

Exact Testing of Many Moment Inequalities Against Multiple Violations

Nick Koning^a and Paul Bekker^b

Faculty of Economics and Business

University of Groningen

May 27, 2020

Abstract

This paper considers the problem of testing many moment inequalities, where the number of moment inequalities (p) is possibly larger than the sample size (n). Chernozhukov et al. (2019) proposed asymptotic tests for this problem using the maximum t statistic. We observe that such tests can have low power if multiple inequalities are violated. As an alternative, we propose novel randomization tests based on a maximum non-negatively weighted combination of t statistics. We provide a condition guaranteeing size control in large samples. Simulations show that the tests control size in small samples ($n = 30$, $p = 1000$), and often has substantially higher power against alternatives with multiple violations than tests based on the maximum t statistic.

Key words and phrases: Many moment inequalities; High-dimensional inference; Randomization test; Symmetry-based inference.

JEL classification: C12, C14, C55.

^aCorresponding author. Nick Koning, Faculty of Economics and Business, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands. Email: n.w.koning@rug.nl.

^bPaul Bekker, Faculty of Economics and Business, University of Groningen, Postbus 800, 9700 AV Groningen, The Netherlands. Email: p.a.bekker@rug.nl.

1 Introduction

As discussed by Chernozhukov et al. (2019), henceforth CCK, the moments inequalities framework has developed into a powerful tool for inference on causal and structural parameters in partially identified models. In such models, the parameters of interest may be restricted to a subset of the parameter space defined by a collection of moment inequalities. The simultaneous testing of these moment inequalities provides inference about the true underlying parameter values. CCK provide an excellent review of the literature with detailed motivating examples.

They point out that many economic models give rise to problems where the number of moment inequalities p may be much larger than the number of observations n . While there exists a large literature on testing moment inequalities, traditional methods are not well equipped for dealing with *many* moment inequalities.¹ In order to test the moment inequalities against the alternative where at least one of them is violated, CCK use the maximum of p t values as a test statistic. They find critical values by using asymptotic theory and bootstrap methods. Additionally, a first-stage inequality selection step is included to improve the power of their tests. Allen (2018) suggests a refinement of the selection step. Bugni et al. (2016) consider a generalization of the same problem and use Lasso for the first-stage selection.

Our contribution to the framework of testing many moment inequalities is threefold. First, we propose two novel test statistics. Notice that the maximum of p t values is invariant to the size of the second largest t value. If the alternative hypothesis allows for no more than one violation, inference based on the maximum may be powerful, but it would discard power against alternatives where multiple moment inequalities are violated.

In order to retain this power, we propose an enhanced maximum statistic, where we

¹For work on unconditional moment inequalities see, e.g., Canay (2010); Andrews and Barwick (2012); Andrews and Guggenberger (2009); Chernozhukov et al. (2007); Rosen (2008); Romano and Shaikh (2008). Andrews and Shi (2013) notice that conditional moment inequalities can be viewed as an infinite number of unconditional moment inequalities. Contributions to conditional moment inequalities are found in Chernozhukov et al. (2013); Lee et al. (2013, 2018); Armstrong (2014, 2015); Armstrong and Chan (2016).

maximize over an artificially expanded set of t values. This expanded set of t values is constructed by adding an extra set of t values to the set of p t values. Such extra t values are created by taking a non-negatively weighted combination of sample means divided by the square-root of its combined sample variance. This results in a test statistic that is also large if such a weighted combination is large, and not just if any of the original t values is large. Interestingly, this expanded set of t values is still a set of t values (though correlated by construction). Therefore, the results of CCK can also be applied to this new test statistic, if the set of extra t values is finite.

In addition, we consider adding all non-negative combinations, so that the set of extra t values is infinite. This results in a test statistic that can be viewed as a one-sided version of Hotelling's T^2 statistic. An interesting feature of our statistic is that, unlike the Hotelling's T^2 statistic, it can also be used in high dimensional settings. This is because its existence condition is weaker than non-singularity of the sample covariance matrix (Koning, 2019). This condition is related to the positive eigenvalue condition (Meinshausen et al., 2013; Slawski et al., 2013) or copositivity of the sample covariance matrix. Like the positive eigenvalue condition, it may be satisfied if many elements of the sample covariance matrix are positive. However, it is also satisfied if many of the sample means are negative. In a simulation study, we find substantial increases in power compared to the statistic proposed by CCK, against alternatives where multiple moment inequalities are violated. However, a limitation of the test seems to be that it loses power against alternatives where the number of not satisfied moment inequalities is large compared to n and the underlying covariance matrix exhibits weak correlations, as this interferes with the existence condition.

Secondly, we propose using randomization tests in the moment inequalities framework, by imposing a symmetry assumption on the errors. Randomization tests originate with Fisher (1935) and have been widely used in the literature.² An advantage of randomization tests is that they are exact if all moment inequalities are satisfied. In addition, we provide

²See, for example, Lehmann and Romano (2006), Maritz (1995), Romano (1990), Bekker (2002) and Bekker and Lawford (2008)

a sufficient condition to guarantee our tests control size for sufficiently large samples. This result also holds for the statistic that is the maximum over an infinite set of t values.

In our simulation experiments we find that the randomization tests perform similarly or slightly better compared to the empirical bootstrap proposed by CCK in large samples, and better in small samples. While randomization tests work well in small samples, they do require the validity of a symmetry assumption. However, in the simulation experiments where the symmetry assumption does not hold, we find, to our surprise, that the randomization test has better control of size than the empirical bootstrap procedure proposed by CCK, even if the sample is large.

Finally, we describe a first stage selection step in order to improve power by eliminating non-binding moment inequalities. The approach is based on symmetry and it does not depend on a specific selection method. Simulation results show that the tests with selection control size, and have increased power in the presence of non-binding moment inequalities compared to tests without selection.

2 The model and two test statistics

Let observations in the $n \times p$ matrix \mathbf{X} satisfy

$$\mathbf{X} = \boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E}, \tag{1}$$

where $\boldsymbol{\nu}_n$ is an n -vector of ones, $\boldsymbol{\mu}$ is a p -vector of parameters and \mathbf{E} is a random matrix that satisfies a symmetry assumption. Specifically, we assume that any row of \mathbf{E} can be multiplied by -1 without affecting the distribution of \mathbf{E} . This assumption is satisfied if the rows of \mathbf{E} are independently symmetrically distributed about zero. Moments need not exist and the rows of \mathbf{E} need not be identically distributed. We describe exact inference based on the finite sample.

Following CCK, we are interested in testing the hypothesis $H_0: \boldsymbol{\mu} \leq \mathbf{0}$ against the

alternative $H_1: \boldsymbol{\mu} \not\leq \mathbf{0}$. That is, we want to test whether all elements of $\boldsymbol{\mu}$ are non-positive, against the alternative that at least one is positive. We will refer to the null hypothesis as the *moment inequalities*. The j th moment inequality is said to be *violated* if $\mu_j > 0$, it is *binding* if $\mu_j = 0$, and it is *strictly satisfied* if $\mu_j < 0$.

To define the test statistics, let $\mathbf{I}_n = (\mathbf{e}_1, \dots, \mathbf{e}_n)$ be the $n \times n$ identity matrix, and define $\mathbf{P}_{\boldsymbol{\nu}_n} = \boldsymbol{\nu}_n \boldsymbol{\nu}_n' / \boldsymbol{\nu}_n' \boldsymbol{\nu}_n$. Define $\hat{\boldsymbol{\mu}} = \mathbf{X}' \boldsymbol{\nu}_n / n$ with elements $\hat{\mu}_j$, and $\hat{\boldsymbol{\Sigma}} = \mathbf{X}' (\mathbf{I}_n - \mathbf{P}_{\boldsymbol{\nu}_n}) \mathbf{X} / n$ with diagonal elements $\hat{\sigma}_j^2$. We assume $\hat{\boldsymbol{\mu}} \not\leq \mathbf{0}$, otherwise we would not reject H_0 , and consider test statistics of the form

$$T_{\mathcal{U}} = n^{1/2} \max_{\boldsymbol{\lambda} \in \mathcal{U}} \frac{\hat{\boldsymbol{\mu}}' \boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}' \hat{\boldsymbol{\Sigma}} \boldsymbol{\lambda}}},$$

where $\mathcal{U} \subseteq \mathbb{R}^p$ and Condition 1 holds (Koning, 2019).

Condition 1. $\boldsymbol{\lambda}' \hat{\boldsymbol{\Sigma}} \boldsymbol{\lambda} > 0$, for all $\boldsymbol{\lambda} \in \mathcal{U}$, $\boldsymbol{\lambda} \neq \mathbf{0}$, $\hat{\boldsymbol{\mu}}' \boldsymbol{\lambda} > 0$.

The test statistic has been used before for particular choices of \mathcal{U} . Here we propose new choices as well.

In order to test H_0 against H_1 , CCK use the maximum t value, $t_j = n^{1/2} \hat{\mu}_j / \hat{\sigma}_j$, for $j = 1, \dots, p$. We will denote this statistic by $t_{\max} = t_{\max}(\mathbf{X}) = \max_{1 \leq j \leq p} t_j = T_{\mathcal{C}^p}$, where $\mathcal{C}^p = \{\mathbf{e}_1, \dots, \mathbf{e}_p\}$. By ignoring the sizes of all but the largest t value, the t_{\max} statistic may perform well against alternatives with one or few violations, but lack power when testing against alternatives with multiple violations. To shift power towards multiple violations, one might expand \mathcal{U} by adding vectors from the non-negative orthant, $\mathbf{a}_i \in \mathbb{R}_+^p$, $i = 1, \dots, m$. Let $\mathcal{U} = \mathcal{C}^p \cup \{\mathbf{a}_1, \dots, \mathbf{a}_m\}$, then $T_{\mathcal{U}}$ may outperform t_{\max} in terms of power under alternatives with multiple violations.

Remark 1. *Interestingly, if \mathcal{U} is finite, the $T_{\mathcal{U}}$ statistic fits entirely within the framework of CCK. This implies that their results regarding the bootstrap and self-normalized methods can be immediately applied if the suitable regularity conditions hold.*

Remark 2. If the off-diagonals of $\widehat{\Sigma}$ are positive, expanding the set \mathcal{U} may be less effective. In the extreme case that $\widehat{\Sigma} = \boldsymbol{\nu}_p \boldsymbol{\nu}'_p$, the $T_{\mathcal{U}}$ and t_{\max} statistics coincide if $\mathcal{U} \subset \mathbb{R}_+^p$ is finite and $\mathcal{C}^p \subset \mathcal{U}$. That is to say, $T_{\mathcal{U}} \geq t_{\max}$ by construction, and

$$n^{-1/2}T_{\mathcal{U}} = \max_{\boldsymbol{\lambda} \in \mathcal{U}} \frac{\widehat{\boldsymbol{\mu}}' \boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}' \widehat{\Sigma} \boldsymbol{\lambda}}} = \max_{\boldsymbol{\lambda} \in \mathcal{U}} \frac{\widehat{\boldsymbol{\mu}}' \boldsymbol{\lambda}}{\boldsymbol{\nu}'_p \boldsymbol{\lambda}} = \max_{\substack{\boldsymbol{\lambda} \in \mathcal{U} \\ \boldsymbol{\nu}'_p \boldsymbol{\lambda} = 1}} \widehat{\boldsymbol{\mu}}' \boldsymbol{\lambda} \leq \max_{\substack{\boldsymbol{\lambda} \geq \mathbf{0} \\ \boldsymbol{\nu}'_p \boldsymbol{\lambda} = 1}} \widehat{\boldsymbol{\mu}}' \boldsymbol{\lambda} = \max_{\boldsymbol{\lambda} \in \mathcal{C}^p} \widehat{\boldsymbol{\mu}}' \boldsymbol{\lambda} = n^{-1/2}t_{\max}.$$

Alternatively, we consider infinite sets \mathcal{U} . In particular, we use $\mathcal{U} = \mathbb{R}_+^p$ and define $T_+ = T_{\mathbb{R}_+^p}$.³ The T_+ statistic can be viewed as a one-sided version of the square-root of Hotelling's T^2 -statistic, since $T_{\mathbb{R}^p}^2 = T^2 = n \widehat{\boldsymbol{\mu}}' \widehat{\Sigma}^{-1} \widehat{\boldsymbol{\mu}}$, which requires $\widehat{\Sigma}$ to be invertible. For the T_+ statistic $\widehat{\Sigma}$ need not be invertible. Instead, a weaker condition suffices, which is a special case of Condition 1.

Condition 2. $\boldsymbol{\lambda}' \widehat{\Sigma} \boldsymbol{\lambda} > 0$, for all $\boldsymbol{\lambda} \geq \mathbf{0}$, $\boldsymbol{\lambda} \neq \mathbf{0}$, $\widehat{\boldsymbol{\mu}}' \boldsymbol{\lambda} > 0$.

This condition is weaker than (strict) copositivity of $\widehat{\Sigma}$, which requires that $\boldsymbol{\lambda}' \widehat{\Sigma} \boldsymbol{\lambda} > 0$ for all $\boldsymbol{\lambda} \geq \mathbf{0}$, $\boldsymbol{\lambda} \neq \mathbf{0}$ (see e.g. Hiriart-Urruty and Seeger, 2010), or the equivalent positive eigenvalue condition formulated by Meinshausen et al. (2013).⁴

An example where copositivity holds is if all elements of $\widehat{\Sigma}$ are strictly positive (i.e. $\widehat{\Sigma}$ is *positive*). Meinshausen et al. (2013) also provides an example where $\widehat{\Sigma}$ is allowed to contain some negative entries. In particular, if \mathcal{A}^c contains the indices of the the largest positive principal submatrix of $\widehat{\Sigma}$ and \mathcal{A} is its complement, then the condition only has to hold on the principal submatrix of $\widehat{\Sigma}$ with indices in \mathcal{A} . If the number of elements in \mathcal{A} is small compared to n , this need not be very restrictive.

The difference between Condition 2 and copositivity lies in the role played by $\widehat{\boldsymbol{\mu}}$. Condition 2 is equivalent to copositivity if $\widehat{\boldsymbol{\mu}} > \mathbf{0}$. At the other extreme, if $\widehat{\boldsymbol{\mu}} \leq \mathbf{0}$ the set

³To handle cases with multiple violations, CCK mention the sum statistic $\sum_{i=1}^p (\sqrt{n} \max\{\widehat{\mu}_i / \widehat{\sigma}_i, 0\})^2$ and the quasi likelihood ratio test statistic $\min_{\boldsymbol{\mu} \leq \mathbf{0}} n (\widehat{\boldsymbol{\mu}} - \boldsymbol{\mu})' \widehat{\Sigma}^{-1} (\widehat{\boldsymbol{\mu}} - \boldsymbol{\mu})$ as possibilities. Only if $\widehat{\Sigma}$ is diagonal, these statistics fit within the $T_{\mathcal{U}}$ framework. In that case the statistics are equal to T_+ . However, the sum statistic ignores the covariances, and the quasi likelihood-ratio statistic requires $\widehat{\Sigma}$ to be invertible.

⁴The positive eigenvalue condition $\min_{\boldsymbol{\lambda} \geq \mathbf{0}, \boldsymbol{\lambda} \neq \mathbf{0}} \boldsymbol{\lambda}' \widehat{\Sigma} \boldsymbol{\lambda} / \|\boldsymbol{\lambda}\|_1^2 > 0$ is equivalent to $\boldsymbol{\lambda}' \widehat{\Sigma} \boldsymbol{\lambda} > 0$, for all $\boldsymbol{\lambda} \geq \mathbf{0}$, $\boldsymbol{\lambda} \neq \mathbf{0}$.

$\{\boldsymbol{\lambda} \mid \boldsymbol{\lambda} \geq \mathbf{0}, \boldsymbol{\lambda} \neq \mathbf{0}, \boldsymbol{\mu}'\boldsymbol{\lambda} \geq 0\}$ is empty, so that Condition 2 holds, whether or not copositivity holds. If the elements of $\hat{\boldsymbol{\mu}}$ have comparable positive or negative sizes, then Condition 2 tends to become less restrictive the more negative elements there are. So, intuitively, even if $\hat{\boldsymbol{\Sigma}}$ is not copositive, Condition 2 may still hold, in particular if many moment inequalities are strictly satisfied. This reasoning is confirmed in the simulation experiments, where the test based on T_+ performs well even when $n = 30$ and $p = 1000$, and only 10 moment inequalities are violated.

To compute T_+ we use a recent result by Koning (2019), who shows

$$T_+ = n^{1/2} \max_{\boldsymbol{\lambda} \geq \mathbf{0}} \frac{\hat{\boldsymbol{\mu}}'\boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}'\hat{\boldsymbol{\Sigma}}\boldsymbol{\lambda}}} = n^{1/2} \frac{\hat{\boldsymbol{\mu}}'\hat{\boldsymbol{\lambda}}}{\sqrt{\hat{\boldsymbol{\lambda}}'\hat{\boldsymbol{\Sigma}}\hat{\boldsymbol{\lambda}}}}, \quad \text{where } \hat{\boldsymbol{\lambda}} = \arg \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \|\boldsymbol{\iota}_n - \mathbf{X}\boldsymbol{\lambda}\|_2^2.$$

This problem can be solved by the non-negative least squares algorithm of Lawson and Hanson (1995).⁵ An interesting additional feature is that the resulting maximizing weight vector $\hat{\boldsymbol{\lambda}}$ may contain zeros, resulting in a selection of moment inequalities that are ‘suspected’ to violate the null.

3 Symmetry based inference

The randomization tests that we consider are based on the symmetry assumption. For completeness, we include this section explain how the assumption leads to exact randomization tests, and show how they can be applied in practice. For other descriptions of randomization tests, see e.g. Lehmann and Romano (2006), Maritz (1995) and Bekker and Lawford (2008).

Let $\mathcal{R} = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N\}$ be the set of $N = 2^n$ diagonal $n \times n$ matrices with diagonal elements in $\{-1, 1\}$, where \mathbf{R}_1 denotes the identity matrix. This set constitutes a finite reflection group under matrix multiplication. The group \mathcal{R} determines a partitioning \mathcal{E} of $\mathbb{R}^{n \times p}$ into equivalence classes (“orbits”) denoted by $\mathcal{R}_{\mathbf{E}^*} = \{\mathbf{E}^*, \mathbf{R}_2\mathbf{E}^*, \dots, \mathbf{R}_N\mathbf{E}^*\}$,

⁵This algorithm is implemented in the R package ‘nnls’ and the MATLAB function ‘lsqnonneg’.

where $\mathbf{E}^* \in \mathbb{R}^{n \times p}$ acts as a representative of the class. For simplicity, we only consider orbits $\mathcal{R}_{\mathbf{E}^*} \in \mathcal{E}$, for which the cardinality satisfies $|\mathcal{R}_{\mathbf{E}^*}| = |\mathcal{R}| = N$. This is equivalent to assuming that we only consider the subset of orbits $\mathcal{E}^0 \subset \mathcal{E}$ where the elements of $\mathcal{R}_{\mathbf{E}^*} \in \mathcal{E}^0$ have no rows that equal zero.

The basic assumption that permits the construction of randomization tests is a distributional invariance assumption on the error term under group transformations. In particular, we assume that $\mathbf{R}\mathbf{E}$ and \mathbf{E} have the same distribution, for all $\mathbf{R} \in \mathcal{R}$. Equivalently, we assume the conditional distribution of \mathbf{E} , given an orbit $\mathbf{E} \in \mathcal{R}_{\mathbf{E}^*}$, is uniform for all orbits $\mathcal{R}_{\mathbf{E}^*} \in \mathcal{E}^0$.

Let $g : \mathcal{R}_{\mathbf{E}^*} \rightarrow \mathbb{R} \cup \{\infty\}$ be a mapping. Notice that if g is injective on all orbits $\mathcal{R}_{\mathbf{E}^*} \in \mathcal{E}^0$, then the symmetry assumption implies that the conditional distribution of $g(\mathbf{E})$, given an orbit $\mathcal{R}_{\mathbf{E}^*} \in \mathcal{E}^0$, is uniform just as well. Furthermore, the function $p(\mathbf{E}) = |\{\mathbf{R} \in \mathcal{R} \mid g(\mathbf{R}\mathbf{E}) \geq g(\mathbf{E})\}|/N$ has a conditional distribution, given an orbit, that is uniform over $\{1/N, 2/N, \dots, 1\}$. As this holds for all $\mathcal{R}_{\mathbf{E}^*} \in \mathcal{E}^0$, the function $p(\mathbf{E})$ has an unconditional uniform distribution over $\{1/N, 2/N, \dots, 1\}$. This leads to the following results.

Proposition 1. $P(p(\mathbf{E}) \leq \alpha) \leq \alpha$, for $\alpha \in [0, 1]$. If g is injective on all $\mathcal{R}_{\mathbf{E}^*} \in \mathcal{E}^0$, then $P(p(\mathbf{E}) \leq \alpha) = \alpha$, for $\alpha \in \{1/N, 2/N, \dots, 1\}$.

Following Bekker and Lawford (2008), we will refer to g as an *inferential function*. As the cardinality of \mathcal{R} grows exponentially in n , the computation of $p(\mathbf{E})$ is intractable even in small samples. Therefore, we provide the following result to allow for sampling from \mathcal{R} .⁶

Proposition 2. Let \mathcal{R}^M contain the identity matrix and $M - 1$ other elements drawn randomly without replacement from $\mathcal{R} \setminus \{\mathbf{R}_1\}$. Let $p_M(\mathbf{E}) = |\{\mathbf{R} \in \mathcal{R}^M \mid g(\mathbf{R}\mathbf{E}) \geq g(\mathbf{E})\}|/M$, then $P(p_M(\mathbf{E}) \leq \alpha) \leq \alpha$, for $\alpha \in [0, 1]$ and, if g is injective, then $P(p_M(\mathbf{E}) \leq \alpha) = \alpha$ for $\alpha \in \{1/M, 2/M, \dots, 1\}$.

⁶Bekker and Lawford (2008) also describe sampling with replacement.

4 Randomization tests for moment inequalities

To test the moment inequalities of model (1), we use randomization tests based on the $T_{\mathcal{U}}$ statistic with $\mathcal{U} \subseteq \mathbb{R}_+^p$, such as t_{\max} and T_+ . We will use the convention that $T_{\mathcal{U}} = 0$ if $\hat{\boldsymbol{\mu}} \leq \mathbf{0}$ and $T_{\mathcal{U}} = \infty$ if Condition 1 does not hold. Let the test statistic be generically denoted as T . The tests are exact if the moment inequalities are binding. In that case $\mathbf{X} = \mathbf{E}$ and T may be seen as an inferential function $g(\mathbf{E}) = T(\mathbf{E})$. So, Proposition 2 can be used based on $p_M(\mathbf{E}) = p_M(\mathbf{X}) = |\{\mathbf{R} \in \mathcal{R}^M \mid T(\mathbf{R}\mathbf{X}) \geq T(\mathbf{X})\}|/M$. In particular, we use the following testing procedure.

Algorithm 1 (Symmetry Randomization Test).

1. Create the set \mathcal{R}^M consisting of the \mathbf{I}_n and $M - 1$ matrices drawn from $\mathcal{R} \setminus \{\mathbf{I}_n\}$ without replacement.
2. Generate the transformed dataset $\mathbf{R}\mathbf{X}$ and compute $T(\mathbf{R}\mathbf{X})$, for all $\mathbf{R} \in \mathcal{R}_M$.
3. Compute $p_M(\mathbf{X})$ by counting the proportion that exceeds $T(\mathbf{X})$.
4. Reject the null if $p_M(\mathbf{X}) \leq \alpha$.

If $\boldsymbol{\mu} = \mathbf{0}$ and T is injective and $\alpha \in \{1/N, 2/N, \dots, 1\}$, this test is exact (note that point accumulation could occur at 0 or ∞).

In case some moment inequalities are assumed to be strictly satisfied, the test is no longer exact, but we observe in all our simulations that the size is controlled if the symmetry assumption holds. To prove this formally is another matter. If, under the null hypothesis $\boldsymbol{\mu} \leq \mathbf{0}$, it can be verified that $T(\mathbf{R}\mathbf{X}) \geq T(\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})$, then the test is conservative.⁷ In that case we find, that $T(\mathbf{X}) = T(\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E})$ is an inferential function, $g(\mathbf{E}) = T(\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E})$,

⁷Notice that for the numerator of $T_+(\mathbf{R}\mathbf{X})$ it holds that $\boldsymbol{\nu}'_n \mathbf{R}\mathbf{X}\boldsymbol{\lambda} = \boldsymbol{\nu}'_n \mathbf{R}\boldsymbol{\nu}_n \boldsymbol{\mu}' \boldsymbol{\lambda} + \boldsymbol{\nu}'_n \mathbf{R}\mathbf{E}\boldsymbol{\lambda} \geq \boldsymbol{\nu}'_n (\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})\boldsymbol{\lambda}$.

so that

$$p^T(\mathbf{X}) = |\{\mathbf{R} \in \mathcal{R} \mid T(\mathbf{R}\mathbf{X}) \geq T(\mathbf{X})\}| \geq |\{\mathbf{R} \in \mathcal{R} \mid g(\mathbf{R}\mathbf{E}) \geq g(\mathbf{E})\}| = p(\mathbf{E}),$$

which implies $P(p^T(\mathbf{X}) \leq \alpha) \leq P(p(\mathbf{E}) \leq \alpha) \leq \alpha$. Unfortunately, we have not been able to formally prove whether or not the condition holds for our test statistics. However, a condition that the test is conservative for sufficiently large samples is easy.

Condition 3. As $n \rightarrow \infty$ and p is fixed, $\mathbf{E}'\boldsymbol{\nu}_n/n \xrightarrow{a.s.} \mathbf{0}$ and $\mathbf{E}'\mathbf{E}/n \xrightarrow{a.s.} \boldsymbol{\Sigma}$.

To verify $T(\mathbf{R}\mathbf{X}) \geq T(\boldsymbol{\nu}_n\boldsymbol{\mu}' + \mathbf{R}\mathbf{E})$, first notice that this inequality holds trivially if $\mathbf{R} = \mathbf{I}_n$ or $\boldsymbol{\mu}'\boldsymbol{\lambda} = 0$. So, we can restrict ourselves to the case that $\mathbf{R} \neq \mathbf{I}_n$ and $\boldsymbol{\mu}'\boldsymbol{\lambda} < 0$, as $\boldsymbol{\mu}'\boldsymbol{\lambda} \leq 0$ for $\boldsymbol{\lambda} \geq \mathbf{0}$. Furthermore, notice that

$$T^*(\mathbf{X}) = \max_{\boldsymbol{\lambda} \in \mathcal{U}} \frac{\boldsymbol{\nu}'_n \mathbf{X} \boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}' \mathbf{X}' \mathbf{X} \boldsymbol{\lambda}}} = \frac{T(\mathbf{X})}{\sqrt{1 + T(\mathbf{X})^2/n}},$$

is a strict monotone transformation of T , so that we can instead look at T^* . Consider the function of $\boldsymbol{\lambda} \in \mathcal{U} \subseteq \mathbb{R}_+^p$ that is maximized at $T^*(\mathbf{R}\mathbf{X})$,

$$\begin{aligned} \frac{\boldsymbol{\nu}'_n \mathbf{R} \mathbf{X} \boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}' \mathbf{X}' \mathbf{X} \boldsymbol{\lambda}}} &= \frac{\boldsymbol{\nu}'_n \mathbf{R} \boldsymbol{\nu}_n \boldsymbol{\mu}' \boldsymbol{\lambda} + \boldsymbol{\nu}'_n \mathbf{R} \mathbf{E} \boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}' (\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E})' (\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E}) \boldsymbol{\lambda}}} > \frac{n \boldsymbol{\mu}' \boldsymbol{\lambda} + \boldsymbol{\nu}'_n \mathbf{R} \mathbf{E} \boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}' (\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E})' (\boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E}) \boldsymbol{\lambda}}} \\ &= \frac{n \boldsymbol{\mu}' \boldsymbol{\lambda} + \boldsymbol{\nu}'_n \mathbf{R} \mathbf{E} \boldsymbol{\lambda}}{\sqrt{n(\boldsymbol{\mu}' \boldsymbol{\lambda})^2 + 2 \boldsymbol{\lambda}' \mathbf{E}' \boldsymbol{\nu}_n \boldsymbol{\mu}' \boldsymbol{\lambda} + \boldsymbol{\lambda}' \mathbf{E}' \mathbf{E} \boldsymbol{\lambda}}}, \end{aligned}$$

as $\boldsymbol{\nu}'_n \mathbf{R} \boldsymbol{\nu}_n < n$. Consequently, due to Condition 3, as \mathbf{E} and $\mathbf{R}\mathbf{E}$ follow the same distribution,

$$n^{-1/2} \frac{\boldsymbol{\nu}'_n \mathbf{R} \mathbf{X} \boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}' \mathbf{X}' \mathbf{X} \boldsymbol{\lambda}}} > \frac{\boldsymbol{\mu}' \boldsymbol{\lambda}}{(\boldsymbol{\mu}' \boldsymbol{\lambda})^2 + \boldsymbol{\lambda}' \boldsymbol{\Sigma} \boldsymbol{\lambda}} + o_{a.s.}(0).$$

Furthermore, $T^*(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})$ maximizes

$$\frac{\boldsymbol{\iota}'_n(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})\boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}'(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})'(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})\boldsymbol{\lambda}}} = \frac{n\boldsymbol{\mu}'\boldsymbol{\lambda} + \boldsymbol{\iota}'_n \mathbf{R}\mathbf{E}\boldsymbol{\lambda}}{\sqrt{n(\boldsymbol{\mu}'\boldsymbol{\lambda})^2 + 2\boldsymbol{\lambda}'\mathbf{E}'\mathbf{R}\boldsymbol{\iota}_n \boldsymbol{\mu}'\boldsymbol{\lambda} + \boldsymbol{\lambda}'\mathbf{E}'\mathbf{E}\boldsymbol{\lambda}}},$$

so that

$$n^{-1/2} \frac{\boldsymbol{\iota}'_n(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})\boldsymbol{\lambda}}{\sqrt{\boldsymbol{\lambda}'(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})'(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})\boldsymbol{\lambda}}} = \frac{\boldsymbol{\mu}'\boldsymbol{\lambda}}{(\boldsymbol{\mu}'\boldsymbol{\lambda})^2 + \boldsymbol{\lambda}'\boldsymbol{\Sigma}\boldsymbol{\lambda}} + o_{a.s.}(0).$$

Let $\boldsymbol{\lambda}^*$ be the maximizer of $T^*(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})$, then

$$n^{1/2}T^*(\mathbf{R}\mathbf{X}) \geq n^{-1/2} \frac{\boldsymbol{\iota}'_n \mathbf{R}\mathbf{X}\boldsymbol{\lambda}^*}{\sqrt{\boldsymbol{\lambda}^{*\prime} \mathbf{X}'\mathbf{X}\boldsymbol{\lambda}^*}} > n^{-1/2}T^*(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E}) + o_{a.s.}(0).$$

As a result, the condition $T^*(\mathbf{R}\mathbf{X}) \geq T^*(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})$ holds, and therefore $T(\mathbf{R}\mathbf{X}) \geq T(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{R}\mathbf{E})$, so the test is conservative if n is sufficiently large.

5 Symmetry based inference with pre-selection

As the tests are found to be conservative in case some moment equalities are strictly satisfied, there may be room for power improvements. CCK propose an inequality selection step aiming at removal of such strict moment inequalities j . In particular, they remove moment inequalities j for which $t_j < c$, where c is some chosen cut-off value. They propose several techniques to select this cut-off value such that the asymptotic testing procedures remain applicable, up to a small correction of the significance level. A similar way to select the cut-off value using Lasso is proposed by Bugni et al. (2016).

Our approach for inference with pre-selection is similar to the approach of the previous section where there was no pre-selection. We use symmetry based inference. Given any inequality selection rule we describe tests that are exact or conservative if the moment inequalities are binding. That is to say, if $\boldsymbol{\mu} = \mathbf{0}$, then $T(\mathbf{X}) = T(\mathbf{E})$ and $p^T(\mathbf{X}) = |\{\mathbf{R} \in \mathcal{R} \mid T(\mathbf{R}\mathbf{X}) \geq T(\mathbf{X})\}| \leq \alpha$.

Let $\mathcal{J}(\mathbf{X}) \subset \{1, \dots, p\}$ be an index selection subset, and let $\mathbf{X}_{\mathcal{J}(\mathbf{X})}$ be the submatrix of \mathbf{X} consisting of columns with indexes in $\mathcal{J}(\mathbf{X})$. As the selection depends on \mathbf{X} , it is not guaranteed that $p^T(\mathbf{X}_{\mathcal{J}(\mathbf{X})}) = |\{\mathbf{R} \in \mathcal{R} \mid T(\mathbf{R}\mathbf{X}_{\mathcal{J}(\mathbf{X})}) \geq T(\mathbf{X}_{\mathcal{J}(\mathbf{X})})\}| \leq \alpha$. However, if $\boldsymbol{\mu} = \mathbf{0}$, then $T(\mathbf{X}_{\mathcal{J}(\mathbf{X})}) = T(\mathbf{E}_{\mathcal{J}(\mathbf{E})}) = g(\mathbf{E})$ is an inferential function and

$$\begin{aligned} p_{\text{sel}}^T(\mathbf{X}_{\mathcal{J}(\mathbf{X})}) &= |\{\mathbf{R} \in \mathcal{R} \mid T(\mathbf{R}\mathbf{X}_{\mathcal{J}(\mathbf{R}\mathbf{X})}) \geq T(\mathbf{X}_{\mathcal{J}(\mathbf{X})})\}| \\ &= |\{\mathbf{R} \in \mathcal{R} \mid g(\mathbf{R}\mathbf{E}) \geq g(\mathbf{E})\}| = p(\mathbf{E}), \end{aligned}$$

which implies $P(p_{\text{sel}}^T(\mathbf{X}_{\mathcal{J}(\mathbf{X})}) \leq \alpha) = P(p(\mathbf{E}) \leq \alpha) \leq \alpha$, for $\alpha \in [0, 1]$ and, if g is injective, then $P(p_M(\mathbf{E}) \leq \alpha) = \alpha$ for $\alpha \in \{1/M, 2/M, \dots, 1\}$, as in Proposition 2. Therefore, we use the following testing procedure.

Algorithm 2 (Symmetry randomization test with inequality selection).

1. Create the set \mathcal{R}^M consisting of the \mathbf{I}_n and $M - 1$ matrices drawn from $\mathcal{R} \setminus \{\mathbf{I}_n\}$ without replacement.
2. Compute the selected inequalities $\mathcal{J}(\mathbf{R}\mathbf{X})$, for each $\mathbf{R} \in \mathcal{R}_M$.
3. Compute $T(\mathbf{R}\mathbf{X}_{\mathcal{J}(\mathbf{R}\mathbf{X})})$ for all $\mathbf{R} \in \mathcal{R}_M$.
4. Compute $p_{\text{sel}}^T(\mathbf{X})$, by counting the proportion that exceeds $T(\mathbf{X}_{\mathcal{J}(\mathbf{X})})$.
5. Reject the null if $p_{\text{sel}}^T(\mathbf{X}) \leq \alpha$.

If $\boldsymbol{\mu} \leq \mathbf{0}$ and $\boldsymbol{\mu} \neq \mathbf{0}$, the test is conservative if $T(\mathbf{R}\mathbf{X}_{\mathcal{J}(\mathbf{R}\mathbf{X})}) \geq T((\boldsymbol{\nu}\boldsymbol{\mu}' + \mathbf{R}\mathbf{E})_{\mathcal{J}(\boldsymbol{\nu}\boldsymbol{\mu}' + \mathbf{R}\mathbf{E})})$. The power of the test varies with the value of $\boldsymbol{\mu}$ and the selection method. We follow CCK and use $\mathcal{J}(\mathbf{X}) = \{j \in \{1, \dots, p\} \mid t_j > c\}$, where the selection constant is chosen using their empirical bootstrap procedure. The simulations show that the size of the tests is controlled when $\boldsymbol{\mu} \leq \mathbf{0}$.

6 Simulations

In this section, we present Monte Carlo simulation results. We use setups based on the simulation experiments presented by CCK and Bugni et al. (2016). The tests described in this paper are compared to the Empirical Bootstrap (EB) tests described in CCK.⁸ All experiments were implemented in R and code for the tests will be made available at <https://github.com/nickwkoning/>.

Data generation

The data is created as follows: we generate $n \times p$ matrix $\mathbf{X} = \boldsymbol{\nu}_n \boldsymbol{\mu}' + \mathbf{E}\mathbf{A}$, where $\boldsymbol{\nu}_n$ is an n -vector of ones, $n \times p$ matrix \mathbf{E} has i.i.d. elements drawn from a distribution F , and \mathbf{A} is defined such that $\boldsymbol{\Sigma} = \mathbf{A}'\mathbf{A}$, where $\boldsymbol{\Sigma}$ has elements $\sigma_{ij} = \rho^{|i-j|}$.

The parameters we use are $n \in \{30, 400\}$, $p \in \{200, 500, 1000\}$ and $\rho \in \{0, 0.5, 0.9\}$. For the vector $\boldsymbol{\mu}$, we consider four different designs. The signs of the elements of $\boldsymbol{\mu}$ are chosen to represent four general cases: all moment inequalities are binding (Design 1), most moment inequalities are strictly satisfied and some are binding (Design 2), all moment inequalities are violated (Design 3), and some moment inequalities are violated while some are strictly satisfied (Design 4). The values of $\boldsymbol{\mu}$ for each combination of n and p were selected to ensure that the rejection probabilities are bounded away from 1.

For the errors, we use a symmetric and two asymmetric distributions: one left-skewed and one right-skewed. As symmetric distribution, we use $t(4)/\sqrt{2}$, where $t(4)$ is the Student's t distribution with 4 degrees of freedom and we divide by $\sqrt{2}$ so that the variance is 1. As asymmetric distributions we use the skew-normal distribution with mean 0, variance equal to 1 and two configurations for the skewness: $\gamma = -0.667$ and $\gamma = 0.667$. A density plot comparing these skewed distributions to the standard normal distribution is provided in Figure 1. For the sake of brevity, the asymmetric error distributions were only

⁸We only compare to the EB tests described in CCK, as they find that the EB tests perform similarly to the Multiplier Bootstrap tests and better than the self-normalized tests.

considered for Design 1.

This setup is similar to the setup used by an earlier working paper of CCK.⁹ The differences are that we do not consider equicorrelated data or uniformly distributed errors, but instead consider small samples ($n = 30$) and asymmetric error distributions. In addition, the positive values of $\boldsymbol{\mu}$ are substantially decreased to account for the higher power of tests based on the T_+ statistic.

Tests

We use symmetry based randomization (SR) tests based on the test statistics t_{\max} and T_+ , with and without pre-selection. In addition, we also consider the $t_{\max}^{\iota} = T_{\mathcal{C}_p \cup \{\iota_p\}}$ statistic, where the maximum is taken over just a single extra vector compared to t_{\max} . As a comparison, we include the EB tests described by CCK with and without pre-selection for both the t_{\max} and t_{\max}^{ι} statistics.

For each setting, 1000 realizations of \mathbf{X} were generated and the proportion of rejections was recorded. For the EB tests 1000 bootstrap samples were used and for the SR tests we used 1000 reflection samples. The significance level was fixed at $\alpha = 0.05$. The selection constant for the tests that include pre-selection is chosen using the EB procedure with $\beta = 0.001$ and 1000 resamples as described in CCK. For the small sample experiments with $n = 30$, we do not use pre-selection as the selection constant depends on the asymptotic properties of the EB test.

In terms of computation time for the largest settings, a single test without pre-selection using the T_+ statistic on a single core of a standard 13-inch 2017 MacBook Pro takes approximately 80-270 seconds, depending on the design. For the t_{\max} and t_{\max}^{ι} statistics this is approximately 5-7 seconds.

⁹This earlier version can be found at <https://arxiv.org/abs/1312.7614v4>.

6.1 Results

The results of the simulation experiments can be found in Tables 1 to 5, and are discussed separately for each design.

Design 1: $\boldsymbol{\mu} = \mathbf{0}$.

In Design 1, all elements of $\boldsymbol{\mu}$ are equal to zero. Therefore, the rejection probability for the tests should be at most $\alpha = 0.05$. The results for Design 1 are displayed in Table 1 for the symmetric error distribution, and Table 5 for the asymmetric error distributions.

In Table 1, for the case that the number of observations is large ($n = 400$), we see that the EB tests for both t_{\max} and t_{\max}^t reject with probability close to 0.05, which is in line with Remark 1. For the symmetry randomization (SR) tests, Proposition 2 states that all rejection probabilities should be at most 0.05, which is confirmed in the simulations. In addition, if the test statistic is injective on all orbits, then the rejection probability is equal to α . One exception found to this in the simulations is the configuration $p = 1000$ and $\rho = 0$, for the statistic T_+ , where the proportion of rejections is 0. Further inspection shows that for this configuration without pre-selection, the test statistics $T_+(\mathbf{X})$ and critical values, defined as the $1 - \alpha$ quantile of the set $\{T_+(\mathbf{R}\mathbf{X}) \mid \mathbf{R} \in \mathcal{R}_M\}$, are infinite, as Condition 2 is not satisfied. Therefore, the function T_+ is not injective on the orbit and the rejection probability can be strictly smaller than α , according Proposition 2. A similar problem of non-injectivity occurs for the case with pre-selection. In line with the relation between Condition 2 and copositivity, the issue diminishes if ρ is increased.

If the number of observations is small ($n = 30$), then for the t_{\max} and t_{\max}^t statistics, the rejection rates for the SR tests remain approximately α . In contrast, the EB tests over-reject. For the SR tests based on the T_+ statistic, the rejection rate is frequently zero due to the non-injectivity phenomenon described above.

Table 5 shows the rejection rate for Design 1 when the error distributions are left-skewed ($\gamma = -0.667$) and right-skewed ($\gamma = 0.667$). As the symmetry assumption is violated, it

is not surprising that the SR tests have a rejection rate different from α . In particular, we find that the rejection rate for the left-skewed error distribution is larger than α , while the rejection rate for the right-skewed error distribution is smaller than α .

Surprisingly, even though the sample size is large ($n = 400$), the EB tests over-reject even more than the SR tests under the left-skewed distribution. This suggests that the EB tests may also benefit from errors being symmetrically distributed in finite samples. Although these simulation results are by no means exhaustive, they suggest that one may sometimes be better off using an SR test than an EB test even if it is known that the true error distribution is not symmetric.

Design 2: $\mu \leq 0, \mu \neq 0$.

In Design 2, for $n = 400$, the first $0.1p$ elements of μ are equal to 0 and the remaining elements are equal to -0.8 . For $n = 30$, the first 10 elements are 0 and the remaining elements are equal to -5 . As the data for Design 2 is generated under H_0 , the rejection rates should be smaller than α (up to sampling errors) if no pre-selection is used, and close to α if pre-selection is used. The results for Design 2 can be found in Table 2.

From Table 2, it can be observed that the rejection rates of the tests without pre-selection is smaller than α . For the SR tests without pre-selection and $n = 400$, this is in line with the conclusions in Section 4 as n is large. For $n = 30$, the rejection rates are closer to, but still smaller than α . In addition, pre-selection lifts the rejection rates close to the nominal size α .

Design 3: $\mu > 0$.

In Design 3, all elements of μ are positive. In particular, for $n = 400$ they are equal to 0.01 and for $n = 30$ they are equal to 0.03. As the data is generated under the alternative, the rejection probability should be as large as possible. The results for Design 3 are found in Table 3.

In Table 3, it can be seen that for $n = 400$ the power of the EB and SR tests is similar, and inequality selection has no noticeable effect. The T_+ statistic generally leads to the highest power. One exception is the case where $p = 1000$ and $\rho = 0$, which is the most challenging case with respect to Condition 2. The power difference between the t_{\max}^t and t_{\max} statistics is quite remarkable, as the t_{\max}^t statistic maximizes over just a single extra vector. However, in Design 4 we will see that the statistic is more sensitive to the alternative than T_+ . Furthermore, the difference in the power between the statistics is largest under weak correlations, which is in line with Remark 2.

For $n = 30$, the results for the EB tests should be ignored, as the tests do not control size. Although there are multiple violations, the T_+ statistic performs worst. This is expected as Condition 2 is unlikely to be satisfied here. So the same phenomenon observed under Design 1 occurs, where both the test statistic and critical value are infinite. Therefore, it is recommended that the test statistic and critical value are inspected before the outcome of the test is interpreted. If both are infinite, then the test based on the T_+ statistic should not be used. In addition, the t_{\max}^t statistic generally performs as well as or better than the t_{\max} statistic, especially if $\rho = 0$.

Design 4: $\boldsymbol{\mu} \not\leq \mathbf{0}$, $\boldsymbol{\mu} \not\geq \mathbf{0}$.

In Design 4, for $n = 400$, the first $0.1p$ elements of $\boldsymbol{\mu}$ are equal to 0.02 and the remaining elements are equal to -0.75 . For $n = 30$, the first 10 elements of $\boldsymbol{\mu}$ are equal to 0.3 and the remaining elements are equal to -5 . The data is generated under the alternative, so the rejection rates should be as large as possible. The results for Design 4 are presented in Table 4.

The results in Table 4 for $n = 400$ show that the tests with pre-selection have substantially higher power than the tests without pre-selection. When considering the tests with pre-selection, tests based on the T_+ statistic have much more power than the tests based on the t_{\max} and t_{\max}^t statistics. The power for the EB and SR tests based on the t_{\max} statistic

are similar if pre-selection is used, but the SR tests seem more powerful if pre-selection is not used. In addition, the performance of the t_{\max}^u and t_{\max} statistics is similar, unlike in Design 3. Further inspection shows that the pre-selection often fails to eliminate a few strictly satisfied moment inequalities, which causes the added weight vector \boldsymbol{t}_p to rarely be the maximizing vector.

Even though the EB tests do not control size for $n = 30$, the SR tests seem to be more powerful. Comparing the different statistics for the SR tests, it can be observed that the T_+ statistic performs well compared to the t_{\max} statistic if $\rho = 0$, even if $p = 1000$ and $n = 30$. So, the issue regarding Condition 1 is not observed here. This is in line with the intuitions obtained in Section 2 which suggest that the condition may still be satisfied if many of the moment inequalities are strictly satisfied. The t_{\max}^u statistic performs best overall for $n = 30$.

References

- R. Allen. Testing moment inequalities: Selection versus recentering. *Economics Letters*, 162:124–126, 2018.
- D. W. Andrews and P. J. Barwick. Inference for parameters defined by moment inequalities: A recommended moment selection procedure. *Econometrica*, 80(6):2805–2826, 2012.
- D. W. Andrews and P. Guggenberger. Validity of subsampling and plug-in asymptotic inference for parameters defined by moment inequalities. *Econometric Theory*, 25(3): 669–709, 2009.
- D. W. Andrews and X. Shi. Inference based on conditional moment inequalities. *Econometrica*, 81(2):609–666, 2013.
- T. B. Armstrong. Weighted ks statistics for inference on conditional moment inequalities. *Journal of Econometrics*, 181(2):92–116, 2014.
- T. B. Armstrong. Asymptotically exact inference in conditional moment inequality models. *Journal of Econometrics*, 186(1):51–65, 2015.
- T. B. Armstrong and H. P. Chan. Multiscale adaptive inference on conditional moment inequalities. *Journal of Econometrics*, 194(1):24–43, 2016.
- P. A. Bekker. Exact inference for the linear model with groupwise heteroscedastic spherical disturbances. *Journal of Econometrics*, 111(2):285 – 302, 2002. ISSN 0304-4076.

- P. A. Bekker and S. Lawford. Symmetry-based inference in an instrumental variable setting. *Journal of Econometrics*, 142(1):28–49, 2008.
- F. A. Bugni, M. Caner, A. B. Kock, and S. Lahiri. Inference in partially identified models with many moment inequalities using lasso. *arXiv preprint arXiv:1604.02309*, 2016.
- I. A. Canay. El inference for partially identified models: Large deviations optimality and bootstrap validity. *Journal of Econometrics*, 156(2):408–425, 2010.
- V. Chernozhukov, H. Hong, and E. Tamer. Estimation and confidence regions for parameter sets in econometric models 1. *Econometrica*, 75(5):1243–1284, 2007.
- V. Chernozhukov, S. Lee, and A. M. Rosen. Intersection bounds: Estimation and inference. *Econometrica*, 81(2):667–737, 2013.
- V. Chernozhukov, D. Chetverikov, and K. Kato. Inference on causal and structural parameters using many moment inequalities. *The Review of Economic Studies*, 86(5):1867–1900, 2019.
- R. A. Fisher. *The Design of Experiments*. Oliver and Boyd, London, 1935.
- J.-B. Hiriart-Urruty and A. Seeger. A variational approach to copositive matrices. *SIAM review*, 52(4):593–629, 2010.
- N. Koning. Directing power towards conic parameter subspaces. *arXiv preprint arXiv:1907.05077*, 2019.
- C. L. Lawson and R. J. Hanson. *Solving least squares problems*, volume 15. Siam, 1995.
- S. Lee, K. Song, and Y.-J. Whang. Testing functional inequalities. *Journal of Econometrics*, 172(1):14–32, 2013.
- S. Lee, K. Song, and Y.-J. Whang. Testing for a general class of functional inequalities. *Econometric Theory*, 34(5):1018–1064, 2018.
- E. L. Lehmann and J. P. Romano. *Testing statistical hypotheses*. Springer Science & Business Media, 2006.
- J. S. Maritz. *Distribution-free statistical methods*, volume 17. CRC Press, 1995.
- N. Meinshausen et al. Sign-constrained least squares estimation for high-dimensional regression. *Electronic Journal of Statistics*, 7:1607–1631, 2013.
- J. P. Romano. On the behavior of randomization tests without a group invariance assumption. *Journal of the American Statistical Association*, 85(411):686–692, 1990.
- J. P. Romano and A. M. Shaikh. Inference for identifiable parameters in partially identified econometric models. *Journal of Statistical Planning and Inference*, 138(9):2786–2807, 2008.

- A. M. Rosen. Confidence sets for partially identified parameters that satisfy a finite number of moment inequalities. *Journal of Econometrics*, 146(1):107–117, 2008.
- M. Slawski, M. Hein, et al. Non-negative least squares for high-dimensional linear models: Consistency and sparse recovery without regularization. *Electronic Journal of Statistics*, 7:3004–3056, 2013.




Figure 1: A density plot of the left-skewed (left) and right-skewed (right) normal distribution with mean 0, standard deviation 1 and skewness -0.667 and 0.667 , respectively, overlaid on the density of the standard normal distribution.

| n | p | ρ | t_{\max} | | | | t_{\max}^t | | | | T_+ | |
|------|-----|--------|------------|------|-----------|-----------|--------------|------|-----------|-----------|-------|-----------|
| | | | EB | SR | EB sel | SR sel | EB | SR | EB sel | SR sel | SR | SR sel |
| 200 | 500 | 0 | .045 | .046 | .032 | .052 | .042 | .050 | .041 | .050 | .040 | .043 |
| | | 0.5 | .044 | .038 | .046 | .057 | .050 | .048 | .042 | .061 | .055 | .046 |
| | | 0.9 | .047 | .049 | .052 | .040 | .052 | .051 | .075 | .060 | .068 | .050 |
| 400 | 500 | 0 | .044 | .054 | .043 | .032 | .057 | .042 | .041 | .051 | .054 | .061 |
| | | 0.5 | .043 | .054 | .034 | .069 | .051 | .049 | .038 | .049 | .054 | .045 |
| | | 0.9 | .059 | .051 | .053 | .046 | .060 | .050 | .057 | .059 | .048 | .049 |
| 1000 | 500 | 0 | .043 | .042 | .051 | .063 | .036 | .060 | .036 | .053 | .000 | .000 |
| | | 0.5 | .044 | .063 | .051 | .052 | .052 | .057 | .038 | .052 | .049 | .054 |
| | | 0.9 | .058 | .057 | .042 | .052 | .046 | .040 | .047 | .041 | .039 | .051 |
| 30 | 500 | 0 | .094 | .055 | - | - | .111 | .061 | - | - | .000 | - |
| | | 0.5 | .122 | .056 | - | - | .117 | .048 | - | - | .000 | - |
| | | 0.9 | .128 | .050 | - | - | .153 | .045 | - | - | .048 | - |
| 1000 | 500 | 0 | .114 | .040 | - | - | .124 | .045 | - | - | .000 | - |
| | | 0.5 | .142 | .044 | - | - | .141 | .049 | - | - | .000 | - |
| | | 0.9 | .167 | .052 | - | - | .173 | .058 | - | - | .000 | - |
| 1000 | 500 | 0 | .137 | .049 | - | - | .155 | .045 | - | - | .000 | - |
| | | 0.5 | .170 | .063 | - | - | .174 | .051 | - | - | .000 | - |
| | | 0.9 | .210 | .049 | - | - | .204 | .060 | - | - | .000 | - |

Table 1: Monte Carlo rejection probabilities with 1000 repetitions for Design 1: $\mu = \mathbf{0}$, with symmetrically distributed errors $(t(4)/\sqrt{2})$. The columns represent the Empirical Bootstrap (EB) and Symmetry Randomization (SR) tests, based on the t_{\max} , t_{\max}^t and T_+ statistics, both with pre-selection (sel) and without pre-selection.

| n | p | ρ | t_{\max} | | | | t_{\max}^t | | | | T_+ | |
|------|-----|--------|------------|------|-----------|-----------|--------------|------|-----------|-----------|-------|-----------|
| | | | EB | SR | EB sel | SR sel | EB | SR | EB sel | SR sel | SR | SR sel |
| 200 | | 0 | .003 | .009 | .038 | .037 | .002 | .009 | .033 | .053 | .000 | .049 |
| | | .5 | .005 | .006 | .031 | .045 | .006 | .004 | .045 | .048 | .001 | .055 |
| | | .9 | .006 | .009 | .040 | .054 | .004 | .011 | .041 | .038 | .000 | .060 |
| 400 | 500 | 0 | .000 | .011 | .035 | .057 | .003 | .003 | .043 | .032 | .002 | .042 |
| | | .5 | .008 | .005 | .059 | .063 | .004 | .008 | .046 | .037 | .001 | .047 |
| | | .9 | .005 | .005 | .039 | .049 | .005 | .009 | .052 | .050 | .000 | .048 |
| 1000 | | 0 | .002 | .009 | .025 | .050 | .004 | .003 | .032 | .038 | .000 | .057 |
| | | .5 | .004 | .012 | .040 | .053 | .003 | .007 | .031 | .050 | .000 | .055 |
| | | .9 | .007 | .012 | .047 | .049 | .003 | .010 | .041 | .030 | .001 | .050 |
| 200 | | 0 | .004 | .034 | - | - | .006 | .044 | - | - | .009 | - |
| | | .5 | .006 | .020 | - | - | .013 | .046 | - | - | .012 | - |
| | | .9 | .010 | .023 | - | - | .014 | .033 | - | - | .014 | - |
| 30 | 500 | 0 | .002 | .021 | - | - | .011 | .042 | - | - | .008 | - |
| | | .5 | .005 | .028 | - | - | .019 | .040 | - | - | .010 | - |
| | | .9 | .002 | .037 | - | - | .017 | .045 | - | - | .007 | - |
| 1000 | | 0 | .002 | .035 | - | - | .014 | .037 | - | - | .003 | - |
| | | .5 | .003 | .022 | - | - | .024 | .032 | - | - | .013 | - |
| | | .9 | .004 | .024 | - | - | .029 | .042 | - | - | .006 | - |

Table 2: Monte Carlo rejection probabilities with 1000 repetitions for Design 2: $\boldsymbol{\mu} \leq \mathbf{0}$, with symmetrically distributed errors ($t(4)/\sqrt{2}$). For the cases that $n = 400$, $\mu_j = 0$ if $j \leq 0.1p$ and $\mu_j = -0.8$ if $j > 0.1p$, and for the cases where $n = 30$, $\mu_j = 0$ if $j \leq 10$ and $\mu_j = -5$ if $j > 10$. The columns represent the Empirical Bootstrap (EB) and Symmetry Randomization (SR) tests, based on the t_{\max} , t_{\max}^t and T_+ statistics, both with pre-selection (sel) and without pre-selection.

| n | p | ρ | t_{\max} | | | | t_{\max}^{μ} | | | | T_+ | |
|------|-----|--------|------------|------|-----------|-----------|------------------|------|-----------|-----------|-------|-----------|
| | | | EB | SR | EB sel | SR sel | EB | SR | EB sel | SR sel | SR | SR sel |
| 200 | | 0 | .093 | .099 | .094 | .099 | .307 | .333 | .313 | .340 | .668 | .667 |
| | | .5 | .099 | .087 | .094 | .084 | .135 | .125 | .141 | .117 | .377 | .357 |
| | | .9 | .099 | .119 | .099 | .099 | .086 | .114 | .078 | .083 | .129 | .144 |
| 400 | 500 | 0 | .103 | .108 | .094 | .112 | .755 | .806 | .773 | .799 | .919 | .926 |
| | | .5 | .106 | .107 | .100 | .097 | .206 | .216 | .207 | .237 | .649 | .653 |
| | | .9 | .116 | .083 | .111 | .094 | .092 | .109 | .092 | .108 | .215 | .230 |
| 1000 | | 0 | .095 | .106 | .094 | .100 | .985 | .991 | .991 | .994 | .000 | .000 |
| | | .5 | .099 | .116 | .094 | .095 | .447 | .458 | .461 | .453 | .848 | .864 |
| | | .9 | .103 | .117 | .111 | .099 | .113 | .119 | .129 | .111 | .365 | .342 |
| 200 | | 0 | .189 | .107 | - | - | .297 | .202 | - | - | .000 | - |
| | | .5 | .190 | .101 | - | - | .231 | .094 | - | - | .000 | - |
| | | .9 | .167 | .086 | - | - | .168 | .075 | - | - | .129 | - |
| 30 | 500 | 0 | .225 | .092 | - | - | .610 | .411 | - | - | .000 | - |
| | | .5 | .253 | .100 | - | - | .298 | .121 | - | - | .000 | - |
| | | .9 | .275 | .099 | - | - | .264 | .068 | - | - | .000 | - |
| 1000 | | 0 | .242 | .111 | - | - | .896 | .818 | - | - | .000 | - |
| | | .5 | .302 | .091 | - | - | .429 | .209 | - | - | .000 | - |
| | | .9 | .350 | .086 | - | - | .315 | .115 | - | - | .000 | - |

Table 3: Monte Carlo rejection probabilities with 1000 repetitions for Design 3: $\boldsymbol{\mu} > \mathbf{0}$, with symmetrically distributed errors ($t(4)/\sqrt{2}$). For the cases that $n = 400$, $\mu_j = 0.01$ and for the cases that $n = 30$, $\mu_j = 0.03$, for all j . The columns represent the Empirical Bootstrap (EB) and Symmetry Randomization (SR) tests, based on the t_{\max} , t_{\max}^{μ} and T_+ statistics, both with pre-selection (sel) and without pre-selection.

| n | p | ρ | t_{\max} | | | | t_{\max}^t | | | | T_+ | |
|-----|------|--------|------------|------|-----------|-----------|--------------|------|-----------|-----------|-------|-----------|
| | | | EB | SR | EB sel | SR sel | EB | SR | EB sel | SR sel | SR | SR sel |
| | 200 | 0 | .017 | .033 | .135 | .133 | .018 | .034 | .109 | .124 | .025 | .394 |
| | | .5 | .026 | .032 | .136 | .139 | .010 | .035 | .147 | .129 | .006 | .242 |
| | | .9 | .020 | .027 | .113 | .122 | .011 | .034 | .123 | .111 | .002 | .130 |
| 400 | 500 | 0 | .018 | .031 | .099 | .103 | .024 | .029 | .107 | .136 | .143 | .675 |
| | | .5 | .025 | .038 | .141 | .136 | .022 | .034 | .138 | .156 | .015 | .367 |
| | | .9 | .026 | .037 | .117 | .146 | .015 | .022 | .137 | .126 | .002 | .159 |
| | 1000 | 0 | .022 | .049 | .091 | .119 | .020 | .039 | .082 | .131 | .397 | .898 |
| | | .5 | .022 | .041 | .153 | .151 | .014 | .032 | .126 | .164 | .087 | .550 |
| | | .9 | .030 | .037 | .135 | .131 | .021 | .033 | .140 | .143 | .003 | .214 |
| | 200 | 0 | .509 | .828 | - | - | .734 | .925 | - | - | .976 | - |
| | | .5 | .417 | .709 | - | - | .645 | .846 | - | - | .721 | - |
| | | .9 | .277 | .486 | - | - | .421 | .571 | - | - | .319 | - |
| 30 | 500 | 0 | .342 | .813 | - | - | .887 | .982 | - | - | .966 | - |
| | | .5 | .298 | .701 | - | - | .810 | .937 | - | - | .673 | - |
| | | .9 | .224 | .499 | - | - | .550 | .693 | - | - | .318 | - |
| | 1000 | 0 | .278 | .802 | - | - | .952 | .996 | - | - | .951 | - |
| | | .5 | .237 | .687 | - | - | .930 | .980 | - | - | .642 | - |
| | | .9 | .156 | .466 | - | - | .653 | .799 | - | - | .296 | - |

Table 4: Monte Carlo rejection probabilities with 1000 repetitions for Design 4: $\boldsymbol{\mu} \not\leq \mathbf{0}$, $\boldsymbol{\mu} \not\geq \mathbf{0}$, with symmetrically distributed errors ($t(4)/\sqrt{2}$). For the cases that $n = 400$, $\mu_j = 0.02$ if $j \leq 0.1p$ and $\mu_j = -0.75$ if $j > 0.1p$, and for the cases that $n = 30$, $\mu_j = 0.3$ if $j \leq 10$ and $\mu_j = -5$ if $j > 10$. The columns represent the Empirical Bootstrap (EB) and Symmetry Randomization (SR) tests, based on the t_{\max} , t_{\max}^t and T_+ statistics, both with pre-selection (sel) and without pre-selection.

| γ | p | ρ | t_{\max} | | | | t_{\max}^t | | | | T_+ | |
|----------|------|--------|------------|------|-----------|-----------|--------------|------|-----------|-----------|-------|-----------|
| | | | EB | SR | EB sel | SR sel | EB | SR | EB sel | SR sel | SR | SR sel |
| | 200 | .0 | .116 | .069 | .104 | .089 | .086 | .077 | .082 | .106 | .082 | .081 |
| | | .5 | .107 | .083 | .098 | .063 | .081 | .061 | .072 | .083 | .084 | .073 |
| | | .9 | .069 | .053 | .075 | .050 | .056 | .048 | .057 | .059 | .070 | .059 |
| -.667 | 500 | .0 | .119 | .087 | .115 | .093 | .124 | .084 | .081 | .135 | .081 | .098 |
| | | .5 | .120 | .086 | .097 | .082 | .102 | .058 | .080 | .093 | .074 | .090 |
| | | .9 | .064 | .058 | .074 | .071 | .060 | .056 | .058 | .059 | .056 | .048 |
| | 1000 | .0 | .136 | .087 | .146 | .093 | .153 | .091 | .106 | .136 | .000 | .000 |
| | | .5 | .109 | .075 | .112 | .093 | .109 | .104 | .089 | .121 | .068 | .085 |
| | | .9 | .065 | .058 | .087 | .056 | .059 | .065 | .074 | .067 | .054 | .052 |
| | 200 | .0 | .031 | .025 | .034 | .024 | .017 | .034 | .032 | .020 | .043 | .027 |
| | | .5 | .038 | .034 | .034 | .031 | .033 | .040 | .042 | .042 | .035 | .036 |
| | | .9 | .042 | .037 | .050 | .043 | .045 | .048 | .034 | .042 | .045 | .053 |
| .667 | 500 | .0 | .020 | .019 | .023 | .033 | .018 | .024 | .017 | .024 | .021 | .019 |
| | | .5 | .033 | .024 | .037 | .041 | .026 | .036 | .027 | .028 | .023 | .027 |
| | | .9 | .030 | .037 | .039 | .042 | .049 | .048 | .054 | .050 | .039 | .047 |
| | 1000 | .0 | .033 | .025 | .021 | .030 | .019 | .028 | .016 | .027 | .000 | .000 |
| | | .5 | .029 | .036 | .029 | .021 | .034 | .031 | .025 | .028 | .038 | .034 |
| | | .9 | .059 | .042 | .036 | .048 | .043 | .052 | .051 | .041 | .034 | .046 |

Table 5: Monte Carlo rejection probabilities with 1000 repetitions for Design 1: $\boldsymbol{\mu} = \mathbf{0}$, $n = 400$ observations, asymmetrically distributed errors (skew-normal with skewness parameter γ). The columns represent the Empirical Bootstrap (EB) and Symmetry Randomization (SR) tests, based on the t_{\max} , t_{\max}^t and T_+ statistics, both with pre-selection (sel) and without pre-selection.