

Inference for Linear Conditional Moment Inequalities*

Isaiah Andrews Jonathan Roth Ariel Pakes

September 17, 2019

Abstract

We consider inference based on linear conditional moment inequalities, which arise in a wide variety of economic applications, including many structural models. We show that linear conditional structure greatly simplifies confidence set construction, allowing for computationally tractable projection inference in settings with nuisance parameters. Next, we derive least favorable critical values that avoid conservativeness due to projection. Finally, we introduce a conditional inference approach which ensures a strong form of insensitivity to slack moments, as well as a hybrid technique which combines the least favorable and conditional methods. Our conditional and hybrid approaches are new even in settings without nuisance parameters. We find good performance in simulations based on Wollmann (2018), especially for the hybrid approach.

Keywords: Moment Inequalities, Subvector Inference, Uniform Inference

JEL Codes: C12

1 Introduction

Moment inequalities are an important tool in empirical economics, enabling researchers to use the most direct implications of utility or profit maximization for inference in

*We thank Gary Chamberlain, Ivan Canay, Kirill Evdokimov, Jerry Hausman, Bulat Gafarov, Hiroaki Kaido, Adam McCloskey, Francesca Molinari, Whitney Newey, Ashesh Rambachan, Jesse Shapiro, Brit Shari, Xiaoxia Shi, Joerg Stoye, and participants at several seminars for helpful comments, and thank Thomas Wollmann for helpful discussion of his application. Andrews gratefully acknowledges financial support from the NSF under Grant 1654234. Roth gratefully acknowledges financial support from an NSF Graduate Research Fellowship under Grant DGE1144152. Andrews: iandrews@fas.harvard.edu. Roth: jonathanroth@g.harvard.edu. Pakes: apakes@fas.harvard.edu

both single-agent settings and games. Moment inequalities have also been used to weaken parametric, behavioral, measurement, and selection assumptions in a range of problems.¹

Inference based on moment inequalities raises a number of challenges. First, calculating tests and confidence sets can be computationally taxing in settings with more than a few nuisance parameters (for instance, coefficients on control variables). Second, a simple approach to inference in settings with nuisance parameters is to use projection, but this can yield imprecise results. Finally, it is often unclear ex-ante which of the many moments implied by an economic model will be informative, and inclusion of uninformative or slack moments yields wide confidence sets for some procedures.

This paper proposes new methods which address these three implementation challenges for an important class of moment inequalities, which we term *linear conditional moment inequalities*. These are conditional moment inequalities that (a) are linear in nuisance parameters and (b) have conditional variance (given the instruments) that does not depend on the nuisance parameters. Such inequalities arise naturally when the nuisance parameters enter the moments linearly and interact only with exogenous variables. This occurs, for example, in regression and instrumental variables settings with interval-valued outcomes and exogenous controls. Linear conditional structure also appears in a number of structural applications of moment inequalities in the empirical literature. We next discuss how linear conditional structure allows us to overcome the challenges discussed above and construct powerful, tractable inference procedures.

The first challenge discussed above, computational burden, often arises from the

¹For recent overviews of research involving moment inequalities, and partial identification more broadly, see Ho & Rosen (2017) and Molinari (2019). For the behavioral and measurement assumptions underlying the use of moment inequalities in problems where agents are assumed to maximize utility or profit see Pakes (2010) and Pakes et al. (2015). For examples of inequalities generated by first order conditions see Dickstein & Morales (2018) on export decisions and Holmes (2011) on Walmart’s location decisions. For examples of inequalities generated by Nash equilibrium conditions see Ciliberto & Tamer (2009), Eizenberg (2014), or Wollman (2018) on entry and exit decisions. For examples of the use of inequalities to weaken assumptions see Haile & Tamer (2003) on auctions, Chetty (2012) on labor supply, and Kline & Tartari (2016) on a welfare reform experiment. For moment inequalities used to overcome measurement problems see Manski & Tamer (2002) on interval-valued outcome variables and Ho & Pakes (2014) on errors in regressors in discrete choice models. For the use of inequalities to overcome selection problems see Blundell et al. (2007) on changes in inequality and Kreider et al. (2012) on take-up of SNAP. There is also closely related work in other fields, for example on computation of bounds for competing risk models (e.g. Honore & Lleras-Muney (2006) on the war on cancer).

use of test inversion. A common approach evaluates a test statistic on a grid of parameter values and defines the confidence set as the set where the statistic falls below a parameter-specific critical value. The computational cost of this approach scales with the number of grid points, which typically grows exponentially in the dimension of the parameter vector. Hence, these methods become very difficult to apply in problems with more than a few parameters. We show that in settings with linear conditional moment inequalities the nuisance parameters can be eliminated by solving a simple linear program, so it suffices to specify a grid for the parameters of interest. This allows us to easily compute confidence sets for the parameters of interest even in cases where the dimension of the nuisance parameters renders traditional grid-based techniques impractical.²

The second challenge discussed above stems from the fact that many existing techniques deliver joint confidence sets for all parameters entering the moment inequalities, which must then be projected to obtain confidence sets for lower-dimensional parameters of interest. For examples of projection in the theoretical and empirical literature, see Canay & Shaikh (2017). As discussed by Bugni et al. (2017) and Kaido et al. (2019a), however, projection can yield very conservative tests and confidence sets. We show that in settings with linear conditional moment inequalities, it is straightforward to derive computationally tractable least favorable critical values that account for the presence of nuisance parameters, and so construct non-conservative confidence sets for the parameters of interest.³

The final challenge discussed above, sensitivity to slack moments, arises from the fact that the distribution of moment inequality test statistics depends on the (unknown) degree to which the moments are slack. As discussed by D. Andrews & Soares (2010), the degree of slackness cannot be uniformly consistently estimated, so the least favorable approach calculates critical values under the worst-case assumption that all moments bind. The resulting procedures may have low power when this assumption is false and many moments are slack. We show that in settings with linear conditional moment inequalities, one can derive tests that condition on the set of binding moments in the data. Conditional tests are simple to implement and insensitive to slack

²Note, however, that our asymptotic results (developed in the appendix) hold the number of parameters and moments fixed. Hence, our analysis does not address settings that are “high-dimensional” in the sense that the number of parameters or moments grows with the sample size.

³In cases where some nuisance parameters enter the moments nonlinearly, these techniques deliver confidence sets for the parameters of interest together with the nonlinear nuisance parameters.

moments in the strong sense that, as a subset of the moments becomes arbitrarily slack, the conditional test converges to the test that drops these moments ex-ante. Unlike the approach of e.g. D. Andrews & Soares (2010), conditional tests achieve this insensitivity without a sequence of sample size-dependent tuning parameters. To improve power in cases where conditional tests underperform, we further introduce hybrid tests that combine least favorable and conditional techniques. Tests based on a similar hybridization are used by Andrews et al. (2018) for inference following a data-driven choice of a target parameter.

For simplicity of exposition, the main text develops our results in a finite sample normal model motivated as an asymptotic approximation. In the supplement, we translate these finite sample results to uniform asymptotic results. We show that our least favorable approach (with its critical value increased by an infinitesimal uniformity factor as in D. Andrews & Shi (2013)) is uniformly asymptotically valid under minimal conditions. Under additional conditions, which still allow for any combination of binding and nonbinding moments in the population, we show uniformity for the least favorable approach without the infinitesimal uniformity factor, and for versions of the conditional and hybrid approaches which do not reject when the moments are far from binding.

To explore the numerical performance of our methods, we apply our techniques in simulations calibrated to Wollman (2018)’s study of the US auto bailout. We consider designs with up to ten nuisance parameters, and find that our approach remains tractable throughout. We find substantial power improvements for our least favorable critical values relative to the projection method. We find further improvements for our conditional approach at most parameter values. Finally, we find that our hybrid approach performs well, with power never substantially below and often exceeding the other procedures considered. Hence, we recommend the hybrid approach.

Related Literature Uniform inference on subsets of parameters based on linear moment inequalities was previously studied by Cho & Russell (2019) and Gafarov (2019). Flynn (2019) further allows for the possibility of a continuum of linear moments. Unlike our approach these papers consider unconditional moment inequalities, but do not discuss the case where the parameters of interest may enter the moments nonlinearly. Hsieh et al. (2017) propose a conservative form of projection inference for settings which include linear unconditional moment inequalities. Kaido et al.

(2019a) develop techniques for eliminating projection conservativeness, while Bugni et al. (2017) develop an alternative approach for inference on subsets of parameters, and Belloni et al. (2018) build on this approach to develop results for subset inference with high-dimensional unconditional moments. All three techniques are more widely applicable than those we develop, requiring neither linearity nor conditional moment inequalities. At the same time, all can be computationally intensive in settings with a large number of nuisance parameters.⁴ Chernozhukov et al. (2015) develop techniques for subset inference based on conditional moment inequalities, which unlike our approach do not require linearity. Romano & Shaikh (2008) discuss subvector inference based on subsampling. Chen et al. (2018) discuss confidence sets for the identified set for subvectors based on a quasi-posterior Monte Carlo approach.

Finally, there is a large literature on techniques which seek to reduce sensitivity to the inclusion of slack moments in settings without nuisance parameters, including D. Andrews & Soares (2010), D. Andrews & Barwick (2012), Romano et al. (2014a), and Cox & Shi (2019). Chernozhukov et al. (2015), Bugni et al. (2017), Belloni et al. (2018), and Kaido et al. (2019a) build on related ideas to reduce sensitivity to slack moments in models with nuisance parameters. If applied in our setting, however, these techniques would eliminate the linear structure which simplifies computation. Even in settings without nuisance parameters our conditioning approach appears to be new, and a small set of simulations without nuisance parameters (described in Appendix F) finds our hybrid approach neither dominates nor is dominated by the test proposed by Romano et al. (2014a).

Preview of Paper The next section introduces our linear conditional setting. Section 3 develops a conditional asymptotic approximation that motivates our analysis, and discusses the relationship between our approach and the literature on conditional moment inequalities. Section 4 introduces projection and least favorable tests, while Section 5 introduces conditional and hybrid tests. Section 6 discusses the practical details of implementing our approach, while Section 7 reports Monte Carlo results. A reader looking to apply our methods but not interested in the theory can skip from Section 3 to Section 6. Additional technical results are stated in Appendices A and B, while proofs for all results in the main text are provided in Appendix C. Uniform

⁴Kaido et al. (2019a) propose the use of a response surface technique to facilitate computation, and find that it yields substantial improvements. See Kaido et al. (2019a) and Gafarov (2019) for further evidence on computational performance.

asymptotic results are stated Appendix D and proved in Appendix E. Results from a small simulation study without nuisance parameters are reported in Appendix F, while additional details and results for the simulations in the main text are reported in Appendix G.

2 Linear Conditional Moment Inequalities

Throughout the paper, we assume that we observe independent and identically distributed data D_i , $i = 1, \dots, n$ drawn from a distribution P . We are interested in parameters identified by k -dimensional conditional moment inequalities

$$E_P [g(D_i, \beta, \delta) | Z_i] \leq 0 \text{ almost surely} \quad (1)$$

assumed to hold at the true parameter value, for $g(D_i, \beta, \delta)$ a known function of the data and parameters. Going forward we leave the “almost surely” implicit for brevity. We seek tests and confidence sets for β , while the p -dimensional vector δ is a nuisance parameter. Formally, we want to test the null that a given value β_0 belongs to the identified set, $\tilde{H}_0 : \beta_0 \in B_I(P)$, where

$$B_I(P) = \{\beta : \text{there exists } \delta \text{ such that } E_P [g(D_i, \beta, \delta) | Z_i] \leq 0\}$$

is the set of all values β such that there exists δ which makes (1) hold.

We assume that the moment function $g(D_i, \beta, \delta)$ is of the form

$$g(D_i, \beta, \delta) = g(D_i, \beta, 0) - X(Z_i, \beta) \delta \quad (2)$$

for some $k \times p$ matrix-valued function $X(Z_i, \beta)$ of the instruments and the parameter of interest β . This imposes two key restrictions. First, (2) requires that the nuisance parameter δ enter the moments linearly. Since linear models are widely used in economics, this holds in a wide variety of applications. Second, (2) requires that the derivative of the moments with respect to δ be non-random conditional on the instruments Z_i . Stated differently, we require that the moment inequalities (1) hold conditional on the Jacobian of the moments with respect to δ . This implies that

$$Var_P(g(D_i, \beta, \delta) | Z_i) = Var_P(g(D_i, \beta, 0) | Z_i),$$

so the conditional variance of the moments does not depend on δ . This condition plays a crucial role in the asymptotic approximation developed in Section 3 below.

We call moment inequalities satisfying (1) and (2) *linear conditional moment inequalities*. They can be understood as a generalization of the linear model with exogenous regressors and outcome Y_i^* ,

$$Y_i^* = X_i' \delta + \varepsilon_i \text{ where } E_P [\varepsilon_i | X_i] = 0, \quad (3)$$

to the moment inequality setting. Specifically, for linear conditional moment inequalities we can define

$$(Y_i, X_i) = (g(D_i, \beta_0, 0), X(Z_i, \beta_0)) \quad (4)$$

for β_0 again the null value of β . If $\beta_0 \in B_I(P)$, then we can write

$$Y_i = X_i \delta + \varepsilon_i \text{ where } E_P [\varepsilon_i | Z_i] \leq 0. \quad (5)$$

Thus, the linear conditional moment inequality model resembles a generalization of the traditional linear regression model, where we (a) allow the possibility that there are instruments Z_i beyond the regressors X_i and (b) relax the conditional moment restriction on the errors ε_i to an inequality. We show below that the restriction to linear conditional moment inequalities yields important simplifications in the problem of testing $\tilde{H}_0 : \beta_0 \in B_I(P)$. Before developing these results, we motivate our study of linear conditional moment inequalities by showing that moment inequalities of this form arise in a variety of economic examples.

Example 1 Linear conditional moment inequalities arise naturally from the linear regression model (3), and its instrumental variables generalization, when we only observe bounds on the outcome Y_i^* . Consider the model

$$Y_i^* = T_i \beta + V_i' \delta + \varepsilon_i, \quad E_P [\varepsilon_i | Z_i] = 0$$

where V_i is exogenous in the sense that it is a function of Z_i , while T_i may be endogenous. For instance, β may be a causal effect of interest, whereas V_i represents a set of control variables. This is a linear instrumental variables model where the error is mean-independent of the instrument.

As in e.g. Manski & Tamer (2002), suppose that rather than observing Y_i^* , we

instead observe bounds Y_i^L and Y_i^U where $Y_i^L \leq Y_i^* \leq Y_i^U$ with probability one. The linear model (2) implies that $E[Y_i^L - T_i\beta - V_i'\delta|Z_i] \leq 0$ and $E[T_i\beta + V_i'\delta - Y_i^U|Z_i] \leq 0$, so we obtain conditional moment inequalities. To cast these inequalities into our framework, suppose we are interested in inference on β , and for any vector of non-negative functions of the instruments $f(Z_i)$ let $Y_i(\beta) = (Y_i^L - T_i\beta, T_i\beta - Y_i^U)' \otimes f(Z_i)$, and $X_i = (V_i \otimes (1, 1)') \otimes f(Z_i)$, for “ $\otimes$ ” the Kroneker product. This yields the moments $E[Y_i(\beta) - X_i\delta|Z_i] \leq 0$, as desired.⁵ $\triangle$

Example 2 Katz (2007) studies the impact of travel time on supermarket choice. Katz assumes that agent utilities are additively separable in utility from the basket of goods bought (B_i), the travel time to the supermarkets chosen ($T_{i,s}$), and the cost of the basket ($\pi(B_i, s)$). Normalizing coefficient on cost to one, agent i ’s realized utility is assumed to be

$$U_i(B_i, s) = U_i(B_i) + C'_s\delta - (\beta + \nu_i)T_{i,s} - \pi(B_i, s),$$

where C_s are observed characteristics of the supermarket, $T_{i,s}$ is the travel time for i going to s , and $\beta + \nu_i$ is its impact on utility, where ν_i has mean zero given supermarket characteristics and travel times.

Katz assumes travel times and store characteristics are known to the shopper. For $\tilde{s}$ a supermarket with $T_{i,\tilde{s}} > T_{i,s}$ that also marketed B_i , he divides the difference $U_i(B_i, s) - U_i(B_i, \tilde{s})$ by $T_{i,s} - T_{i,\tilde{s}}$ and notes that a combination of expected utility maximization and revealed preference implies that $E[Y_i(\beta) - X_i\delta|Z_i] \leq 0$, for

$$Y_i(\beta) \equiv -\beta - \frac{[\pi(B_i, s) - \pi(B_i, \tilde{s})]}{T_{i,s} - T_{i,\tilde{s}}}, \quad X_i \equiv -\frac{C_s - C_{\tilde{s}}}{T_{i,s} - T_{i,\tilde{s}}}, \quad \text{and } Z_i \equiv (T_{i,s}, T_{i,\tilde{s}}, C_s, C_{\tilde{s}})'.$$

By adding an analogous inequality which uses a store closer to the agent, Katz obtains both upper and lower bounds for β .

A similar approach can be used in any ordered choice problem, including those with interacting agents; see Pakes et al. (2015), who also provide a way to handle the

⁵Our approach to this application relies on the conditional moment restriction $E_P[\varepsilon_i|Z_i] = 0$. As discussed by Ponomareva & Tamer (2011), this means that the identified set may be empty if the linear model is incorrect. For $Z_i = (T_i, V_i)'$, Beresteanu & Molinari (2008) assume only that $E[\varepsilon_i Z_i] = 0$, and their approach yields inference on the (necessarily nonempty) set of best linear predictors. Bontemps et al. (2012) study identification and inference, including specification tests, for a class of linear models with unconditional moment restrictions.

boundaries of the choice set (as would occur in Katz’s case if there were no closer supermarket for some observations). $\triangle$

Example 3 Wollman (2018) considers the bailout of GM and Chrysler’s commercial truck divisions during the 2008 financial crisis and asks what would have happened had they instead been allowed to either fail or merge with another firm. This example is the basis for our simulations below.

Merger analysis focuses on price differences pre- and post-merger. Wollmann notes that some commercial truck production is modular (it is possible to connect different cab types to different trailers), so some products would likely have been repositioned after the change in the environment. To analyze product repositioning he requires estimates for the fixed costs of marketing a product. His estimated demand and cost systems enable him to estimate counterfactual profits from adding or deleting products. Assuming firms maximize expected profits, differences in the expected profits from adding or subtracting products imply bounds on fixed costs.

To illustrate, let $J_{f,t}$ be the set of models that firm f marketed in year t and let $J_{f,t}/j$ be that set excluding product j , while $\Delta\pi(J_{f,t}, J_{f,t}/j)$ is the difference in expected profits between marketing $J_{f,t}$ and $J_{f,t}/j$. Denote the fixed cost to firm f of marketing product j at time t by $X_{j,f,t}(\beta)\delta$ where the X ’s are product characteristics and β is a scalar which differentiates between marketing costs for products that were and were not marketed in the prior year. Then if $Z_{f,t}$ represents a set of variables known to the the firm when marketing decisions were made (which includes the variables used to form $X_{j,f,t}(\beta)$), the equilibrium condition ensures that

$$E[Y_{j,f,t} - X_{j,f,t}(\beta)\delta | Z_{f,t}] \geq 0 \text{ for all } j,$$

where

$$Y_{j,f,t} \equiv \Delta\pi(J_{f,t}, J_{f,t}/j) \cdot 1\{j \in J_{f,t}, j \in J_{f,t-1}\}, \quad X_{j,f,t}(\beta) \equiv X_{f,j}(\beta) \cdot 1\{j \in J_{f,t}, j \in J_{f,t-1}\}$$

and $1\{A\}$ is an indicator for the event A . Additional inequalities can be added for marketing a product that was not marketed in the prior period, for withdrawing products, and for combining the withdrawal of one product with adding another. See Section 7 below for details. $\triangle$

Other recent applications that use linear conditional moment inequalities include

Ho & Pakes (2014), who study the effect of physician incentives on hospital referrals, and Morales et al. (2019), who develop and estimate an extended gravity model of trade flows. As the variety of examples illustrates, linear conditional moment inequalities arise in a range of economic contexts.

3 Conditional Asymptotics

In this section we derive a normal asymptotic approximation that motivates the procedures developed in the rest of the paper. For $(Y_i, X_i) = (g(D_i, \beta_0, 0), X(Z_i, \beta_0))$ as in (4), recall that we can write the moments evaluated at β_0 as $g(D_i, \beta_0, \delta) = Y_i - X_i\delta$. We consider procedures that test $\tilde{H}_0 : \beta_0 \in B_I(P)$ based on the scaled sample average of the moments evaluated at β_0 ,

$$g_n(\beta_0, \delta) = \frac{1}{\sqrt{n}} \sum_i g(D_i, \beta_0, \delta) = Y_n - X_n\delta,$$

for $Y_n = \frac{1}{\sqrt{n}} \sum_i Y_i$ and $X_n = \frac{1}{\sqrt{n}} \sum_i X_i$. As in Bugni et al. (2017), we will form confidence sets for β by testing a grid of values β_0 . Hence, for the moment we fix a null value β_0 and suppress dependence on β_0 in our notation, deferring further discussion of test inversion to Section 6 below.

Similar to Abadie et al. (2014) we consider asymptotic approximations that condition on the instruments $\{Z_i\} = \{Z_i\}_{i=1}^\infty$. If we define

$$\mu_i = \mu(Z_i) = E_P[Y_i|Z_i]$$

as the conditional mean of Y_i given Z_i , and $\mu_n = \frac{1}{\sqrt{n}} \sum_i \mu_i$ as the scaled sample average of μ_i , then under $\tilde{H}_0 : \beta_0 \in B_I(P)$ there exists a value δ such that $\mu_n - X_n\delta \leq 0$ (for almost every $\{Z_i\}$). Since μ_n and X_n are nonrandom once we condition on $\{Z_i\}$, to test $\beta_0 \in B_I(P)$ we will test the implied hypothesis $H_0 : \mu_n \in \mathcal{M}_0$ for

$$\mathcal{M}_0 = \{\mu_n : \text{There exists } \delta \text{ such that } \mu_n - X_n\delta \leq 0\}. \quad (6)$$

Note that $\beta_0 \in B_I(P)$ implies that $\mu_n \in \mathcal{M}_0$ for almost every $\{Z_i\}$, so tests of $H_0 : \mu_n \in \mathcal{M}_0$ with correct size also control size as tests of $\tilde{H}_0 : \beta_0 \in B_I(P)$. Note further that $H_0 : \mu_n \in \mathcal{M}_0$ holds trivially conditional on $\{Z_i\}$ if the column span of X_n contains a strictly negative vector. Hence, going forward we assume that $X_n\delta$ has

at least one non-negative element for all δ .

To derive asymptotic approximations useful for testing H_0 , note that $Y_n - \mu_n$ has mean zero conditional on $\{Z_i\}$ by construction. Thus, under mild conditions we can apply the central limit theorem conditional on $\{Z_i\}$.

Lemma 1 (*Lindeberg-Feller*) Suppose that as $n \rightarrow \infty$, conditional on $\{Z_i\}$ we have

$$\frac{1}{n} \sum_i E_P [Y_i Y_i' | Z_i] 1 \left\{ \frac{1}{\sqrt{n}} \|Y_i\| > \varepsilon \right\} \rightarrow 0 \text{ for all } \varepsilon > 0,$$

$$\frac{1}{n} \sum_i \text{Var}_P (Y_i | Z_i) \rightarrow \Sigma = E_P [\text{Var}_P (Y_i | Z_i)].$$

Then $Y_n - \mu_n \rightarrow_d N(0, \Sigma)$.

The first condition of Lemma 1 requires that the average of Y_i given Z_i not be dominated by a small number of large observations, while the second requires that the average conditional variance converge.

Under these conditions, Lemma 1 suggests the normal approximation

$$Y_n - X_n \delta | \{Z_i\} \approx^d N(\mu_n - X_n \delta, \Sigma), \quad (7)$$

where we use $\approx^d$ to denote approximate equality in distribution, and we have used that X_n is non-random conditional on $\{Z_i\}$ to put it on the right hand side in (7). In the next three sections we assume this approximation holds exactly for known Σ and derive finite-sample results. We return to the issue of approximation error in Appendix D. There, we show that we can consistently estimate Σ , and that the finite-sample properties of our procedures in the normal model translate to uniform asymptotic properties over large classes of data generating processes.

Choice of Moments Our asymptotic approximations focus on a fixed choice of moments $g(D_i, \beta, \delta)$, which we take as given. This is common in practice, including in all of the empirical papers using conditional moment inequalities that we discuss above, and is without loss of generality if the instruments Z_i have finite support.

For Z_i continuously distributed, however, a single conditional moment inequality implies an uncountable family of possible moments. Specifically, given a moment

function $\tilde{g}(D_i, \beta, \delta)$ that satisfies (1), for $f(Z_i)$ non-negative

$$g(D_i, \beta, \delta) = \tilde{g}(D_i, \beta, \delta)f(Z_i)$$

also satisfies (1). To obtain consistent tests (that is, tests that reject all values $\beta_0 \notin B_I(P)$ with probability going to one as $n \rightarrow \infty$), one may need to consider an infinite number of inequalities in large samples.⁶ Motivated by this fact, the literature on conditional moment inequalities, including D. Andrews & Shi (2013), Armstrong (2014b) and Chetverikov (2018), has primarily focused on consistent and efficient inference on (β, δ) jointly, based on checking (at least asymptotically) an infinite number of inequality restrictions. More recently, Chernozhukov et al. (2015) have developed results that can be used for subvector inference with conditional moment inequalities. Whether one can combine the results we develop here with results from the previous literature on conditional moment inequalities to obtain tests that are consistent in settings with continuously distributed Z_i is an interesting topic for future work.

3.1 Comparison to Unconditional Approximation

In many empirical applications using conditional moment inequalities, inference is based on asymptotic approximations that do not condition on $\{Z_i\}$. This section explores the relationship between such unconditional asymptotic approximations and our conditional approach.

Lemma 2 *Suppose that $E_P[Y_i Y_i']$ and $E_P[X_i X_i']$ are both finite. Then for all δ ,*

$$Y_n - X_n \delta - E_P[Y_n - X_n \delta] \rightarrow_d N(0, \Omega(\delta))$$

for $\Omega(\delta) = \text{Var}_P(Y_i - X_i \delta)$.

This suggests the approximation

$$Y_n - X_n \delta \approx^d N(E_P[Y_n - X_n \delta], \Omega(\delta)) \quad (8)$$

where $\tilde{H}_0 : \beta_0 \in B_I(P)$ implies that $E_P[Y_n - X_n \delta] \leq 0$ for some δ . Many commonly-used approaches to testing joint hypotheses on (β, δ) , including D. Andrews & Soares

⁶In particular, for a fixed, finite set of moments we may have $\mu_n \in \mathcal{M}_0$ with probability approaching one even though $\beta_0 \notin B_I(P)$.

(2010), D. Andrews & Barwick (2012), and Romano et al. (2014a), can be interpreted as applications of this approximation.⁷

Both (7) and (8) imply that the moments $g_n(\delta) = Y_n - X_n\delta$ are approximately normal, but the means and variances differ. Considering first the mean vectors, note that by the law of iterated expectations

$$E_P[\mu_n - X_n\delta] = E_P[Y_n - X_n\delta].$$

Thus, the mean vectors in (7) and (8) coincide on average, but the mean vector in (7) is random from an unconditional perspective while that in (8) is fixed.

Turning next to the variance matrices, by the law of total variance

$$\begin{aligned}\Omega(\delta) &= E_P[Var_P(Y_i - X_i\delta|Z_i)] + Var_P(E_P[Y_i - X_i\delta|Z_i]) \\ &= \Sigma + Var_P(\mu_i - X_i\delta).\end{aligned}$$

Hence, we see that $\Omega(\delta)$ is always weakly larger than Σ in the usual matrix order, and will typically be strictly larger. Thus, using the conditional approximation (7) we obtain a smaller variance matrix.⁸ While the smaller variance matrix in the conditional approximation (7) will often lead to more powerful tests, one can show that this is not universally the case for the procedures we consider.⁹ Critically for our results, however, Σ does not depend on δ , whereas $Var_P(\mu_i - X_i\delta)$ does.

4 Least Favorable Tests

Recall that we are interested in testing the hypothesis $H_0 : \mu_n \in \mathcal{M}_0$ under the linear normal model (7). The unknown parameter δ appears in the null hypothesis, and is a nuisance parameter that needs to be dealt with to allow testing. A common approach to handling nuisance parameters in moment inequality settings is the pro-

⁷The main text in Romano et al. (2014a) uses bootstrap critical values, but the appendix, Romano et al. (2014b), develops results for the normal model.

⁸Conditional variances were previously considered by e.g. Chetverikov (2018) for inference with conditional moment inequalities, and by Kaido et al. (2019b) and Barseghyan et al. (2019) for settings with a discrete instrument. We discuss estimation of Σ in Section 6 below.

⁹Though the diagonal terms in Σ are smaller than those in $\Omega(\delta)$, and this will lead to larger values of the test statistics introduced below, their off diagonal correlations also differ, which can generate larger critical values.

jection method (see Canay & Shaikh 2017 for examples). We begin by describing the projection method in our setting. We then explain why linear conditional structure allows us to eliminate the computational problems which can arise for the projection method. Finally, to avoid the conservativeness of the projection method, we derive (non-conservative) least-favorable critical values.

4.1 A Projection Method Test

The projection method tests the family of hypotheses

$$H_0(\delta) : \mu_n - X_n\delta \leq 0, \quad \delta \in \mathbb{R}^p \quad (9)$$

and rejects $H_0 : \mu_n \in \mathcal{M}_0$ if and only if we reject $H_0(\delta)$ for all δ . Provided our tests of $H_0(\delta)$ control size the projection method test does as well, since one of the hypotheses tested corresponds to the true δ .

Note that under $H_0(\delta)$, $Y_n - X_n\delta$ is normally distributed with a weakly negative mean. Thus, testing $H_0(\delta)$ reduces to testing that the mean of a multivariate normal vector is less than or equal to zero. A number of tests have been proposed for this hypothesis, but here we focus on tests that reject for large values of the max statistic

$$S(Y_n - X_n\delta, \Sigma) = \max_j \left\{ (Y_{n,j} - X_{n,j}\delta) / \sqrt{\Sigma_{jj}} \right\}$$

where $Y_{n,j} - X_{n,j}\delta$ denotes the j th element of the vector $Y_n - X_n\delta$ and Σ_{jj} is the j th diagonal element of Σ , which we assume throughout is strictly positive for all j .¹⁰ This choice of test statistic will allow us to compute projection tests of the composite hypothesis $H_0 : \mu_n \in \mathcal{M}_0$ via linear programming. That said, many of the results of this section (though not those in the following section) extend directly to other statistics $S(\cdot, \cdot)$ that are elementwise increasing in the first argument.

To test $H_0(\delta)$ based on $S(Y_n - X_n\delta, \Sigma)$, we need a critical value. As discussed in e.g. Rosen (2008) and D. Andrews & Guggenberger (2009), to ensure correct size we can compare $S(Y_n - X_n\delta, \Sigma)$ to the maximum of its $1 - \alpha$ quantile over data generating processes consistent with $H_0(\delta)$. Formally, let $c_\alpha(\gamma, \Sigma)$ be the $1 - \alpha$ -quantile of $S(\xi, \Sigma)$

¹⁰Desirable properties for tests based on this statistic are discussed by Armstrong (2014a).

for $\xi \sim N(\gamma, \Sigma)$. The least favorable critical value is then

$$c_{\alpha, LFP}(\Sigma) = \sup_{\gamma \leq 0} c_{\alpha}(\gamma, \Sigma) = c_{\alpha}(0, \Sigma),$$

where the fact that the sup is achieved at $\gamma = 0$ follows from the fact that S is elementwise increasing in its first argument. We subscript by LFP to emphasize that this is the least favorable critical value for testing $H_0(\delta)$, which is in turn part of the projection test for H_0 .

If we define the test of $H_0(\delta)$ to reject when $S(Y_n - X_n\delta, \Sigma)$ exceeds $c_{\alpha, LFP}(\Sigma)$,

$$\phi_{LF}(\delta) = 1\{S(Y_n - X_n\delta, \Sigma) > c_{\alpha, LFP}(\Sigma)\},$$

where we use $\phi = 1$ and $\phi = 0$ to denote rejection and non-rejection respectively, then it follows from the argument above that $\phi_{LF}(\delta)$ has size α as a test of $H_0(\delta)$:

$$\sup_{\mu_n: \mu_n - X_n\delta \leq 0} E_{\mu_n}[\phi_{LF}(\delta)] = \alpha.$$

The least favorable projection test of H_0 rejects if and only if $\phi_{LF}(\delta)$ rejects for all δ ,

$$\phi_{LFP} = \inf_{\delta} \phi_{LF}(\delta) = 1 \left\{ \min_{\tilde{\delta}} S(Y_n - X_n\tilde{\delta}, \Sigma) > c_{\alpha, LFP}(\Sigma) \right\}.$$

For any $\mu_n \in \mathcal{M}_0$ we know that there exists $\delta(\mu_n)$ such that $\mu_n - X_n\delta(\mu_n) \leq 0$, so

$$\sup_{\mu_n \in \mathcal{M}_0} E_{\mu_n}[\phi_{LFP}] \leq \alpha.$$

As we now show, the fact that neither $\min_{\tilde{\delta}} S(Y_n - X_n\tilde{\delta}, \Sigma)$ nor the critical value $c_{\alpha, LFP}(\Sigma) = c_{\alpha}(0, \Sigma)$ depends on δ makes ϕ_{LFP} particularly easy to compute.

Lemma 3 *We can write*

$$\phi_{LFP} = 1 \{ \hat{\eta} > c_{\alpha, LFP}(\Sigma) \}$$

for $\hat{\eta}$ the solution to

$$\begin{aligned} & \min_{\eta, \delta} \eta \\ & \text{subject to } (Y_{n,j} - X_{n,j}\delta) / \sqrt{\Sigma_{jj}} \leq \eta \quad \forall j. \end{aligned} \tag{10}$$

Thus, to calculate ϕ_{LFP} we need only solve a linear programming problem and calculate $c_{\alpha,LFP}(\Sigma)$. Hence, ϕ_{LFP} remains tractable even when the dimension of δ is large.¹¹ The linear normal model (7) plays a key role in this result in two ways, first through linearity in δ and second, perhaps less obviously, through the fact that the covariance Σ (and thus the critical value $c_{\alpha,LFP}(\Sigma)$) does not depend on δ .

If we instead considered projection tests based on the unconditional normal approximation (8), this corresponds to substituting $\Omega(\delta)$ for Σ in our expressions for $\phi_{LF}(\delta)$ and ϕ_{LFP} , and implies the unconditional projection method test

$$\phi_{LFP}^U = 1 \left\{ \min_{\delta} (S(Y_n - X_n\delta, \Omega(\delta)) - c_{\alpha,LFP}(\Omega(\delta))) > 0 \right\}.$$

The dependence of $\Omega(\delta)$ on δ means that evaluating this test requires nonlinear optimization. While this problem can be solved numerically when the dimension of δ is low, when the dimension is high this becomes computationally taxing.¹²

Thus, we see that the linear conditional structure we assume allows us to easily calculate the least favorable projection method test ϕ_{LFP} . As discussed by Bugni et al. (2017) and Kaido et al. (2019a), however, projection method tests are typically conservative,

$$\sup_{\mu_n \in \mathcal{M}_0} E_{\mu_n} [\phi_{LFP}] < \alpha,$$

and can be severely so when the dimension of the nuisance parameter δ is large.

4.2 A Least Favorable Test

To see why the projection test ϕ_{LFP} is conservative, recall that its critical value is calculated as the $1 - \alpha$ quantile of $S(\xi, \Sigma)$ where $\xi \sim N(0, \Sigma)$. By contrast, $\hat{\eta}$ is equal to $\min_{\delta} S(Y_n - X_n\delta, \Sigma)$. Hence, $c_{\alpha,LFP}(\Sigma)$ does not account for minimization over δ . In this section we use the structure of the normal linear model (7) to derive smaller, non-conservative least favorable critical values that account for minimization over δ .

¹¹Other recent applications of linear programming in set-identified settings include Mogstad et al. (2018), Khan et al. (2019), Tebaldi et al. (2019), and Torgovitsky (2019).

¹²Kaido et al. (2019a) discuss a response surface approach to speed this optimization in a more general setting.

Specifically, define $c_\alpha(\mu_n, X_n, \Sigma)$ as the $1 - \alpha$ quantile of

$$\begin{aligned} & \min_{\eta, \delta} \eta \\ & \text{subject to } (\xi_j - X_{n,j}\delta)/\sqrt{\Sigma_{jj}} \leq \eta \quad \forall j. \end{aligned} \tag{11}$$

when $\xi \sim N(\mu_n, \Sigma)$. The (non-conservative) least favorable critical value is

$$c_{\alpha, LF}(X_n, \Sigma) = \sup_{\mu_n \in \mathcal{M}_0} c_\alpha(\mu_n, X_n, \Sigma).$$

Note that the least favorable projection critical value $c_{\alpha, LFP}(\Sigma)$ corresponds to setting $\delta = 0$ in (11), rather than minimizing. Hence, by construction $c_{\alpha, LF}(X_n, \Sigma) \leq c_{\alpha, LFP}(\Sigma)$. If we define the least favorable test to reject when the max statistic exceeds $c_{\alpha, LF}(X_n, \Sigma)$,

$$\phi_{LF} = \left\{ \min_{\delta} S(Y_n - X_n\delta, \Sigma) > c_{\alpha, LF}(X_n, \Sigma) \right\} = \{\hat{\eta} > c_{\alpha, LF}(X_n, \Sigma)\},$$

then provided $\hat{\eta}$ is continuously distributed this test has size α ,

$$\sup_{\mu_n \in \mathcal{M}_0} E_{\mu_n} [\phi_{LF}] = \alpha.$$

If instead the distribution of $\hat{\eta}$ has point mass, the size is bounded above by α .

While describing the least favorable critical value $c_{\alpha, LF}(X_n, \Sigma)$ is conceptually straightforward, to derive it in practice we need to maximize the quantile $c_\alpha(\mu_n, X_n, \Sigma)$ over the set of μ_n values consistent with the null. The linear structure of the problem implies that the maximum is attained at $\mu_n = 0$.

Proposition 1

$$c_{\alpha, LF}(X_n, \Sigma) = c_\alpha(0, X_n, \Sigma).$$

This result follows immediately from the observations that (i) $c_\alpha(\mu_n, X_n, \Sigma)$ is invariant to shifting μ_n by $X_n\tilde{\delta}$, in the sense that for all $\tilde{\delta}$,

$$c_\alpha(\mu_n, X_n, \Sigma) = c_\alpha(\mu_n + X_n\tilde{\delta}, X_n, \Sigma),$$

(ii) that $c_\alpha(\mu_n, X_n, \Sigma)$ is non-decreasing in μ_n , and (iii) that for every $\mu_n \in \mathcal{M}_0$ there exists $\delta(\mu_n)$ such that $\mu_n - X_n\delta(\mu_n) \leq 0$.

To calculate the LF critical value we can simulate draws $\xi \sim N(0, \Sigma)$, solve the

linear programming problem (11) for each draw, and take the $1 - \alpha$ quantile of the resulting optimized values. While the need to repeatedly solve the problem (11) means that this approach requires more computation than the projection method, it remains highly tractable and yields a non-conservative test.

5 Conditional and Hybrid Tests

While less conservative than the projection approach, least favorable critical values still assume that all the moments are binding, $\mu_n = 0$. In practice we may suspect that some of the moments are far from binding, and the data may be informative about this. Motivated by this fact, D. Andrews & Soares (2010), D. Andrews & Barwick (2012), Romano et al. (2014a), and related papers propose techniques that use information from the data to either select moments or shift the mean of the distribution from which the critical values are calculated. This allows them to construct tests with higher power in empirically relevant cases where many of the moments are slack.

In our setting one can test $H_0 : \mu_n \in \mathcal{M}_0$ by first using one of the aforementioned approaches to test $H_0(\delta)$ as defined in (9) for all δ and then applying the projection method. This yields a conservative test, but Kaido et al. (2019a) show how to eliminate this conservativeness when considering projections based on D. Andrews & Soares (2010). Unfortunately, however, projection tests based on moment-selection procedures break the linear structure discussed in the last section. Implementing these approaches consequently requires solving a nonlinear, non-convex optimization problem.

To obtain procedures which both perform well when we have slack moments and preserve linearity, we introduce a novel conditional testing approach. When there is a unique, non-degenerate solution in the linear program (10), exactly $p + 1$ of the inequality constraints bind at the optimum. We propose tests which condition on the identity of these binding moments, and on a sufficient statistic for the slackness of the remaining moments. These tests control size both conditional on the set of binding moments and unconditionally, and are highly computationally tractable. Moreover, these tests are insensitive to the presence of slack moments in the sense that as a subset of the moments grows arbitrarily slack the conditional test approaches the test which drops the slack moments ex-ante. Conditional tests thus automatically incorporate a strong form of moment selection.

When the solution to (10) is non-unique or degenerate the set of binding moments is no longer uniquely defined, which would seem to pose a problem for the conditional test as described above. We show, however, that a reformulation of the conditional approach based on the dual linear program continues to apply in such settings. This approach is equivalent to conditioning on the set of binding moments in (10) when there is a unique, non-degenerate solution but remains valid and easy to implement even when these conditions fail.

In what follows, we first introduce the test in a special case where there are no nuisance parameters δ before turning to our results for the general case with a unique, non-degenerate solution. Results for the formulation based on the dual linear program, which allow for non-unique or degenerate solutions, are discussed in Section 5.3 and formally developed in Appendix A.

5.1 Special Case: No Nuisance Parameters

To develop intuition for our conditional approach we first consider a model without nuisance parameters δ . To further simplify, we assume that the variance is equal to the identity matrix, $\Sigma = I$. Our problem then reduces to that of testing $\mu_n \leq 0$ based on $Y_n \sim N(\mu_n, I)$, which has been well-studied in the previous literature.

In this setting, $\hat{\eta}$ is simply the max of the moments, $\hat{\eta} = S(Y_n, I) = \max_j \{Y_{n,j}\}$. With probability one there is a unique binding constraint in the linear program (10), corresponding to the largest moment. Once we condition on the identity of the largest moment, $\hat{j} = \arg \max_j Y_{n,j}$, the problem becomes one of inference based on a normal vector conditional on the max occurring at a particular location, $\hat{j} = j$.

Unfortunately, the distribution of $\hat{\eta} = Y_{n,\hat{j}}$ conditional on $\hat{j} = j$ still depends on the full vector μ_n . This dependence comes from the fact that $\hat{j} = j$ if and only if $Y_{n,j} \geq \max_{\tilde{j} \neq j} Y_{n,\tilde{j}}$, where the distribution of the lower bound depends on $\{\mu_{n,\tilde{j}} : \tilde{j} \neq j\}$. To eliminate this dependence, we further condition on the value of the second largest moment. Once we condition on $\hat{j} = j$ and on the value of the second largest moment, say $\max_{\tilde{j} \neq j} Y_{n,\tilde{j}} = \mathcal{V}^{lo}$, $\hat{\eta}$ follows a truncated normal distribution

$$\hat{\eta} | \left\{ \hat{j} = j \ \& \ \max_{\tilde{j} \neq j} Y_{n,\tilde{j}} = \mathcal{V}^{lo} \right\} \sim \xi | \mathcal{V}^{lo} \leq \xi$$

for $\xi \sim N(\mu_{n,j}, 1)$.

Lemma A.1 of Lee et al. (2016) shows that this truncated normal distribution is increasing in $\mu_{n,j}$, so since $\mu_{n,j} \leq 0$ under the null, the $1 - \alpha$ quantile of the conditional distribution under $\mu_{n,j} = 0$ is a valid conditional critical value. We denote this conditional critical value by $c_{\alpha,C}(j, \mathcal{V}^{lo}, I)$. The conditional test

$$\phi_C = 1 \left\{ \hat{\eta} > c_{\alpha,C} \left(\hat{j}, \max_{\tilde{j} \neq \hat{j}} Y_{n,\tilde{j}}, I \right) \right\}$$

has maximal rejection probability equal to α under the null, conditional on $\hat{j} = j$ and $\max_{\tilde{j} \neq j} Y_{n,\tilde{j}} = \mathcal{V}^{lo}$. By the law of iterated expectations its unconditional rejection probability under the null is thus bounded above by α as well, and this bound is achieved at $\mu_n = 0$. Thus, ϕ_C is a size α test of $H_0 : \mu_n \leq 0$.

The simplicity of the present setting allows us to highlight some important features of the conditional test. When the second largest element of μ_n , say $\max_{\tilde{j} \neq j} \mu_{n,\tilde{j}}$, is very negative while the largest element ($\mu_{n,j}$) is not, $\hat{j} = j$ with high probability. In this case, the lower truncation point is very small with high probability, so the truncated normal critical value $c_{\alpha,C} \left(\hat{j}, \max_{\tilde{j} \neq \hat{j}} Y_{n,\tilde{j}}, I \right)$ is close to the level $1 - \alpha$ standard normal critical value with high probability. Thus, when the largest element of μ_n is well separated from the remaining elements, the conditional test closely resembles the test which limits attention to the j th moment ex-ante, $\phi_j = 1\{Y_{n,j} > c_\alpha\}$ for c_α the level $1 - \alpha$ standard normal critical value. The power of ϕ_j lies on the power envelope for tests of $H_0 : \mu_n \leq 0$ when all the other elements of μ_n are negative (see Romano et al. 2014b). Thus, the conditional test has power approaching the power envelope when we take all moments but one to be slack. More broadly, Proposition 3 below shows that if we take a subset of elements of μ_n to $-\infty$, the conditional test converges to the conditional test which drops the corresponding moments ex-ante.

The only other test that we know of which shares this strong insensitivity property, while also controlling size in the finite sample normal model, is that of Cox & Shi (2019).¹³ In particular, while the tests of D. Andrews & Barwick (2012) and Romano et al. (2014a) are relatively insensitive to the presence of slack moments, they are both affected by the addition of slack moments.¹⁴ While the test of Cox & Shi (2019) is

¹³Specifically, the baseline test discussed in that paper, not the modification discussed in their Remark 3. Interestingly, this test is also based on conditioning, though in the present example their approach conditions on the identity of the non-negative moments, $\{j : Y_j > 0\}$, while we condition on the identity of the largest moment and the value of the second-largest moment.

¹⁴Through the size correction factor in D. Andrews & Barwick (2012), and the first-stage critical

strongly insensitive to slack moments, its power does not in general converge to the power envelope in the case where all moments but one are slack.

This example also highlights a less desirable feature of our conditional test. When the largest element of μ_n is not well-separated, $\mu_{n,j} \approx \max_{\tilde{j} \neq j} \mu_{n,\tilde{j}}$, the second-largest moment $\max_{\tilde{j} \neq \hat{j}} Y_{n,\tilde{j}}$ will often be nearly as large as the largest moment. Since the conditional critical value $c_\alpha(\hat{j}, \max_{\tilde{j} \neq \hat{j}} Y_{n,\tilde{j}}, I)$ is always strictly larger than $\max_{\tilde{j} \neq \hat{j}} Y_{n,\tilde{j}}$, this can lead to poor power for the conditional test. We illustrate this issue in simulation in Appendix F.

Hybrid Tests To address power declines for the conditional test when the largest element of μ_n is not well-separated we introduce what we call a hybrid test. This modifies the conditional test to reject whenever the max statistic $\hat{\eta}$ exceeds a level $\kappa \in (0, \alpha)$ least-favorable critical value, $c_{\kappa,LF}(I)$. If $\hat{\eta} \leq c_{\kappa,LF}(I)$ we then consider a conditional test, where we (i) further condition on the event that $\hat{\eta} \leq c_{\kappa,LF}(I)$ and (ii) modify the level of the conditional test to reflect the first step. By the arguments above the distribution of $\hat{\eta}$, conditional on not rejecting in the first stage, is again truncated normal, now truncated both from below and above,

$$\hat{\eta} | \left\{ \hat{j} = j, \max_{\tilde{j} \neq j} Y_{n,\tilde{j}} = \mathcal{V}^{lo} \ \& \ \hat{\eta} \leq c_{\kappa,LF}(I) \right\} \sim \xi | \mathcal{V}^{lo} \leq \xi \leq c_{\kappa,LF}(I)$$

for $\xi \sim N(\mu_{n,j}, 1)$. For $c_{\tilde{\alpha},H}(j, \mathcal{V}^{lo}, I)$ the $1 - \tilde{\alpha}$ quantile of this distribution,

$$\inf_{\mu_n \leq 0} Pr_{\mu_n} \left\{ \hat{\eta} \leq c_{\tilde{\alpha},H}(j, \mathcal{V}^{lo}, I) | \hat{j} = j, \max_{\tilde{j} \neq j} Y_{n,\tilde{j}} = \mathcal{V}^{lo}, \hat{\eta} \leq c_{\kappa,LF}(I) \right\} = 1 - \tilde{\alpha}.$$

To form hybrid tests, we set $\tilde{\alpha} = \frac{\alpha - \kappa}{1 - \kappa}$ to account for the first-step comparison to the least favorable critical value. Since $c_{\tilde{\alpha},H}(j, \mathcal{V}^{lo}, I) \leq c_{\kappa,LF}$ by definition, we can thus write the hybrid test as

$$\phi_H = 1 \left\{ \hat{\eta} > c_{\frac{\alpha - \kappa}{1 - \kappa},H} \left(\hat{j}, \max_{\tilde{j} \neq \hat{j}} Y_{n,\tilde{j}}, I \right) \right\}.$$

This test again has rejection probability under the null bounded above by α , and this bound is attained at $\mu_n = 0$. By construction this test rejects whenever the level κ least favorable test does, which improves power relative to the conditional test when

value in Romano et al. (2014a).

the largest element of μ_n is not well-separated. While the hybrid test retains many of the properties of the conditional test, its dependence on the least-favorable critical value means that it is affected by the inclusion of even arbitrarily slack moments. Similar to the test of Romano et al. (2014a), however, the impact is small when κ is close to zero.

To illustrate the performance of hybrid tests in the present simplified setting, Appendix F reports simulation results for cases with two, ten, and fifty moments. We also calculate results for the test proposed by Romano et al. (2014a) for comparison. We find that the hybrid approach improves power relative to the conditional test in the poorly-separated case, while still improving power relative to the least favorable test in the well-separated case. Neither the hybrid test nor the test of Romano et al. (2014a) dominates the other: the test of Romano et al. (2014a) has better performance in the poorly-separated case, while the hybrid test has slightly higher power when the largest moment is moderately well-separated. Unlike the test of Romano et al. (2014a) however, the hybrid and conditional tests easily extend to the case with nuisance parameters δ . Simulation results based on Wollman (2018), reported in Section 7, demonstrate that the power gains of the hybrid test are borne out in more realistic settings with nuisance parameters.

5.2 Conditional Tests with Nuisance Parameters

We next discuss our conditional approach in the case with nuisance parameters δ and a covariance matrix Σ which may not equal the identity. In this section we assume that the linear program (10) has a unique, non-degenerate solution with probability one, while Appendix A develops an alternative formulation for the conditioning approach, based on the dual linear program, that does not impose these conditions. The primal and dual approaches are numerically equivalent when the solution to (10) is unique and non-degenerate (as we expect will often be the case in applications), so we focus on the primal approach here for ease of exposition.¹⁵

To define our conditional approach, note that we can rewrite (10) as

$$\begin{aligned} & \min_{\eta, \delta} \eta \\ & \text{subject to } Y_n - W_n(\eta, \delta')' \leq 0. \end{aligned} \tag{12}$$

¹⁵Degeneracy means that for W_n as defined below, the rows of W_n corresponding to binding constraints are linearly dependent. See Section 10.4 of Schrijver (1986).

for W_n the matrix with row j equal to $W_{n,j} = \begin{pmatrix} \sqrt{\Sigma_{jj}} & X_{n,j} \end{pmatrix}$. Let $(\hat{\eta}, \hat{\delta})$ denote the optimal values in (12), which we assume for the moment are unique, and let $\hat{B} \subseteq \{1, \dots, k\}$ collect the indices corresponding to the binding constraints at these optimal solutions, so $Y_{n,j} - W_{n,j}(\hat{\eta}, \hat{\delta})' = 0$ if and only if $j \in \hat{B}$. Let $Y_{n,\hat{B}}$ and $W_{n,\hat{B}}$ collect the corresponding rows of Y_n and W_n .

Lemma 4 *If the solution to (12) is unique and non-degenerate, $|\hat{B}| = p + 1$, and $W_{n,\hat{B}}$ has full rank.*

Since $Y_{n,\hat{B}} - W_{n,\hat{B}}(\hat{\eta}, \hat{\delta})' = 0$ by the definition of $\hat{B}$, Lemma 4 implies that $(\hat{\eta}, \hat{\delta})' = W_{n,\hat{B}}^{-1} Y_{n,\hat{B}}$. Thus, given a particular set of binding moments $\hat{B} = B$, we can write $\hat{\eta}$ as a linear function of Y_n ,

$$\hat{\eta} = \gamma'_{n,B} Y_n = e_1' W_{n,B}^{-1} Y_{n,B},$$

for e_1 the first standard basis vector.

We next consider under what conditions there exists a solution with moments B binding.

Lemma 5 *For $B \subseteq \{1, \dots, k\}$ such that $W_{n,B}$ is a square, full-rank matrix, there exists a solution with the moments B binding if and only if*

$$Y_n - W_n W_{n,B}^{-1} Y_{n,B} \leq 0. \quad (13)$$

Thus we see that there exists a solution with the moments B binding if and only if the implied $(\hat{\eta}, \hat{\delta})'$ make the constraints in (12) hold.

Our conditional test will condition on the existence of a solution with the moments B binding and reject when $\hat{\eta}$ is large relative to the resulting conditional distribution under the null. The set of values Y_n such that (13) holds is a polytope (a convex set with flat sides, also known as a polyhedron— see Schrijver 1986 pages 87-88), and as noted above we can write $\hat{\eta}$ as a linear function of Y_n conditional on this event. Thus, we are interested in the distribution of a linear function of a normal vector conditional on that vector falling in a polytope. Lee et al. (2016) consider problems of this form, and we can use their results to derive conditional critical values. We first calculate the range of possible values for $\hat{\eta}$ conditional on Y_n falling in this polytope. We then determine the distribution of $\hat{\eta}$ over this range conditional on a sufficient statistic for the part of μ_n not corresponding to $\hat{\eta}$.

To this end we use the following result, which is an immediate consequence of Lemma 5.1 of Lee et al. (2016).

Lemma 6 *Let M_B be the selection matrix which selects rows corresponding to B . Suppose that $W_{n,B}$ is a square, full-rank matrix, and let $\gamma_{n,B}$ be the vector with $M_B \gamma_{n,B} = W_{n,B}'^{-1} e_1$, and zeros elsewhere. Assume $\gamma_{n,B}' \Sigma \gamma_{n,B} > 0$. Let $\Lambda_{n,B} = I - W_n W_{n,B}^{-1} M_B$, and define*

$$\Delta_{n,B} = \frac{\Sigma \gamma_{n,B}}{\gamma_{n,B}' \Sigma \gamma_{n,B}},$$

and $S_{n,B} = (I - \Delta_{n,B} \gamma_{n,B}') Y_n$. Further define

$$\mathcal{V}^{lo}(S_{n,B}) = \max_{j: (\Lambda_{n,B} \Delta_{n,B})_j < 0} - \frac{(\Lambda_{n,B} S_{n,B})_j}{(\Lambda_{n,B} \Delta_{n,B})_j} \quad (14)$$

$$\mathcal{V}^{up}(S_{n,B}) = \min_{j: (\Lambda_{n,B} \Delta_{n,B})_j > 0} - \frac{(\Lambda_{n,B} S_{n,B})_j}{(\Lambda_{n,B} \Delta_{n,B})_j} \quad (15)$$

$$\mathcal{V}^0(S_{n,B}) = \min_{j: (\Lambda_{n,B} \Delta_{n,B})_j = 0} - (\Lambda_{n,B} S_{n,B})_j.$$

The set of values Y_n such that there exists a solution with the moments B binding is

$$\{Y_n : Y_n - W_n W_{n,B}^{-1} Y_{n,B} \leq 0\} = \{Y_n : \mathcal{V}^{lo}(S_{n,B}) \leq \gamma_{n,B}' Y_n \leq \mathcal{V}^{up}(S_{n,B}), \mathcal{V}^0(S_{n,B}) \geq 0\}.$$

This result shows that there exists a solution with the moments B binding if and only if $\gamma_{n,B}' Y_n$ lies between the data-dependent bounds $\mathcal{V}^{lo}(S_{n,B})$ and $\mathcal{V}^{up}(S_{n,B})$ and, in addition, $\mathcal{V}^0(S_{n,B}) \geq 0$. When such a solution exists, however, our arguments above show that $\hat{\eta} = \gamma_{n,B}' Y_n$. Thus, whenever there exists a solution with the moments B binding, $\hat{\eta}$ lies between $\mathcal{V}^{lo}(S_{n,B})$ and $\mathcal{V}^{up}(S_{n,B})$ by construction.

Lemma 6 assumes that $\gamma_{n,B}' \Sigma \gamma_{n,B} > 0$. This implies that $\hat{\eta}$ has a non-degenerate distribution conditional on the set of binding moments. While not necessary for our conditional testing approach, this simplifies a number of statements in what follows, so going forward we maintain a sufficient condition for $\gamma_{n,B}' \Sigma \gamma_{n,B} > 0$.¹⁶

Assumption 1 *For all γ with $W_n' \gamma = e_1$ and $\gamma \geq 0$, $\gamma' \Sigma \gamma > 0$.*

¹⁶If this condition fails, we can define our conditional test to reject whenever $\gamma_{n,B}' \Sigma \gamma_{n,B} = 0$ and $\hat{\eta} > 0$, but this results in tests with size bounded above by α , rather than exactly correct size.

One can show that $\gamma_{n,B}$ as defined in Lemma 6 has $W'_n \gamma_{n,B} = e_1$ and $\gamma_{n,B} \geq 0$ for any set of binding moments B . A sufficient, but not necessary, condition for Assumption 1 is that the variance matrix Σ is positive-definite.

Lemma 6 clarifies what it means to condition on the existence of a solution with the moments B binding, and thus the inference problem we need to solve. We are interested in the behavior of $\hat{\eta} = \gamma'_{n,B} Y_n$ conditional on the set of binding moments, but as in the simplified example above this conditional distribution depends on the full mean vector μ_n , rather than just on $\gamma'_{n,B} \mu_n$, due to the influence of the bounds $\mathcal{V}^{lo}(S_{n,B})$ and $\mathcal{V}^{up}(S_{n,B})$. Moreover, this conditional distribution is not in general monotonic in μ_n , making it difficult to find least favorable values. To eliminate dependence on μ_n other than through $\gamma'_{n,B} \mu_n$, we thus follow Lee et al. (2016) and further condition on $S_{n,B}$, which is the minimal sufficient statistic for the part of μ_n other than $\gamma'_{n,B} \mu_n$.¹⁷ Note that $\gamma'_{n,B} Y_n$ and $S_{n,B}$ are jointly normal and uncorrelated by construction, and thus independent. Hence, $\hat{\eta}$ follows a truncated normal distribution conditional on $S_{n,B}$ and the set of binding moments.

Lemma 7 *If the solution to (12) is unique and nondegenerate with probability one, the conditional distribution of $\hat{\eta}$ given $\hat{B} = B$ and $S_{n,B} = s$ is truncated normal,*

$$\hat{\eta} \mid \left\{ \hat{B} = B \text{ \& } S_{n,B} = s \right\} \sim \xi \mid \xi \in [\mathcal{V}^{lo}(s), \mathcal{V}^{up}(s)],$$

for $\xi \sim N(\gamma'_{n,B} \mu_n, \gamma'_{n,B} \Sigma \gamma_{n,B})$, provided we consider a value s such that $\mathcal{V}^0(s) \geq 0$.

As in Section 5.1 above, this truncated distribution is increasing in the mean $\gamma'_{n,B} \mu_n$. Since $\gamma_{n,B} \geq 0$, $\gamma'_{n,B} X_n = 0$,¹⁸ and $\mu_n - X_n \delta \leq 0$ under the null, the largest value of $\gamma'_{n,B} \mu_n$ possible under the null is zero. We define the conditional critical value $c_{\alpha,C}(\gamma, \mathcal{V}^{lo}, \mathcal{V}^{up}, \Sigma)$ to equal the $1 - \alpha$ quantile of the truncated normal distribution

$$\xi \mid \xi \in [\mathcal{V}^{lo}, \mathcal{V}^{up}]$$

for $\xi \sim N(0, \gamma'_{n,B} \Sigma \gamma_{n,B})$. We can write this critical value as

$$c_{\alpha,C}(\gamma, \mathcal{V}^{lo}, \mathcal{V}^{up}, \Sigma) = \sqrt{\gamma' \Sigma \gamma} \cdot \Phi^{-1}((1 - \alpha) \zeta^{up} + \alpha \zeta^{lo}) \quad (16)$$

¹⁷In particular, $S_{n,B}$ is minimal sufficient for $(I - \Delta_{n,B} \gamma'_{n,B}) \mu_n$ and μ_n is a one-to-one transformation of $(\gamma'_{n,B} \mu_n, (I - \Delta_{n,B} \gamma'_{n,B}) \mu_n)$, since $\mu_n = (I - \Delta_{n,B} \gamma'_{n,B}) \mu_n + \Delta_{n,B} \gamma'_{n,B} \mu_n$.

¹⁸This follows from Lemma 10 and Proposition 5 in Appendix A, but can also be verified directly using the Kuhn-Tucker conditions for optimality of $(\hat{\eta}, \hat{\delta})$.

for Φ^{-1} the inverse of the standard normal distribution function, and

$$(\zeta^{lo}, \zeta^{up}) = \left(\Phi \left(\mathcal{V}^{lo} / \sqrt{\gamma' \Sigma \gamma} \right), \Phi \left(\mathcal{V}^{up} / \sqrt{\gamma' \Sigma \gamma} \right) \right).$$

Thus, conditional critical values are easy to compute in practice.

Assuming the solution to (12) is unique and nondegenerate with probability one and Assumption 1 holds, the results above imply that the conditional test which compares $\hat{\eta}$ to the conditional critical value,

$$\phi_C = 1 \left\{ \hat{\eta} > c_{\alpha, C} \left(\gamma_{n, \hat{B}}, \mathcal{V}^{lo}(S_{n, \hat{B}}), \mathcal{V}^{up}(S_{n, \hat{B}}), \Sigma \right) \right\}, \quad (17)$$

rejects with probability at most α conditional on $\hat{B} = B$ under the null, and thus has unconditional size α as well.

Proposition 2 *If the solution to (12) is unique and non-degenerate with probability one and Assumption 1 holds, the conditional test ϕ_C has size α both conditional on $\hat{B}$,*

$$\sup_{\mu_n \in \mathcal{M}_0} E_{\mu_n} [\phi_C | \hat{B} = B] = E_0 [\phi_C | \hat{B} = B] = \alpha$$

for all B such that $Pr_{\mu_n} \{ \hat{B} = B \} > 0$, and unconditionally,

$$\sup_{\mu_n \in \mathcal{M}_0} E_{\mu_n} [\phi_C] = E_0 [\phi_C] = \alpha.$$

5.3 Conditional Tests Without Uniqueness

In our discussion of conditional tests so far we have relied on the uniqueness and non-degeneracy of the solution to ensure both that the set of binding moments $\hat{B}$ is uniquely defined and that the matrix $W_{n, B}$ is invertible. While these assumptions allow us to obtain simple expressions for conditional tests, they are not essential. Even when the solution $(\hat{\eta}, \hat{\delta})$ is nonunique or degenerate, $\hat{\eta}$ is unique. Our conditioning approach for the normal model remains valid in such cases, but we need to work with the dual linear program to (12). This dual conditioning approach is numerically equivalent to that described above when the primal solution is unique and non-degenerate. Since formally developing the dual approach requires additional notation and adds little intuition relative to the results above, we defer this development to Appendix A. There we formally establish the numerical equivalence of the primal and dual approaches

when the former is valid, as well as conditional and unconditional size control for our conditional tests based on the dual in the normal model, even when the primal solution may be non-unique or degenerate. To prove asymptotic validity of the conditional approach with non-normal data, our results in Appendix D require that the primal solution be nondegenerate with probability one asymptotically, though it may be non-unique. A sufficient condition for non-degeneracy is that Σ has full rank, so this condition can be made to hold mechanically by adding a small amount of full-rank noise to Y_n .

It is often not obvious whether the solution to (12) will be unique and non-degenerate with probability one in a given setting. Fortunately, the results in Appendix A suggest a simple way to proceed in practice, based on the fact that the widely-used dual-simplex algorithm for solving the primal problem (12) automatically generates a vertex $\hat{\gamma}$ of the dual solution set as well. Proposition 5 in Appendix A shows that so long as $\hat{\gamma}$ has exactly $p + 1$ strictly positive entries, and the rows of W_n corresponding to these positive entries have full rank, we can take $\hat{B}$ to collect the corresponding indices and apply the results developed above. If this condition fails, then we should use the more general expressions developed in Appendix A.

5.4 Performance with Slack Moments

We motivated our study of conditional tests by a desire to reduce sensitivity to slack moments. To formally understand the behavior of conditional tests in cases where some of the moments are slack, we will consider a sequence of mean vectors $\mu_{n,m}$, indexed by m , such that a subset of the moments grow arbitrarily slack as $m \rightarrow \infty$ while the remaining moments are unchanged. This yields the following result, which generalizes the insensitivity to slack moments noted in Section 5.1 for the special case without nuisance parameters to our general setting.

Proposition 3 *Consider a sequence of mean vectors $\mu_{n,m}$ where $\mu_{n,m,j} \equiv \mu_{n,j} \in \mathbb{R}$ for all m if $j \in \mathcal{B}$, while $\mu_{n,m,j} \rightarrow -\infty$ as $m \rightarrow \infty$ if $j \notin \mathcal{B}$. Let us further suppose that there exists $\gamma_{\mathcal{B}} \geq 0$ with $W'_{n,\mathcal{B}}\gamma_{\mathcal{B}} = e_1$. Under Assumption 1, for $Y_{n,m} \sim N(\mu_{n,m}, \Sigma)$, $\phi_{C,m}$ the conditional test based on $(Y_{n,m}, W_n, \Sigma)$, and $\phi_{C,m}^{\mathcal{B}}$ the conditional test based on $(Y_{n,m,\mathcal{B}}, W_{n,\mathcal{B}}, \Sigma_{\mathcal{B}})$, $\phi_{C,m} \rightarrow_p \phi_{C,m}^{\mathcal{B}}$ as $m \rightarrow \infty$.¹⁹*

¹⁹Indeed, the same conclusion holds if there exists a sequence δ_m and a vector δ such that $\mu_{n,m,j} - X_{n,j}\delta_m = \mu_{n,j} - X_{n,j}\delta \in \mathbb{R}$ for all m if $j \in \mathcal{B}$, while $\mu_{n,m,j} - X_{n,j}\delta_m \rightarrow -\infty$ as $m \rightarrow \infty$ if $j \notin \mathcal{B}$.

The restriction on $W_{n,\mathcal{B}}$ ensures that the feasible set in the dual problem based on $(Y_{n,m,\mathcal{B}}, W_{n,\mathcal{B}}, \Sigma_{\mathcal{B}})$ is non-empty, and thus that the solution in the primal problem is finite (see Section 7.4 of Schrijver (1986)). When this condition fails, the optimal value $\hat{\eta}$ diverges to $-\infty$.

Proposition 3 shows that the conditional tests we consider are robust to the presence of slack moments in a very strong sense. In particular, when a subset of moments become arbitrarily slack, the conditional test converges in probability to the test which drops these moments ex-ante. As noted above, even in settings without nuisance parameters the only other test we are aware of with this property in the normal model is that of Cox & Shi (2019), and their approach does not address settings with nuisance parameters (other than through projection).

5.5 Hybrid Tests

In Section 5.1 above, we noted that in the special case without nuisance parameters conditional tests can have poor power in settings where the lower bound used by the conditional test is large with high probability. The same issue arises more broadly, and as in the case without nuisance parameters we can obtain improved performance by considering hybrid tests.

For some $\kappa \in (0, \alpha)$ the hybrid test rejects whenever $\hat{\eta}$ exceeds the level κ least-favorable critical value $c_{\kappa,LF}(X_n, \Sigma)$.²⁰ When $\hat{\eta}$ is less than this conditional critical value, the hybrid test compares $\hat{\eta}$ to a modification of the conditional critical value that also conditions on $\hat{\eta} \leq c_{\kappa,LF}(X_n, \Sigma)$. This reduces $\mathcal{V}^{up}(s)$ to

$$\mathcal{V}^{up,H}(s) = \min \{ \mathcal{V}^{up}(s), c_{\kappa,LF}(X_n, \Sigma) \}.$$

The level α hybrid test rejects whenever $\hat{\eta}$ exceeds the level $\frac{\alpha-\kappa}{1-\kappa}$ conditional critical value based on the modified truncation points, where we define this quantile to equal $-\infty$ if $\mathcal{V}^{lo}$ exceeds $\mathcal{V}^{up,H}$,

$$\phi_H = \left\{ \hat{\eta} > c_{\frac{\alpha-\kappa}{1-\kappa},C} \left(\hat{\gamma}, \mathcal{V}^{lo}(S_{n,\hat{B}}), \mathcal{V}^{up,H}(S_{n,\hat{B}}), \Sigma \right) \right\}.$$

²⁰Similar to Romano et al. (2014a), we consider $\kappa = \alpha/10$ in our simulations below. Either $c_{\alpha,LP}(\Sigma)$ or $c_{\alpha,LF}(X_n, \Sigma)$ could be used here, the tradeoff being that $c_{\alpha,LF}(X_n, \Sigma)$ will provide a smaller critical value but will have a somewhat higher computational burden.

Since

$$c_{\frac{\alpha-\kappa}{1-\kappa}, C} \left(\hat{\gamma}, \mathcal{V}^{lo}(S_{n, \hat{B}}), \mathcal{V}^{up, H}(S_{n, \hat{B}}), \Sigma \right) \leq \mathcal{V}^{up, H}(S_{n, \hat{B}}) \leq c_{\kappa, LF}(X_n, \Sigma),$$

this test always rejects when $\hat{\eta} > c_{\kappa, LF}(X_n, \Sigma)$, as claimed above. The hybrid test has size equal to $\frac{\alpha-\kappa}{1-\kappa}$ conditional on $\hat{\eta} \leq c_{\kappa, LF}$ and the set of binding moments, and unconditional size equal to α .

Proposition 4 *If the solution to (12) is unique and non-degenerate with probability one, and Assumption 1 holds, the hybrid test ϕ_H has size $\frac{\alpha-\kappa}{1-\kappa}$ conditional on $\hat{\eta} \leq c_{\kappa, LF}(X_n, \Sigma)$ and $\hat{B} = B$,*

$$\sup_{\mu_n \in \mathcal{M}_0} E_{\mu_n} \left[\phi_H | \hat{\eta} \leq c_{\kappa, LF}(X_n, \Sigma), \hat{B} = B \right] = E_0 \left[\phi_H | \hat{\eta} \leq c_{\kappa, LF}(X_n, \Sigma), \hat{B} = B \right] = \frac{\alpha - \kappa}{1 - \kappa},$$

for all B such that $Pr_{\mu_n} \left\{ \hat{\eta} \leq c_{\kappa, LF}(X_n, \Sigma), \hat{B} = B \right\} > 0$, and has unconditional size α ,

$$\sup_{\mu_n \in \mathcal{M}_0} E_{\mu_n} [\phi_H] = E_0 [\phi_H] = \alpha.$$

Thus, we see that our hybrid approach yields a non-conservative level α test. Due to the inclusion of the least favorable critical value $c_{\kappa, LF}(X_n, \Sigma)$ this test no longer shares the strong insensitivity to slack moments established for the conditional test by Proposition 3. That said, as a set of moments becomes slack the power of the hybrid test is bounded below by the power of the size $\frac{\alpha-\kappa}{1-\kappa}$ conditional test that drops these moments ex-ante. Moreover, the Monte Carlo results in Section 7 show that the hybrid does noticeably better than both the conditional and least favorable tests in some cases with slack moments. Appendix A establishes size control for hybrid tests based on the dual approach even when the solution to (12) is non-unique or degenerate.

6 Implementation

This section provides guidance for researchers seeking to implement the methods described in this paper. As in our theoretical results above, we assume that the researcher has a moment function

$$g(D_i, \beta, \delta) = Y_i(\beta) - X_i(\beta)\delta \tag{18}$$

for $Y_i(\beta) \in \mathbb{R}^k$, $\delta \in \mathbb{R}^p$, and $X_i(\beta)$ a $k \times p$ matrix. We assume that at the true parameter values $E_P[Y_i(\beta) - X_i(\beta)\delta|Z_i] \leq 0$, where Z_i is a vector of instruments and $X_i(\beta)$ is non-random given Z_i . We suppose the researcher wishes to compute confidence sets for β . This is often done by discretizing the parameter space for β as $\{\beta_1, \dots, \beta_L\}$, and then testing pointwise whether each β_l in the grid is contained in $B_I(P)$. The confidence set then collects the non-rejected points.

Sections 6.1 to 6.4 provide guidance on how to test whether a single value of β is in the identified set, which can then be applied to all points in the grid. Sections 6.5 and 6.6 discuss implementation in extensions of this basic setting, such as when the researcher wishes to conduct inference on (functions of) linear parameters, or when there are non-linear nuisance parameters.

Alternative Procedures While the linear conditional structure assumed in this paper is present in a variety of moment inequality settings, there are practically important cases where our results do not apply but alternatives are available. First, one may have unconditional moment inequalities that are nonetheless linear in the parameters, in which case one can use the approaches of Cho & Russell (2019) or Gafarov (2019). Alternatively, in settings with unconditional moment inequalities that may or may not be linear in the nuisance parameters δ , or where we may be interested in a nonlinear function of the parameters, one can use the approaches of e.g. Bugni et al. (2017) and Kaido et al. (2019a). For more discussion of the comparison among these options, see Kaido et al. (2019a) and Gafarov (2019). Other alternatives include the procedures discussed by Romano & Shaikh (2008) and Chen et al. (2018).

Asymptotic validity for the procedures discussed above (and for the present paper – see Appendix D) are established under the assumption that the number of moments is fixed as the sample size tends to infinity. This assumption may yield unsatisfactory performance if the number of moments is large relative to the sample size. By contrast, the approach of Belloni et al. (2018) gives guarantees even in high-dimensional settings, while the approach of Flynn (2019) allows a continuum of moments. Finally, the results of Chernozhukov et al. (2015) apply in conditional moment settings where the moments may be nonlinear in the nuisance parameters, and dimension of $g(D_i, \beta, \delta)$ may be large.

6.1 Estimating Σ

All of the tests for whether $\beta \in B_I(P)$ described in this paper require an estimate of the average conditional variance $\Sigma(\beta) = E_P[Var_P(Y_i(\beta)|Z_i)]$. It is important to note that $\Sigma(\beta)$ depends on the non-linear parameter β , and thus must be estimated at each grid point; for ease of exposition, however, we fix β and drop the explicit dependence of Σ , Y , and X on β for the remainder of the section.

The average conditional variance Σ can be estimated using the matching procedure proposed by Abadie et al. (2014). To do this, define $\Sigma_Z = \widehat{Var}(Z_i)$.²¹ For each i , find the nearest neighbor using the Mahalanobis distance in Z_i :

$$\ell_Z(i) = \operatorname{argmin}_{j \in \{1, \dots, n\}, j \neq i} (Z_i - Z_j)' \Sigma_Z^{-1} (Z_i - Z_j).$$

The estimate of Σ is then:

$$\widehat{\Sigma} = \frac{1}{2n} \sum_{i=1}^n (Y_i - Y_{\ell_Z(i)}) (Y_i - Y_{\ell_Z(i)})'.$$

Proposition 10 in Appendix D proves that, under additional assumptions, $\widehat{\Sigma}$ consistently estimates Σ .

6.2 Implementing the LF and LFP Tests

We can test whether a particular value β is in the identified set using the LF or LFP tests by solving the linear program (10) and rejecting if and only if the optimal value $\hat{\eta}$ exceeds a critical value.

To compute the least-favorable projection critical value via simulation, draw a $k \times S$ matrix Ξ of independent standard normals.²² Let Ξ^{max} denote the $S \times 1$ vector where the s th element is the maximum of the s th column of $\widehat{\Sigma}^{1/2} \Xi$. Set $c_{\alpha, LFP}(\widehat{\Sigma})$ to the $1 - \alpha$ quantile of Ξ^{max} .

²¹The matching procedure described below assumes that $\widehat{Var}(Z_i)$ is invertible. In certain applications, such as in our Monte Carlo, elements of Z_i may be linearly dependent by construction, leading $\widehat{Var}(Z_i)$ to be singular. In this case conditioning on a maximal linearly independent subset of Z_i is equivalent to conditioning on the full vector, so one can drop dependent elements from Z_i until $\widehat{Var}(Z_i)$ is invertible.

²²Note that Ξ need only be drawn once, and can be reused for many iterations of the LFP test, as well as for the LF test. Holding the simulation draws fixed as we vary β is likely to produce confidence sets with smoother boundaries and may ease the computational burden.

Similarly, to compute the least favorable critical value, again let Ξ be a $k \times S$ matrix of independent standard normal draws. Denote by ξ_s the s th column of $\widehat{\Sigma}^{1/2} \Xi$. For each $s = 1, \dots, S$, calculate

$$\eta_s = \begin{array}{ll} \min_{\eta, \delta} \eta \\ \text{subject to } (\xi_s - X_{n,j} \delta) / \sqrt{\widehat{\Sigma}_{jj}} \leq \eta \quad \forall j. \end{array}$$

Set $c_{\alpha, LF}(\widehat{\Sigma})$ to the $1 - \alpha$ quantile of $\{\eta_1, \dots, \eta_S\}$.

6.3 Implementing the Conditional Test

To implement the conditional test in practice, we recommend taking the following steps:

1. Solve the primal LP (10) using the dual-simplex method, which generates as a byproduct multipliers $\hat{\gamma}$ corresponding to a vertex of the solution set in the dual problem (see Appendix A).
2. Check whether there are exactly $p+1$ positive multipliers in $\hat{\gamma}$, and if so, whether the rows of the constraint matrix corresponding with the positive multipliers, $W_{n,B}$, are full-rank.
3. If the conditions checked in step 2 hold, compute $\mathcal{V}^{lo}$ and $\mathcal{V}^{up}$ using the analytical formulas in (14) and (15), replacing Σ by $\widehat{\Sigma}$. Otherwise, $\mathcal{V}^{lo}$ and $\mathcal{V}^{up}$ must be calculated using the definition in (22) and (23) in Appendix A. This can be done using a bisection method, which we describe in Appendix H.
4. Compute the $1 - \alpha$ quantile of the truncated standard normal distribution with truncation points $\mathcal{V}^{lo} / \sqrt{\gamma' \widehat{\Sigma} \gamma}$ and $\mathcal{V}^{up} / \sqrt{\gamma' \widehat{\Sigma} \gamma}$.²³ Reject the null if and only if $\hat{\eta} / \sqrt{\gamma' \widehat{\Sigma} \gamma}$ exceeds this critical value.²⁴

²³In our implementation, we do this via simulation using the method of Botev (2017) to efficiently simulate truncated normal draws. The critical value can also be calculated by inverting a normal CDF, as in equation (16), but we found the former method less prone to numerical precision errors.

²⁴To apply the asymptotic uniformity results developed in Appendix D, here and for the hybrid test below we should reject if and only if $\hat{\eta} / \sqrt{\gamma' \widehat{\Sigma} \gamma}$ exceeds the maximum of this critical value and $-C$, for C a user-selected positive constant.

6.4 Implementing the Hybrid Test

To implement the hybrid test, for $\kappa \in (0, \alpha)$ (we use $\kappa = \alpha/10$ in our simulations),

1. Solve the primal LP (10) using the dual-simplex method, which generates as a byproduct multipliers $\hat{\gamma}$ corresponding to a vertex of the solution set to the dual problem.
2. Compare the resulting value $\hat{\eta}$ to $c_{\kappa, LF}(X_n, \hat{\Sigma}(\beta))$, calculated as described in Section 6.2. If $\hat{\eta}$ exceeds this critical value, reject; otherwise continue the procedure.
3. Follow steps 2 and 3 from the conditional approach to compute $\mathcal{V}^{lo}$ and $\mathcal{V}^{up}$.
4. Compute the $1 - \frac{\alpha - \kappa}{1 - \kappa}$ quantile of the truncated standard normal distribution with lower truncation point $\mathcal{V}^{lo}/\sqrt{\gamma' \hat{\Sigma} \gamma}$ and upper truncation point

$$\mathcal{V}^{up, H}/\sqrt{\gamma' \hat{\Sigma} \gamma} = \min \left(\mathcal{V}^{up}, c_{\kappa, LF}(X_n, \hat{\Sigma}(\beta)) \right) / \sqrt{\gamma' \hat{\Sigma} \gamma}.$$

Reject the null if and only if $\hat{\eta}/\sqrt{\gamma' \hat{\Sigma} \gamma}$ exceeds this critical value.

6.5 Inference with Non-Linear Nuisance Parameters

In some cases, we may have moments of the form

$$g(D_i, \beta_1, \beta_2, \delta) = Y_i(\beta_1, \beta_2) - X_i(\beta_1, \beta_2)\delta$$

and be interested in conducting inference only on β_1 . In this case, we can conduct pointwise inference over a grid for $\beta = (\beta_1, \beta_2)$. We then reject for a particular value of β_1 if and only if for all values of β_2 we reject the hypothesis that (β_1, β_2) is in the identified set (that is, we apply the projection method to eliminate β_2 , while applying the methods developed in this paper to eliminate δ). Alternatively, one could use one of the methods discussed above which can directly address nonlinear parameters.

6.6 Inference on Linear Parameters

In certain applications, we may have linear moments of the form $E_P[Y_i - X_i\delta|Z_i] \leq 0$, where Y_i and X_i do not explicitly depend on a non-linear parameter, and we may be

interested in conducting inference on a linear combination of the parameters, $\beta = l'\delta$ (or $l(X_n)'\delta$). For instance we might be interested in constructing confidence intervals for the coefficient on X_j , in which case we would set $l = e_j$, the vector with a 1 in the j th position and zeros elsewhere. If we did this once for every parameter we would obtain confidence intervals for each of the individual coefficients. Linear combinations of δ may be of interest in other settings as well – e.g., in Wollman (2018) and our Monte Carlo, the average cost of marketing a new product is a linear combination of δ .

We first note that we can recast this problem into the standard form (18) and then use any of the methods described above. To see this, let B be a full rank matrix with l in the first row, so that $B\delta = (\beta, \tilde{\delta})'$ for some $\tilde{\delta}$. If we let M_{-1} be the selection matrix that selects all but the first column of a matrix we have²⁵

$$Y - X\delta = Y - X(B^{-1}B)\delta = (Y - XB^{-1}e_1\beta) - XB^{-1}M_{-1}\tilde{\delta} \equiv \tilde{Y}(\beta) - \tilde{X}\tilde{\delta}.$$

Since $Var_P(Y_i - X_i\delta|Z_i)$ does not depend on δ , Σ need only be estimated once and confidence sets for $l'\delta$ using the LF and LFP methods can be obtained from a linear program (there is no need for point-wise grid test inversion). For example to compute the upper bound of the confidence set for $\beta = l'\delta$ one can solve

$$\begin{aligned} & \max_{\delta} l'\delta \\ \text{subject to } & (Y_{n,j} - X_{n,j}\delta)/\sqrt{\hat{\Sigma}_{jj}} \leq c_{\alpha} \quad \forall j, \text{ where } c_{\alpha} \in \{c_{\alpha,LF}, c_{\alpha,LFP}\}. \end{aligned} \tag{19}$$

So far we have discussed the case without non-linear nuisance parameters, but this approach extends to the case where we are interested in $\beta_1 = l'\delta$ and Y and X depend on the non-linear nuisance parameter β_2 . In this case, one can recast the problem as described above so that the moments can be written as $m((\beta_1, \beta_2), \tilde{\delta})$, and then follow the approach in Section 6.5 for non-linear nuisance parameters. Given our assumption that the conditional covariance matrix does not depend on the linear nuisance parameters, computational shortcuts are still available and confidence intervals can be calculated by running a linear program analogous to (19) for each β_2 and taking the maximum of the resulting values as the final upper bound.

²⁵ $M'_{-1} = [0, I_{k-1}]$ where 0 is the zero vector, and I_{k-1} is the $k - 1$ dimensional identity matrix.

7 Simulations

Our simulations are calibrated to Wollman (2018)'s study of the bailouts of GM and Chryslers' truck divisions. To estimate the effect of the bailouts while allowing product repositioning, Wollmann needs to know the fixed cost of marketing a product. He obtains bounds based on conditional moment inequalities.

We adopt the notation of Example 3 above, so $J_{f,i,t}$ is the set of products marketed by firm f in market i in period t , and $\Delta\pi(J_{f,i,t}, J'_{f,i,t})$ is the difference in expected profits from marketing $J_{f,i,t}$ rather than $J'_{f,i,t}$. $J_{f,i,t} \setminus j$ and $J_{f,i,t} \cup j$ are the sets obtained by deleting and adding product j from the set $J_{f,i,t}$ respectively. Following Wollman (2018), the fixed cost to firm f of marketing product j at time t is $\beta(\delta_{c,f} + \delta_g g_j)$ if the product was marketed last year ($j \in J_{f,i,t-1}$), and $\delta_{c,f} + \delta_g g_j$ otherwise. Here $\delta_{c,f}$ is a per-product cost which is constant across products but may differ across firms, while g_j is the gross weight rating of product j .

If we begin with the case where fixed costs are constant across firms ($\delta_{c,f} = \delta_c$ for all f) and again let $1\{\cdot\}$ denote the indicator function, we obtain four conditional moment inequalities by adding and subtracting one product at a time from the set marketed. For instance, similar to the Example 3 above, if firm f markets product j at both $t - 1$ and t , then for

$$m^1(\theta)_{j,f,i,t} \equiv -[\Delta\pi(J_{f,i,t}, J_{f,i,t} \setminus j) - (\delta_c + \delta_g g_j)\beta] \times 1\{j \in J_{f,i,t}, j \in J_{f,i,t-1}\},$$

we must have $E[m^l(\theta)_{j,f,i,t} | V_{f,i,t}] \leq 0$ for all variables $V_{f,i,t}$ in the firm's information set when time- t production decisions were made, since otherwise the firm would have chosen not to market product j in period t . Analogously, considering products that were marketed at time t but not time $t - 1$ yields moment function

$$m^2(\theta)_{j,f,i,t} \equiv -[\Delta\pi(J_{f,i,t}, J_{f,i,t} \setminus j) - \delta_c - \delta_g g_j] \times 1\{j \in J_{f,i,t}, j \notin J_{f,i,t-1}\},$$

while considering products not marketed at time t yields moment functions

$$m^3(\theta)_{j,f,i,t} \equiv -[\Delta\pi(J_{f,i,t}, J_{f,i,t} \cup j) + (\delta_c + \delta_g g_j)\beta] \times 1\{j \notin J_{f,i,t}, j \in J_{f,i,t-1}\},$$

$$m^4(\theta)_{j,f,i,t} \equiv -[\Delta\pi(J_{f,i,t}, J_{f,i,t} \cup j) + \delta_c + \delta_g g_j] \times 1\{j \notin J_{f,i,t}, j \notin J_{f,i,t-1}\}.$$

If the observed data result from a Nash equilibrium then $E[m^l(\theta)_{j,f,i,t} | V_{f,i,t}] \leq 0$ for

$l \in \{1, 2, 3, 4\}$ and all variables $V_{f,i,t}$ in the firm's information set at the time of the decision.

We obtain two further conditional moment inequalities by considering heavier and lighter models than the firm actually marketed. To state them formally, define

$$J^-(j, f, i, t) \equiv \{j' | g_{j'} < g_j, j' \notin J_{f,i,t}, j' \notin J_{f,i,t-1}\},$$

$$J^+(j, f, i, t) \equiv \{j' | g_{j'} > g_j, j' \notin J_{f,i,t}, j' \notin J_{f,i,t-1}\}.$$

and let

$$m_{j,f,i,t}^5(\theta) \equiv - \left(\frac{\sum_{j' \in J^-(j,f,i,t)} [\Delta\pi(J_{f,i,t}, (J_{f,i,t} \setminus j) \cup j') - \delta_g(g_j - g_{j'})]}{\#J^-(j,f,i,t)} \right) \times 1\{j \in J_{f,i,t}, j \notin J_{f,i,t-1}\},$$

$$m_{j,f,i,t}^6(\theta) \equiv - \left(\frac{\sum_{j' \in J^+(j,f,i,t)} [\Delta\pi(J_{f,i,t}, (J_{f,i,t} \setminus j) \cup j') + \delta_g(g_j - g_{j'})]}{\#J^+(j,f,i,t)} \right) \times 1\{j \in J_{f,i,t}, j \notin J_{f,i,t-1}\}.$$

We calibrate our simulation designs using estimates based on Wollmann's data (for details see Appendix G). In each simulation draw we generate data from a cross-section of 500 independent markets.²⁶ This is substantially larger than the 27 observations used by Wollmann, but allows us to consider specifications with a widely varying number of moments. As in Wollmann, $f \in \{1, \dots, F\}$, and there are nine firms so $F = 9$. To generate data we model the expected and observed profits for firm f from marketing product j in market i in period t , denoted by $\pi_{j,f,i,t}^*$ and $\pi_{j,f,i,t}$ respectively, as

$$\pi_{j,f,i,t}^* = \eta_{j,i,t} + \epsilon_{j,f,i,t}, \quad \text{and} \quad \pi_{j,f,i,t} = \pi_{j,f,i,t}^* + \nu_{j,i,t} + \nu_{j,f,i,t},$$

where the ν terms are mean zero disturbances that arise from expectational and measurement error and the η and ϵ terms represent product-, market-, and firm-specific profit shifters known to the firm when marketing decisions are made. The distributions of these errors are calibrated to match moments in Wollmann's data, as described in Appendix G.²⁷

²⁶The data in Wollman (2018) are a time-series but his variance estimates assume no serial correlation, so we adopt a simulation design consistent with this.

²⁷The terms $\eta_{j,i,t}$ and $\nu_{j,i,t}$ reflect product/market/time "shocks" that are known and unknown to the firms, respectively, when they make their decisions. Shocks of this sort are an important aspect of Wollmann's setting. Note that Wollmann also estimates (point-identified) demand and variable cost parameters in a first step, while for simplicity we treat the variable profits $\pi_{j,f,i,t}$ as known to the econometrician.

The moments used to estimate our model are averages (over markets i) of

$$\frac{1}{J} \sum_j \left(m_{j,f,i}^l(\theta) \otimes \tilde{Z}_{j,f,i} \right)', \quad (20)$$

where we also average over all firms assumed to share the same fixed cost $\delta_{f,c}$. Since we consider a single cross-section of markets we suppress the time subscript. We present results both for the case where $\tilde{Z}_{j,f,i}$ includes only a constant and for the case where the last two moments are interacted with a constant but the first four moments are interacted with both a constant and the common profit-shifters η ,

$$\tilde{Z}_{j,f,i} = (1, \eta_{j,i}^+, \eta_{j,i}^-),$$

for $q^+ = \max\{q, 0\}$ and $q^- = -\min\{q, 0\}$. In the model with a single constant term, $\delta_{c,f} = \delta_c$ for all f , this generates 6 and 14 moment inequalities. We also present results when the nine firms are divided into three groups each with a separate constant term, and when each firm has a separate constant term. For each specification we consider the first four moments separately for the firm(s) associated with distinct parameters $\delta_{c,f}$, but average the last two moments across all firms as they do not depend on the constant terms. This generates 14 and 38 moments for the three group classification, and 38 and 110 moments when each firm has a separate constant term. To estimate the conditional variance Σ , in each specification we define the value of the instrument Z_i in market i as the Jacobian of (20) with respect to the linear parameters $(\delta_g, \{\delta_{c,f}\})$.

We consider inference on three parameters of interest: the cost of marketing the truck of mean weight when it was marketed in the prior year;²⁸ the incremental cost of changing the weight of a product, δ_g ; and the non-linear parameter β , where $1 - \beta$ represents the proportional cost savings from marketing a product that was previously marketed relative to a new product. For the first two parameters, each of which can be written as a linear combination of the vector δ , we hold β fixed at its true value to allow us to examine performance in the linear case discussed in Section 6.6. As discussed in Section 6.5, if we instead treated β as unknown we could form joint confidence sets

²⁸When we assume $\delta_{c,f}$ is common across firms this is $\delta_c + \delta_g \mu_g$, where μ_g is the population average weight of trucks. When we allow the estimated δ_c parameters to vary across groups, we estimate $l'\delta$, for $l = (\frac{1}{G}, \dots, \frac{1}{G}, \mu_g)'$, where G denotes the number of groups and $\delta = (\delta_{c,1}, \dots, \delta_{c,G}, \delta_g)'$. Note that since the simulation DGP holds the true value of δ_c constant across groups, the true value of the parameter is the same in all specifications.

for β along with the linear combination of interest, and could form confidence sets for the linear parameter alone by projection. For inference on β we treat the entire vector δ as a nuisance parameter. All results are based on 500 simulation runs.

We begin our discussion of the results with Figure 1, which shows rejection probabilities for the cost of the mean-weight truck.²⁹ The vertical dashed lines denote the conservative estimates for the bounds of the identified set, and the four curves represent the probability that each of the four methods considered rejects a given null value of the parameter of interest.³⁰ There is a clear ranking of the power of the LFP, LF, and Hybrid procedures in Figure 1. In all specifications, the LF test has noticeably higher power than the LFP. The hybrid test has power comparable to or above the LF test in all specifications, with substantial differences emerging in cases with a larger number of moments and parameters. The performance of the conditional test is more nuanced. When the number of moments per parameter is small, the conditional test performs very similarly to the hybrid, and is at least as good as the LF and LFP. When we increase the number of moments holding the number of parameters fixed, the conditional again performs similarly to the hybrid for parameter values close to the identified set bounds, but can have power substantially below any of the other methods far away from the identified set (see for instance Panel (d) of Figure 1).

The power declines for the conditional test reflect that the set of binding moments is not well-separated in this example. In particular, one can show that in this simulation design, when the parameter of interest is the cost of the mean-weight truck we have a multiplicity of solutions to the population moments when the number of moments per parameter is large. As a result, we often have multiple near-solutions to the linear program (10) in sample.³¹ As noted in Section 5, the conditional test may perform poorly in such settings, and this prediction is borne out in this application. Our hybrid test eliminates these problems, as intended.

²⁹Note that all of our simulation results in this section hold the data generating process constant but vary the parameter values considered. Hence, the curves plotted should be interpreted as rejection probabilities for tests of different null hypotheses, or one minus the coverage probability for confidence sets.

³⁰We cannot solve for the true identified set analytically, so we approximate it by the set satisfying the sample (unconditional) moment inequalities based on a simulation run with five million observations. To ensure that our estimate of the identified set is conservative, we follow Chernozhukov et al. (2007) and add a correction factor to the moments of $\log(n)/\sqrt{n} \approx .003$ when $n=5,000,000$. Hence, our estimate of the identified is conservative in these simulations due to both (a) the Chernozhukov et al. (2007) correction factor and (b) the use of unconditional rather than conditional moment inequalities.

³¹Less frequently, we have multiple exact solutions, in which case we apply the dual approach.

Figure 2 reports rejection probabilities for testing hypotheses on the nonlinear parameter β . Unlike in our simulations for the linear parameters, when testing nonlinear parameters it is sometimes the case that no procedure has rejection probability going to one over the grid we consider, though this phenomenon disappears in all but the conditional power curves when we interact the conditional moments with the profit shifters $(\eta_{j,i}^+, \eta_{j,i}^-)$. Regardless, we see that the LF test has higher power than the LFP, and that the power of the hybrid test is higher still. The conditional test performs reasonably well in cases with a small number of moments and parameters (e.g. in Panel (a)) but it has power well below any of the other tests considered at many parameter values in some cases with more moments and/or parameters.

Rejection probabilities for testing hypotheses on δ_g are similar to those for testing the cost of the average weight truck, though with better performance for the conditional test, and so are reported in Appendix G to conserve space. One notable feature of these results is that the identified set for δ_g does not change across specifications, so unlike for our analysis of the other parameters, the specifications with more than six moments are adding moments and nuisance parameters without changing the identified set. The results in this case confirm that the hybrid approach appears less sensitive to the addition of parameters and slack moments than the LF or LFP.

Table 1 reports the size (formally, the maximal null rejection probability over the estimated identified set) for all the tests considered. As expected all tests approximately control size, with the maximal null rejection probabilities for nominal 5% tests bounded above by 8%, and this bound is reached only in cases with 110 moments.³² Our estimates for the identified set are conservative, so those rejection probabilities should, if anything, overestimate the true maximal rejection probability.

Table 1 also reports the median excess length of each confidence set, defined as the difference between the length of the confidence interval and that of the identified set. This provides a summary measure of the extent to which the confidence interval is longer than the identified set. The ranking of confidence sets based on median excess length agrees with the ranking from Figures 1 and 2, with the median excess length for the hybrid comparable to or below that for the LF, which is strictly below that

³²We also ran simulations defining the identified set without the conservative Chernozhukov et al. (2007) correction factor, and the only designs for which this resulted in a difference of maximal rejection probabilities of more than 0.01 were two of the runs with 110 moments, where the bounds with the correction implied probabilities of 0.07 and 0.08, compared to 0.02 and 0.01 without the correction.

Figure 1: Rejection probabilities for 5% tests of fixed cost for truck of mean weight




Figure 2: Rejection probabilities for 5% tests of β




Table 1: Median Excess Length and Size

(a) Parameter: Cost of Mean-Weight Truck

#Parameters	#Moments	Median Excess Length				Max Size			
		LFP	LF	Cond.	Hybrid	LFP	LF	Cond.	Hybrid
2	6	5.32	3.99	4.07	3.75	0	0.02	0.02	0.02
2	14	12.75	10.48	10.49	8.54	0	0	0.02	0.02
4	14	7.56	5.91	4.07	4.37	0	0	0.02	0.02
4	38	19.08	16.33	14.68	11.60	0	0	0.03	0.03
10	38	12.70	10.20	4.71	4.71	0	0	0.02	0.01
10	110	25.61	22.36	17.91	14.34	0	0	0.08	0.07

(b) Parameter: δ_g

#Parameters	#Moments	Median Excess Length				Max Size			
		LFP	LF	Cond.	Hybrid	LFP	LF	Cond.	Hybrid
2	6	6.02	4.28	4.18	3.93	0.01	0.04	0.05	0.06
2	14	6.91	5.40	4.43	4.18	0	0.02	0.05	0.05
4	14	7	5.19	4.43	4.18	0	0.03	0.05	0.06
4	38	7.97	6.68	4.43	4.43	0	0	0.05	0.05
10	38	8.10	6.58	4.43	4.43	0	0.01	0.05	0.05
10	110	9.11	7.69	5.18	5.18	0	0	0.03	0.03

(c) Parameter: β

#Parameters	#Moments	Median Excess Length				Max Size			
		LFP	LF	Cond.	Hybrid	LFP	LF	Cond.	Hybrid
3	6	118.69	61.87	41.67	36.62	0	0	0	0
3	14	0.76	0.56	0.45	0.35	0	0	0.01	0.01
5	14	10.25	7.78	6.01	5.3	0	0	0.01	0.01
5	38	0.86	0.66	0.96	0.45	0	0	0.03	0.03
11	38	1.41	1.01	1.01	0.81	0	0	0.01	0.01
11	110	0.86	0.66	2.57	0.56	0	0	0.05	0.04

of the LFP. The conditional performs comparably to the hybrid in cases with a low number of moments per parameter, but sometimes performs distinctly worse when the number of moments per parameter is higher. We report results for other quantiles of the excess length distribution in Appendix G.

Lastly, Table 2 reports runtimes in minutes to calculate confidence sets for each parameter. Notably, all runtimes, even those for the non-linear parameter when there are eleven parameters in total, are well within acceptable limits for most empirical projects.

A few comparisons between the procedures are worth noting. When conducting inference on the linear parameters, the LF and LFP procedures are substantially faster than the hybrid and conditional approaches. This is because confidence intervals for the former can be computed using linear programming, as described in Section 6.6, whereas the latter approaches rely on test inversion over a grid. All procedures become slower when conducting inference on a non-linear parameter, since they all rely on an estimate of the conditional covariance matrix, which now needs to be computed at each grid point; additionally, the LF and LFP now rely on test inversion over a grid as well. For the non-linear parameter, the LFP and conditional approaches are typically faster than the LF and hybrid, since the former need only calculate one linear program for each grid point, whereas the latter methods require simulating the results of the linear program many times for each grid point.³³ We stress, however, that at least for the simulation designs we consider, all four procedures remain highly tractable, and runtimes could be improved using parallelization.

8 Conclusion

This paper considers the problem of inference based on linear conditional moment inequalities, which arise in a wide variety of economic applications. Using linear conditional structure, we develop inference procedures which remain both computationally tractable and powerful in the presence of nuisance parameters, including conditional and hybrid procedures which are insensitive to the presence of slack moments. We find good performance for our least favorable, conditional, and hybrid procedures under a variety of simulation designs based on Wollman (2018), with especially good

³³If computation times are an issue for the hybrid, the LF first stage can be replaced with a LFP first stage, yielding a faster but somewhat less powerful test.

Table 2: Computation Times

(a) Parameter: Cost of Mean-Weight Truck

		Runtime - Minutes			
#Parameters	#Moments	LFP	LF	Conditional	Hybrid
2	6	0.11	0.18	0.44	0.41
2	14	0.05	0.13	5.44	1.67
4	14	0.08	0.15	6.56	2.12
4	38	0.07	0.15	5.06	1.3
10	38	0.13	0.22	4.77	0.99
10	110	0.13	0.21	4.47	0.92

(b) Parameter: δ_g

		Runtime - Minutes			
#Parameters	#Moments	LFP	LF	Conditional	Hybrid
2	6	0.08	0.16	6.69	2.54
2	14	0.05	0.13	5.2	2.14
4	14	0.08	0.15	6.01	1.96
4	38	0.07	0.15	4.78	1.63
10	38	0.13	0.22	3.77	1.03
10	110	0.13	0.21	3.87	1.04

(c) Parameter: β

		Runtime - Minutes			
#Parameters	#Moments	LFP	LF	Conditional	Hybrid
3	6	5.24	13.68	5.27	13.69
3	14	5.65	14.02	5.68	14.02
5	14	6.9	15.36	6.93	15.37
5	38	8.16	17.12	8.17	17.12
11	38	11.94	20.95	11.95	20.96
11	110	16.78	32.95	16.74	32.96

This table shows runtimes to calculate confidence sets based on one simulated dataset for each specification, without parallelization, on a 2014 Macbook Pro with a 2.6 GHz Intel i5 Processor and 16GB of RAM. For the linear parameters (Panels a and b), the confidence sets for the LF and LFP are computed using linear programming, as described in Section 6.6, and we use a grid of 1,001 parameter values for the hybrid and conditional approaches. For the non-linear parameter β , all four procedures use a grid of length 100. See Appendix G for additional details on the simulation specification.

performance for the hybrid.

References

- Abadie, A., Imbens, G. W. & Zheng, F. (2014), ‘Inference for misspecified models with fixed regressors’, *Journal of the American Statistical Association* **109**(508), 1601–1614.
- Andrews, D. W. & Barwick, P. J. (2012), ‘Inference for parameters defined by moment inequalities: A recommended moment selection procedure’, *Econometrica* **80**(6), 2805–2826.
- Andrews, D. W. & Guggenberger, P. (2009), ‘Hybrid and size-corrected subsampling methods’, *Econometrica* **77**(3), 721–762.
- Andrews, D. W. & Shi, X. (2013), ‘Inference based on conditional moment inequalities’, *Econometrica* **81**(2), 609–666.
- Andrews, D. W. & Soares, G. (2010), ‘Inference for parameters defined by moment inequalities using generalized moment selection’, *Econometrica* **78**(1), 119–159.
- Andrews, I., Kitagawa, T. & McCloskey, A. (2018), Inference on winners. Working Paper.
- Armstrong, T. B. (2014a), A note on minimax testing and confidence intervals in moment inequality models. Working Paper.
- Armstrong, T. B. (2014b), ‘Weighted ks statistics for inference on conditional moment inequalities’, *Journal of Econometrics* **181**(2), 92–116.
- Barseghyan, L., Coughlin, M., Molinari, F. & Teitelbaum, J. C. (2019), Heterogenous choice sets and preferences. Working Paper.
- Belloni, A., Bugni, F. & Chernozhukov, V. (2018), Subvector inference in PI models with many moment inequalities. Working Paper.
- Beresteanu, A. & Molinari, F. (2008), ‘Asymptotic properties for a class of partially identified models’, *Econometrica* **76**(4), 763–814.

- Blundell, R., Gosling, A., Ichimura, H. & Meghir, C. (2007), ‘Changes in the distribution of male and female wages accounting for employment composition’, *Econometrica* **75**, 323–363.
- Bontemps, C., Magnac, T. & Maurin, E. (2012), ‘Set identified linear models’, *Econometrica* **80**(3), 1129–1155.
- Botev, Z. I. (2017), ‘The normal law under linear restrictions: simulation and estimation via minimax tilting’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **79**(1), 125–148.
- Bugni, F., Canay, I. & Shi, X. (2017), ‘Inference for subvectors and other functions of partially identified parameters in moment inequality models’, *Quantitative Economics* **8**(1), 1–38.
- Canay, I. & Shaikh, A. (2017), Practical and theoretical advances in inference for partially identified models, in B. Honore, A. Pakes, M. Piazzesi, & L. Samuelson, eds, ‘Advances in Economics and Econometrics’, Cambridge University Press.
- Chen, X., Christensen, T. & Tamer, E. (2018), ‘Monte carlo confidence sets for identified sets’, *Econometrica* **86**(6), 1965–2018.
- Chernozhukov, V., Hong, H. & Tamer, E. (2007), ‘Estimation and confidence regions for parameter sets in econometric models’, *Econometrica* **75**(5), 1243–1284.
- Chernozhukov, V., Newey, W. & Santos, A. (2015), Constrained conditional moment restriction models. Working Paper.
- Chetty, R. (2012), ‘Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply’, *Econometrica* **80**(3), 969–1018.
- Chetverikov, D. (2018), ‘Adaptive test of conditional moment inequalities’, *Econometric Theory* **34**(1), 186–227.
- Cho, J. & Russell, T. M. (2019), Simple inference on functionals of set-identified parameters defined by linear moments. Working Paper.
- Ciliberto, F. & Tamer, E. (2009), ‘Market structure and multiple equilibria in airline markets’, *Econometrica* **77**(6), 1791–1828.

- Cox, G. & Shi, X. (2019), A simple uniformly valid test for inequalities. Working Paper.
- Dickstein, M. & Morales, E. (2018), ‘What do exporters know?’, *Quarterly Journal of Economics* **133**(4), 1753–1801.
- Eizenberg, A. (2014), ‘Upstream innovation and product variety in the U.S. home pc market’, *Review of Economic Studies* **81**(3), 1003–1045.
- Flynn, Z. (2019), Inference based on continuous linear inequalities via semi-infinite programming. Working Paper.
- Gafarov, B. (2019), Inference in high-dimensional set-identified affine models. Working Paper.
- Haile, P. A. & Tamer, E. (2003), ‘Inference with an incomplete model of english auctions’, *Journal of Political Economy* **111**(1), 1–51.
- Ho, K. & Pakes, A. (2014), ‘Hospital choices, hospital prices and financial incentives to physicians’, *American Economic Review* **104**(12), 3841–84.
- Ho, K. & Rosen, A. (2017), Partial identification in applied research, *in* B. Honore, A. Pakes, M. Piazzesi, & L. Samuelson, eds, ‘Advances in Economics and Econometrics’, Cambridge University Press.
- Holmes, T. (2011), ‘The diffusion of Wal-Mart and economies of density’, *Econometrica* **79**(1), 253–302.
- Honore, B. & Lleras-Muney, A. (2006), ‘Bounds in competing risks models and the war on cancer’, *Econometrica* **74**(6), 1675–1698.
- Hsieh, Y.-W., Shi, X. & Shum, M. (2017), Inference on estimators defined by mathematical programming. Working Paper.
- Kaido, H., Molinari, F. & Stoye, J. (2019a), ‘Confidence intervals for projections of partially identified parameters’, *Econometrica* **87**(4), 1397–1432.
- Kaido, H., Molinari, F. & Stoye, J. (2019b), Online appendix to "confidence intervals for projections of partially identified parameters ". Supplementary Material.

- Katz, M. (2007), Supermarkets and zoning laws. Ph.D. dissertation, Harvard University.
- Khan, S., Ponomareva, M. & Tamer, E. (2019), Identification of dynamic panel binary response models. Working Paper.
- Kline, P. & Tartari, M. (2016), ‘Bounding the labor supply response to a randomized welfare experiment: A revealed preference approach’, *American Economic Review* **106**(4), 972–1014.
- Kreider, B., Pepper, J., Gundersen, C. & Jolliffe, D. (2012), ‘Identifying the effects of snap (food stamps) on child health outcomes when participation is endogenous and misreported’, *Journal of the American Statistical Association* **107**, 958–975.
- Lee, J. D., Sun, D. L., Sun, Y. & Taylor, J. E. (2016), ‘Exact post-selection inference, with application to the lasso’, *Annals of Statistics* **44**(3), 907–927.
- Manski, C. F. & Tamer, E. (2002), ‘Inference on regressions with interval data on a regressor or outcome’, *Econometrica* **70**(2), 519–546.
- Mogstad, M., Santos, A. & Torgovitsky, A. (2018), ‘Using instrumental variables for inference about policy relevant treatment parameters’, *Econometrica* **86**(5), 1589–1619.
- Molinari, F. (2019), Econometrics with partial identification. Working Paper.
- Morales, E., Sheu, G. & Zahler, A. (2019), ‘Extended gravity’, *Review of Economic Studies* **Forthcoming**.
- Pakes, A. (2010), ‘Alternative models for moment inequalities’, *Econometrica* **78**(6), 1783–1822.
- Pakes, A., Porter, J. & Ishii, K. H. J. (2015), ‘Moment inequalities and their application’, *Econometrica* **83**(1), 315–334.
- Ponomareva, M. & Tamer, E. (2011), ‘Misspecification in moment inequality models: Back to moment equalities?’, *Econometrics Journal* **14**(2), 186–203.
- Romano, J. P. & Shaikh, A. (2008), ‘Inference for identifiable parameters in partially identified econometric models’, *Journal of Statistical Planning and Inference* **138**(9), 2786–2807.

- Romano, J. P., Shaikh, A. & Wolf, M. (2014a), ‘A practical two-step method for testing moment inequalities’, *Econometrica* **82**(5), 1979–2002.
- Romano, J. P., Shaikh, A. & Wolf, M. (2014b), ‘Supplement to ‘a practical two-step method for testing moment inequalities’’, *Econometrica* **82**.
- Rosen, A. (2008), ‘Confidence sets for partially identified parameters that satisfy a finite number of moment inequalities’, *Journal of Econometrics* **146**(1), 107–117.
- Schrijver, A. (1986), *Theory of Linear and Integer Programming*, Wiley-Interscience.
- Tebaldi, P., Torgovitsky, A. & Yang, H. (2019), Nonparametric estimates of demand in the california health insurance exchange. Working Paper.
- Torgovitsky, A. (2019), ‘Partial identification by extending subdistributions’, *Quantitative Economics* **10**(1), 105–144.
- Wollman, T. (2018), ‘Trucks without bailouts: Equilibrium product characteristics for commercial vehicles’, *American Economic Review* **108**(6), 1364–1406.