: Miles/dollar	16.31	16.31
All D/F	2.71	2.71
All D/R	5.36	5.55
All S/F	2.54	2.56
All S/R	4.06	6.84
All excluded demand	1.80	1.97
All excluded supply	1.60	1.72
All excluded	1.13	2.56

Notes: The table gives the minimum value of M such that the test of overidentifying restrictions has p -value equal to 0.05. “D/F”: Demand-side instrument based on characteristics of other cars produced by the same firm. “S/F”: Supply-side instrument based on characteristics of other cars produced by the same firm. “D/R”: Demand-side instrument based on characteristics of cars produced by rivals. “S/R”: Supply-side instrument based on characteristics of cars produced by rivals. “All excluded”: All excluded instruments are potentially invalid.

Table 2: Efficiency bounds (in %) for one and two-sided 95% confidence intervals at $c = 0$ under ℓ_p misspecification in the application to [Berry et al. \(1995\)](#).

Instrument set	Two-sided		One-sided	
	$p = 2$	$p = \infty$	$p = 2$	$p = \infty$
D/F: # cars	85.9	85.9	100.0	100.0
S/F: # cars	90.1	90.1	99.8	99.8
Supply: Miles/dollar	85.0	85.0	100.0	100.0
All D/F	85.5	85.7	100.0	100.0
All D/R	94.8	95.3	93.9	95.3
All S/F	88.6	89.1	99.7	99.7
All S/R	89.4	89.2	98.5	99.5
All excluded demand	95.4	96.4	95.0	97.3
All excluded supply	90.3	90.1	98.2	99.6
All excluded	97.0	97.5	99.5	98.2

Notes: For two-sided confidence intervals, the table calculates the ratio of the expected length of a 95% confidence interval that minimizes its length at $c = 0$ relative to the length of the CI in (16), given in (18). For one-sided confidence intervals, the table calculates an analogous bound, given in Appendix C.6, when the confidence interval optimizes the 80% quantile of excess length. Instrument set labels are describe in notes to Table 1.

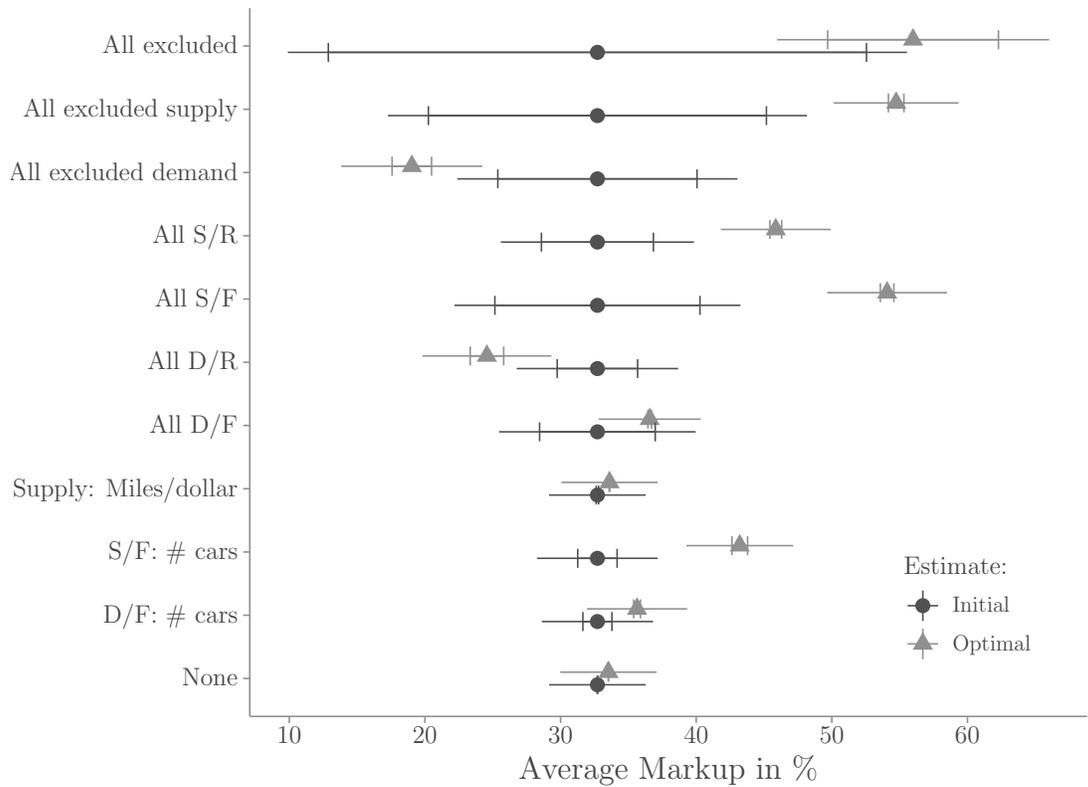


Figure 1: Confidence intervals under ℓ_2 misspecification and $M = 1$ in the application to [Berry et al. \(1995\)](#).

Vertical lines correspond to the estimate \pm the worst case bias, and horizontal lines correspond to 95% confidence intervals. Different rows correspond to assuming that different subsets of instruments are potentially invalid. “None”: correct specification. “D/F”: Demand-side instrument based on characteristics of other cars produced by the same firm. “S/F”: Supply-side instrument based on characteristics of other cars produced by the same firm. “D/R”: Demand-side instrument based on characteristics of cars produced by rivals. “S/R”: Supply-side instrument based on characteristics of cars produced by rivals. “All excluded”: All excluded instruments are potentially invalid.

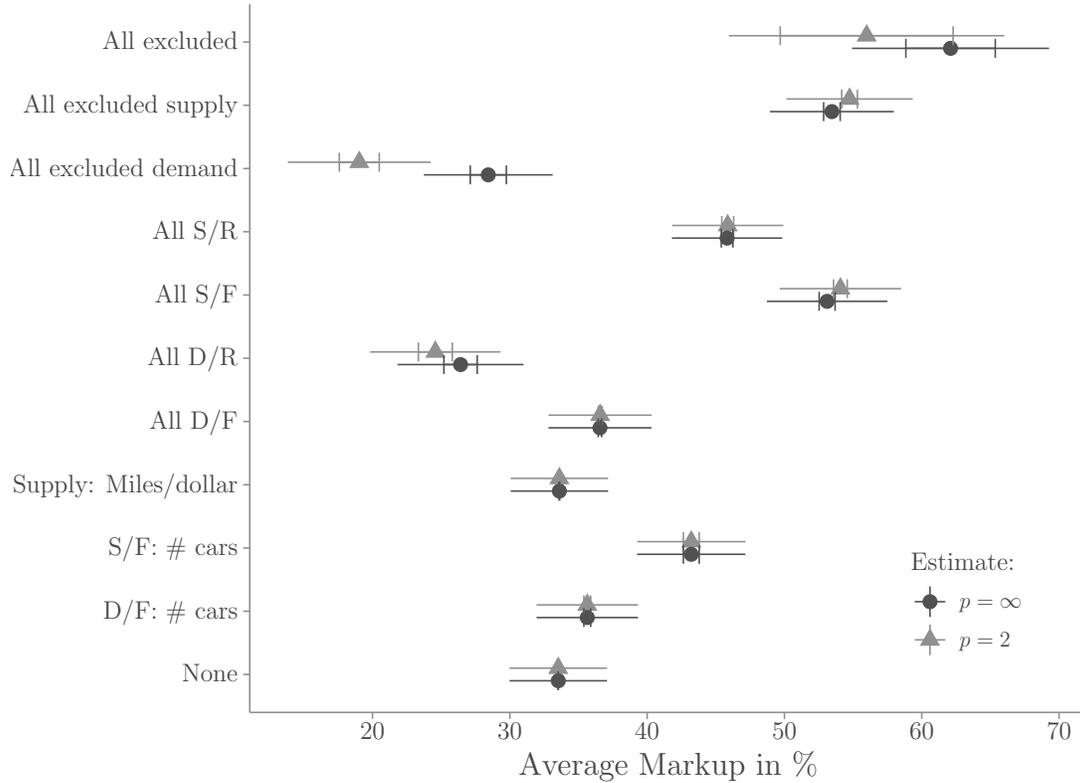


Figure 2: Optimal Confidence intervals under ℓ_2 , and ℓ_∞ misspecification and $M = 1$ in the application to [Berry et al. \(1995\)](#).

Vertical lines correspond to the estimate \pm the worst case bias, and horizontal lines correspond to 95% confidence intervals. Different rows correspond to assuming that different subsets of instruments are potentially invalid. “None”: correct specification. “D/F”: Demand-side instrument based on characteristics of other cars produced by the same firm. “S/F”: Supply-side instrument based on characteristics of other cars produced by the same firm. “D/R”: Demand-side instrument based on characteristics of cars produced by rivals. “S/R”: Supply-side instrument based on characteristics of cars produced by rivals. “All excluded”: All excluded instruments are potentially invalid.

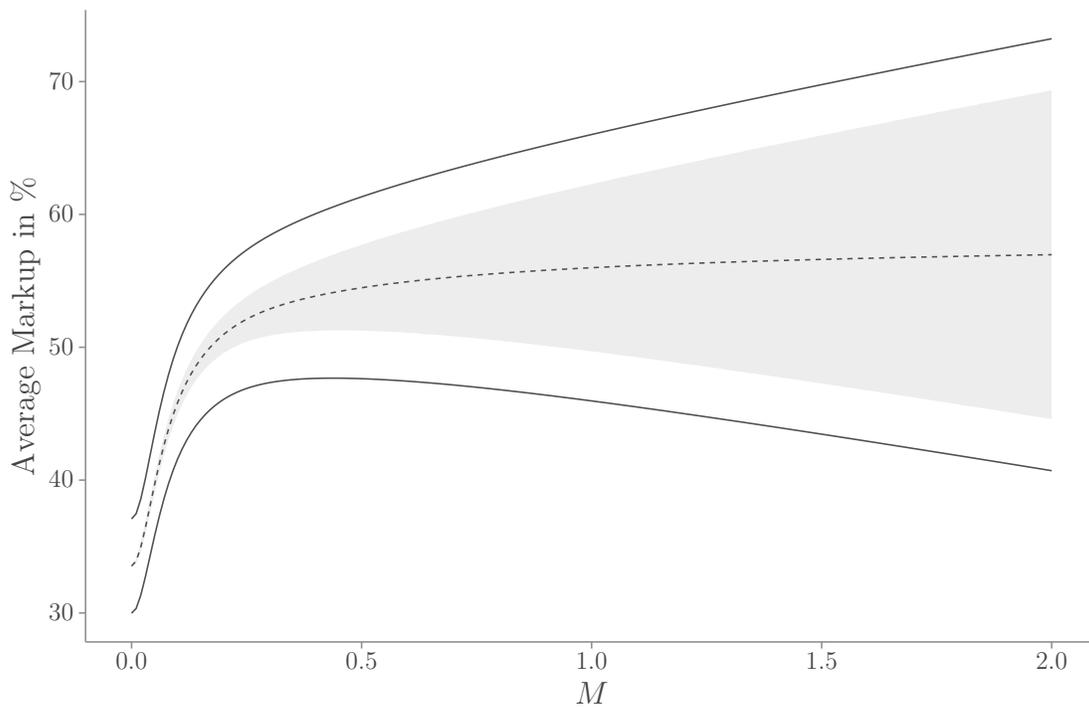


Figure 3: Optimal confidence intervals under ℓ_2 misspecification the application to [Berry et al. \(1995\)](#) as a function of misspecification parameter M , when all excluded instruments are allowed to be potentially invalid.

Dotted line corresponds to point estimate, shaded region denotes the estimate \pm its worst-case bias, and a 95% confidence band is denoted by solid lines.