University of Groningen

Methods for High-Dimensional Data in Econometrics

Koning, Nick

DOI:
[10.33612/diss.168716962](https://doi.org/10.33612/diss.168716962)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2021

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Koning, N. (2021). *Methods for High-Dimensional Data in Econometrics*. [Thesis fully internal (DIV), University of Groningen]. University of Groningen, SOM research school.
<https://doi.org/10.33612/diss.168716962>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Methods for High-Dimensional Data in Econometrics

Nick W. Koning

Published by: University of Groningen, Groningen, The Netherlands.

Printed by: Ipskamp Printing

© 2020 Nick W. Koning

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system of any nature, or transmitted in any form or by any means, electronic, mechanical, now known or hereafter invented, including photocopying or recording, without prior written permission of the publisher.

**rijksuniversiteit
 groningen**

Methods for High-Dimensional Data in Econometrics

Proefschrift

ter verkrijging van de graad van doctor aan de
 Rijksuniversiteit Groningen
 op gezag van de
 rector magnificus Prof. C. Wijmenga
 en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op
 maandag 10 mei 2021 om 16.15 uur

door

Niklaas Wridzer Koning

geboren op 12 januari 1993
 te Groningen

Promotor

Prof. dr. P.A. Bekker

Copromotor

Dr. T. Boot

Beoordelingscommissie

Prof. dr. G.J. van den Berg

Prof. dr. F. Windmeijer

Prof. dr. C. Zhou

Acknowledgements

I am grateful for all the support I have received during the writing of this thesis.

To start, I would like to thank my supervisors Paul and Tom. Paul for his role as my promoter and daily supervisor. Without his encouragement, I would have likely never pursued a PhD in the first place. I am grateful for his attention to detail, the analytical skills he taught me, and his high expectations which pushed me to improve. And also for his afternoon visits after teaching, to chat about his altcoin portfolio or whatever was on his mind. Tom, for his unbreakable enthusiasm, great feedback, inspiring passion for econometrics, and support during my job market process.

I thank the members of the reading committee, consisting Gerard van den Berg, Frank Windmeijer and Chen Zhou for their time and their willingness to assess this thesis. I am especially grateful to Chen Zhou, for his careful examination of several proofs which led to the discovery of an error (which was luckily quickly resolved).

I express my gratitude to my current friends and former colleagues at the Faculty of Economics and Business: to my office nemesis Daniël, for our endless discussions that ripened into constructive conversation over the years; to Christian, for working together to become stronger men; to Ruben, for his genuineness; to Daan, for his contributions to lunch table discussion and war against inefficient climate policy; to Jos, for never being boring; to Mart, for hosting many board game nights and his stimulating ideas; to Steffen, for our teaching adventures and discussions about card games; to Tejas, for his frequent visits; to Johannes, for his volleyball skills; to Anouk, Gianmaria, Maite, Tobi, Daan, Juliette, Stefan, Jerry, Dennis, Albert, Michiel,

and all other colleagues, for creating a great working environment.

I would like to thank my friends for their support: Jochem, for our weekly squash matches, our discussions about machine learning and our friendship; Lennart, for being an overall great friend and our countless evening-long conversations, that have continued despite the large time difference; Arjan, for being my oldest friend, and many good memories; Marissa, for her infinitely broad interests and wit; Jing, for staying in touch and her visits to Groningen.

Finally, I thank my parents, Renger and Marian, for their unconditional support; my brother Jesse, for growing closer over the past years; my uncle Fokko, for his inspiring cheerfulness; my great aunt Marthy, for her interest; my cousin Renger, for his calls; and the rest of my family, who I always enjoy being around.

Nick Koning

Contents

1	Introduction	1
1.1	Thesis outline	2
2	Sparse Unit-Sum Regression	9
2.1	Introduction	9
2.2	Orthogonal X	13
2.2.1	Properties of $Q(\beta^{(P,N,z)})$ as a function of P and N . . .	15
2.2.2	Properties of $Q(\beta^{(p_z,n_z,z)})$ as a function of z for $p_z + n_z \leq k$	16
2.2.3	The case that $p_s + n_s > k$	18
2.2.4	Extension	20
2.3	Sparsity Under Orthogonality	20
2.4	General Case	21
2.4.1	Discrete First-Order Algorithm	22
2.4.2	Mixed-Integer Optimization	24
2.5	Numerical Results	25
2.5.1	Setup Simulation Experiments	25
2.5.2	Implementation and Stopping Criteria	27
2.5.3	Results of Simulation Experiments	27
2.6	Application: Index Tracking	28
2.7	Appendix	31
3	Exact Testing of Many Moment Inequalities against Multiple Violations	37
3.1	Introduction	37

3.2	The model and two test statistics	40
3.3	Symmetry based inference	43
3.4	Randomization tests for moment inequalities	44
3.5	Symmetry based inference with pre-selection	47
3.6	Simulations	48
3.6.1	Results	50
4	Testing Against Specific Alternatives with Conic Statistics	61
4.1	Introduction	61
4.1.1	Comparison to the Power Enhancement Technique	65
4.1.2	Notation	65
4.2	From Quadratic to Conic statistics	66
4.2.1	Existence of T_C	68
4.2.2	Geometric interpretation	68
4.2.3	Computation	69
4.3	Hypotheses	69
4.3.1	Hypotheses	71
4.3.2	Specific alternatives	72
4.3.3	Diagonal covariance matrices and scones	73
4.4	Randomization Test in Location Model	75
4.5	Sparse specific alternatives	77
4.5.1	Computation: full covariance matrix	78
4.5.2	Computation: diagonal covariance matrix	79
4.5.3	Lasso	80
4.6	Simulation experiments	82
4.6.1	Data generation	82
4.6.2	Tests	83
4.6.3	Results	84
4.7	Discussion	85
4.8	Appendix A	86
5	Conclusion and Discussion	95
	Samenvatting (Summary in Dutch)	105

Chapter 1

Introduction

The past decade has seen an incredible rise in the availability of data, not just in the number of observations but also in the level of detail captured by every observation. This data abundance has invigorated economists with a desire for exploration, to look for answer to questions that had previously seemed unanswerable. At the same time, the exploration of this data has lead to numerous new challenges, for which classical methodology originating in an era of data scarcity is often not suitable.

One such challenge is the presence of many variables, typically leading to models with a large number of parameters (p), potentially more than the number of observations (n). Such data is often referred to as being high-dimensional. For high-dimensional data, many traditional methods are not well-defined and inference often remains underdeveloped. In this dissertation, I therefore develop and discuss several techniques that are suited for high-dimensional applications. But before I introduce the chapters, I will briefly describe two concepts that will feature prominently in the chapters.

Perhaps currently the most prominent and best understood method for high-dimensional data is regularized linear regression. In linear regression, the idea of regularization generally refers to reducing the parameter space, based on prior information. This can be done either implicitly (through penalization) or explicitly, without the ex-ante removal of any variables.¹ In

¹ The latter condition is to distinguish it from standard modeling considerations.

practice it may be difficult to determine what parts of the parameter space to cut away. However, there exist several popular regularization methods that feature restrictions of the parameter space that economists may be willing to assume in many practical settings. For example, that the (Euclidean) length of the parameter vector is small (ℓ_2 -regularization), that the total magnitude of the parameter vector's elements is small (ℓ_1 -regularization), that the number of non-zero parameters is small (sparsity / ℓ_0 -regularization), or that the signs of the parameters are known. The latter three types are used extensively in the chapters of this dissertation.

Another common theme encountered when discussing methods designed for high-dimensional data is computational challenges. These challenges can be the result of the sheer size of the data, as some methods may not scale well when the number of variables increases. For regularization methods, this may be because the reduced parameter space is not easy to optimize over. Another computation related challenge is that there may not be computational algorithms available for the method at hand. For example, in Chapter 2 I discuss an existing model that I want to equip with regularization methods in order to make it more suitable for settings with high-dimensional data. However, it is not straightforward to merge the computation algorithm for the existing model with the algorithm for the regularization methods.

1.1 Thesis outline

The heart of this thesis is divided into three chapters that are introduced in the proceeding.

Chapter 2

This chapter discusses sparsity in unit-sum regression. Unit-sum regression is linear regression where the sum of the p regression coefficients must equal 1. It appears in numerous applications across different fields. Some of these applications are as follows.

- *Portfolio optimization.* In financial portfolio optimization one may be interested in constructing a portfolio from a selection of assets, that replicates some particular asset in the sense that it minimizes the squared difference between the portfolio returns and asset returns. Given historical returns data, this problem can be represented as a unit-sum regression problem. Here, the unit-sum restriction represents a budget constraint and the regression coefficients represent the share of investment in each asset.
- *Synthetic control.* To evaluate the outcome of a policy, one would ideally observe a counterfactual universe in which the policy was not implemented, in order to compare the outcomes. As such a counterfactual universe is not observed, researchers may resort to comparing the outcomes of treated units, that were affected by the policy, to control units, that are similar but unaffected by the policy.

In some situations, it may be difficult to find sufficiently similar units in the control group. This led to the development of synthetic control methods, where such control units are synthesized as a (convex) combination of units from the control group. The weights for the combination can be estimated using unit-sum regression, based on historical data.

- *Forecast combinations.* When presented with several conflicting expert forecasts, a good strategy is to average them with equal weights (Timmermann, 2006). If the quality of the forecasters is diverse, then one may want to assign more weight to the higher quality forecasts. Such weights can be estimated using unit-sum regression, based on historical forecasting performance. Here, the unit-sum restriction ensures that if all the individual forecasts are unbiased, then the forecast combination is also unbiased.

Unit-sum regression also appears as a subroutine in various algorithms and estimations procedures. One prominent example of this is the iterative archetypal algorithm (Cutler and Breiman, 1994).

In the chapter, we focus in particular on sparsity in unit-sum regression. The concept of sparsity captures the idea that only some parameters are non-zero, or, more generally, that a process is driven by a small number of large effects. In light of the above, one may know, for example, that their preferred portfolio contains only a small number of assets, or that there are only a few good (or many bad) forecasters in the pool of forecasters. Another benefit is that assuming a sufficient level of sparsity can ensure that a unique solution exists, even if the number of parameters (p) exceeds the number of observations (n); a situation where standard unit-sum regression will never yield a unique solution.

Imposing sparsity can be viewed as a form of regularization (ℓ_0 -type), as it cuts away insufficiently sparse vectors from the parameter space. As ℓ_0 -regularization is traditionally viewed as computationally difficult, especially in high-dimensional settings, it is often replaced in standard linear regression by ℓ_1 -regularization or Lasso. This places an upper bound constraint on the sum of the magnitude of the coefficients. In the presence of a unit-sum restriction, this results in a peculiar problem: because the sum of the coefficients must be 1, the sum of their magnitudes can never be smaller than 1. So, if the sum of the magnitudes in the unregularized problem is already 1, then adding ℓ_1 -regularization has no effect whatsoever. More generally, in unit-sum regression ℓ_1 -regularization may not be able to provide sparsity nor find a unique solution, which were the main motivations to use it in the first place.

As a solution, we propose to combine ℓ_0 - and ℓ_1 -regularization in unit-sum regression. This combination was already used by Mazumder et al. (2017) in standard linear regression, showing promising results. This solves the above described issues encountered with ℓ_1 -regularization, and retains its benefits. In addition, it allows the user to directly specify the desired sparsity level. In the paper, we develop an algorithm that employs state-of-the-art developments from the optimization literature. We explore the merits and demerits of this methodology in a simulation experiment, and find that it is able to achieve a similar or better performance compared to

ℓ_1 regularization, with substantially sparser solutions.

Chapter 3

The second chapter features the testing of moment inequalities. Such tests play an important role in inference for partially identified models.

One application is found in discrete choice models, where the moment inequalities arise from revealed preference. For example, a consumer may choose purchase a good at a particular store, instead of at a cheaper but more distant store or at a closer but more expensive store. If the consumer is aware of the other prices, this implicitly bounds the perceived cost of travel. Pakes (2010) describes how observing many of such consumer choices leads to a collection of moment inequalities, which can be used to partially identify such a perceived cost of travel parameter. A confidence interval can then be constructed by testing these moment inequalities at different values of the parameter, and recording where the test does not reject.

Chernozhukov et al. (2019a) provide further examples, including a market structure model by Ciliberto and Tamer (2009), where the number of moment inequalities equals $p = 2^{m+1}$, where m is the number of firms. They emphasize that the number of inequalities may easily become very large in many applications. For that reason, they propose tests using the maximum of the Studentized sample moments (t statistics) as test statistic, and show the validity even if p grows exponentially in n .

In this chapter, we observe that as a result of only considering the largest t statistic, their approach is well suited for testing against a single violation of the moment inequalities (i.e. a 1-sparse alternative). But by ignoring all but the largest t statistic, it may suffer from low power against alternatives where multiple inequalities are simultaneously violated.

In order to combat this, we propose a test statistic based on a non-negatively weighted combination of t statistics. Within the framework of Chernozhukov et al. (2019a) one way to view this is that we expand the set of moment inequalities with even more moment inequalities, that are implied by the original set of moment inequalities. Alternatively, it can be viewed as in-

creasing the collection of t statistics over which Chernozhukov et al. (2019a) take a maximum, by adding synthetic t statistics that are a non-negative combination of the original t statistics and are large if multiple moment inequalities are violated. We focus on two cases: one where the expanded set of t statistics remains finite, where the results of Chernozhukov et al. (2019a) remain applicable. And another where all non-negative weights are permitted, which results in an uncountably infinite set of t statistics.

We propose inference using a randomization test based on a reflection symmetry assumption and provide a sufficient condition for size control in large samples. In addition, simulation results show that the tests we discuss also control size in small samples ($n = 30$, $p = 1000$), and often have substantially higher power against alternatives with multiple violations than tests based on the maximum t statistic.

Chapter 4

In the final chapter, I expand some ideas from the second chapter beyond the framework of moment inequalities and link them to problems that occur with quadratic statistics when they are used for testing in high-dimensional settings.

Quadratic statistics (such as the Wald statistic) are popular in classical low-dimensional settings. However, in high-dimensional settings such tests are known to have low power against specific alternatives (Fan et al., 2015). This has provoked recent interest in constructing tests that direct power towards specific alternatives of interest (Fan et al., 2015; Kock and Preinerstorfer, 2019). In the chapter, I focus in particular on testing against specific alternatives that can be described as a (not necessarily convex) cone.

Specific alternatives arise if the econometrician has some prior information about the nature of possible violations from the null hypothesis. For example, one may want to test the null hypothesis that a collection of equalities is satisfied, against the alternative that they are not all satisfied. Here, the econometrician could have a prior information that under the alternative, only a small but unknown subset of these equalities is violated. In

practice, the collection of equalities could represent some economic theory, and the unknown subset of potentially violated equalities could represent the presence of group of exceptional individuals.

Such a specific alternative is typically called sparse, as the vector of deviations from the null hypothesis is sparse. A test based on a quadratic statistic is purely based on the (weighted Euclidean) length of the deviations, and therefore typically poor at detecting a sparse violation. In addition, it is not obvious how to incorporate this information about sparsity when testing with a quadratic statistic.

Another issue with quadratic statistics is that they are often weighted by an inverse covariance matrix, which must be estimated in practice. This is problematic, as the standard sample covariance matrix is not invertible in high-dimensional settings, so that such a statistic may not exist. Here, practitioners may have to resort to regularization or other assumptions in order to even have a test statistic to test with.

In this chapter, I attempt to find a practical solution to these issues by considering a generalization of quadratic statistics, which nests many well-known statistics, including the statistics discussed in Chapter 3. The statistic requires the specification of a cone, which I set equal to the specific alternative of interest. The existence of the statistic depends on a restricted eigenvalue condition, similar to the conditions used to prove oracle properties for Lasso. This condition is weaker than non-singularity of the covariance matrix estimator. It can therefore be used in high-dimensional settings, provided that the cone is sufficiently small. In addition, I discuss how the statistic can be computed using cone-restricted quadratic optimization.

For the standard location model, I construct a randomization critical value based on a reflection symmetry assumption and provide an asymptotic condition that guarantees size control, generalizing a condition used in Chapter 3. I illustrate the test by testing against a sparse specific alternative. In a simulation study, the test performs favorably compared to a test based on a quadratic statistic and the power enhancement technique of Fan et al. (2015).

Chapter 2

Sparse Unit-Sum Regression

2.1 Introduction

Linear regression with coefficients that sum to one (henceforth *unit-sum regression*) is used in portfolio optimization and other economic applications such as forecast combinations (Timmermann, 2006) and synthetic control (Abadie et al., 2010).

In this paper, we focus on obtaining a sparse solution (i.e. containing few non-zero elements) to the unit-sum regression problem. A sparse solution may be desirable for a variety of reasons, such as making a model more interpretable, improving estimation efficiency if the underlying parameter vector is known to be sparse, remedying identification issues if the number of variables exceeds the number of observations, or application-specific reasons such as reducing cost by limiting the amount of constituents in a portfolio.

A popular method to produce sparsity is to use regularization. Theoretically, the most straightforward way to obtain a sparse solution is to use ℓ_0 -regularization (also known as best-subset selection), which amounts to restricting the number of non-zero elements in the solution. However, the

This chapter is based on Koning and Bekker (2018).

use of ℓ_0 -regularization is NP-hard (Coleman et al., 2006; Natarajan, 1995) and has traditionally been seen as computationally infeasible for problems with more than about 40 variables, both in unit-sum regression and standard linear regression.

Due to these computational difficulties, ℓ_0 -regularization has often been replaced by ℓ_1 -regularization, also known as Lasso (Tibshirani, 1996). In ℓ_1 -regularization, the ℓ_0 -norm restriction that restricts the number of non-zero elements is replaced by an ℓ_1 -norm restriction that restricts the absolute size of the coefficients. This turns the problem into an easier to solve convex optimization problem. An ℓ_1 -norm restriction shrinks the weights towards zero and, as a consequence of the shrinkage, produces sparsity by setting some weights exactly equal to zero.

The use of ℓ_1 -regularization in the presence of a unit-sum restriction was first considered by DeMiguel et al. (2009) and Brodie et al. (2009) in the context of portfolio optimization. They show that ℓ_1 -regularization is able to produce sparsity in combination with a unit-sum restriction. In addition, they demonstrate that the combination can be viewed as a restriction on the sum of the negative weights. In some applications it is highly desirable to have a parameter that explicitly controls the sum of the negative weights. For example, in a portfolio optimization context negative weights represent potentially costly short positions.

However, the unit-sum restriction can cause a problem when using ℓ_1 -regularization: due to the unit-sum restriction the ℓ_1 -norm of the weights cannot be smaller than 1. This imposes a lower bound on the amount of shrinkage produced by ℓ_1 -regularization. In turn, this places an upper bound on the sparsity produced by ℓ_1 -regularization. This upper bound depends entirely on the data, which makes it difficult to rely on ℓ_1 -regularization if a specific level of sparsity is desired. In addition, due to the bound there does not always exist a value of the tuning parameter that guarantees the existence of a unique solution. Furthermore, Fastrich et al. (2014) point out that a combination of a non-negativity restriction and a unit-sum restriction fixes the ℓ_1 -norm of the weights to 1, which renders ℓ_1 -regularization useless.

In order to address these issues and obtain sparse solutions in unit-sum regression, we use a recent innovation in ℓ_0 -regularization in the standard linear regression setting by Bertsimas et al. (2016). They show that modern Mixed-Integer Optimization (MIO) solvers can find a provably optimal solution to ℓ_0 -regularized regression for problems of practical size. To achieve this, the solver is provided with a good initial solution obtained from a discrete first-order (DFO) algorithm. In a simulation study, they show that ℓ_0 -regularization performs favorably compared to ℓ_1 -regularization in terms of predictive performance and sparsity.

An extended simulation study comparing ℓ_0 - and ℓ_1 -regularization in the standard linear regression setting is performed by Hastie et al. (2017). They find that ℓ_0 -regularization outperforms ℓ_1 -regularization if the signal-to-noise ratio (SNR) is high, while ℓ_1 -regularization performs better if the SNR is low. Additionally, they find that if the tuning parameters are selected to optimize predictive performance, ℓ_0 -regularization yields substantially sparser solutions.

A combination of ℓ_0 - and ℓ_1 -regularization ($\ell_0\ell_1$ -regularization) is studied in the standard linear regression context by Mazumder et al. (2017). They observe that this combination yields a predictive performance similar to ℓ_1 -regularization if the SNR is low, and a predictive performance similar to ℓ_0 -regularization if the SNR is high. In addition, they find that $\ell_0\ell_1$ -regularization produces more sparsity compared to ℓ_1 -regularization, if the tuning parameters are selected in order to optimize predictive performance.

Motivated by the results in the standard linear regression setting, we propose the use of $\ell_0\ell_1$ -regularization in unit-sum regression. Specifically, let $\mathbf{y}$ be a t -vector and let $\mathbf{X}$ be a $t \times m$ matrix, then we consider the problem

$$\min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2, \quad \text{s.t.} \quad \sum_{i=1}^m \beta_i = 1, \quad \|\boldsymbol{\beta}\|_0 \leq k, \quad \|\boldsymbol{\beta}\|_1 \leq 1 + 2s, \quad (2.1)$$

where β_i are the elements of $\boldsymbol{\beta}$, $\|\boldsymbol{\beta}\|_0 = \sum_{i=1}^m 1_{\{\beta_i \neq 0\}}$ is the ℓ_0 -norm of $\boldsymbol{\beta}$, $\|\boldsymbol{\beta}\|_1 = \sum_{i=1}^m |\beta_i|$ is the ℓ_1 -norm of $\boldsymbol{\beta}$, $s \geq 0$ and $1 \leq k \leq m$. Notice that this problem is equivalent to ℓ_0 -regularized unit-sum regression if s is suffi-

ciently large, and equivalent to ℓ_1 -regularized unit-sum regression if $k = m$.

The formulation in (2.1) provides users with explicit control over both the sparsity of the solution and the sum of the negative weights of the solution. In addition, if the tuning parameters are selected in order to maximize predictive performance, we find in a simulation experiment that $\ell_0\ell_1$ -regularization:

- performs better than ℓ_0 -regularization in terms of predictive performance, especially if the signal-to-noise ratio is low.
- performs well compared to ℓ_1 -regularization in terms of predictive performance, especially for higher signal-to-noise ratios, while at the same time producing much sparser solutions.

The main contributions of this paper can be summarized as follows. [1] We propose $\ell_0\ell_1$ -regularization for the unit-sum regression problem. [2] We analyze the problem for orthogonal design matrices and provide an algorithm to compute its solution. [3] We show how the algorithm for the orthogonal design case can be used in finding a solution to the general problem by extending the framework of Bertsimas et al. (2016) to unit-sum regression. [4] We perform a simulation experiment which shows that our approach performs favorably compared to ℓ_0 -regularization or ℓ_1 -regularization. [5] We demonstrate in an application to stock index tracking that we are able to find substantially sparser portfolios with $\ell_0\ell_1$ -regularization than with ℓ_1 -regularization, while maintaining a similar out-of-sample tracking error.

The remainder of the paper is structured as follows. In Section 2.2, problem (2.1) is studied under the assumption that X is orthogonal and an algorithm for the orthogonal case is presented. Section 2.3 analyzes the sparsity production for the orthogonal case and yields some intuitions about the problem. Section 2.4 links the algorithm for the orthogonal case to the framework of Bertsimas et al. (2016) in order to find a solution to the general problem. In Section 2.5, the simulation experiments are presented. Section 2.6 provides an application to index tracking.

2.2 Orthogonal X

As problem (2.1) is difficult to study in its full generality, we first consider the special case that X is orthogonal and $t = m$. We derive properties of a solution to (2.1) under orthogonality and use these properties in order to construct an algorithm that finds a solution. The algorithm is presented at the end of the section. In Section 2.4 this algorithm is used in constructing an algorithm to find a solution to the general problem, by extending the framework of Bertsimas et al. (2016). In Section 2.3 we analyze the sparsity of the solution under orthogonality.

Assume that $X'X = XX' = I_m$, where I_m is the $m \times m$ identity matrix. Let us write $\eta = X'y$, so that minimizing $\|y - X\beta\|_2^2$ in β is equivalent to minimizing $\|X'y - \beta\|_2^2 = \|\eta - \beta\|_2^2 =: Q(\beta)$. Define

$$\mathcal{T} := \left\{ \beta \in \mathbb{R}^m \mid \sum_{i=1}^m \beta_i = 1, \|\beta\|_0 \leq k, \|\beta\|_1 \leq 1 + 2s \right\}. \quad (2.2)$$

Then, problem (2.1) can be written as $\min_{\beta \in \mathcal{T}} Q(\beta)$.

We assume the elements of η are different and $k < m$. Without further loss of generality we assume $\eta_1 > \eta_2 > \dots > \eta_m$. In Section 2.2.4 we relax the assumption that $k < m$ and allow for $k \leq m$.

Let

$$\mathcal{A}_z := \left\{ \beta \in \mathbb{R}^m \mid \sum_{i=1}^m \beta_i = 1, \|\beta\|_0 \leq k, \|\beta\|_1 = 1 + 2z \right\},$$

where $0 \leq z$, so that $\mathcal{T} = \cup_{0 \leq z \leq s} \mathcal{A}_z$. If $\hat{\beta} \in \arg \min_{\beta \in \mathcal{T}} Q(\beta)$, then $\hat{\beta} \in \mathcal{A}_z$ for some $z \leq s$. Let us denote this value of z with $\hat{z}$. We will now show that $\hat{\beta}$ can be computed from the signs of its elements and $\hat{z}$. In order to show this, we first solve a related problem and then show that $\hat{\beta}$ is equal to the solution of a specific case of this related problem.

Let $\mathcal{P} \subseteq \mathcal{M}$ and $\mathcal{N} \subset \mathcal{M}$ be disjoint sets with cardinalities p and n ,

respectively, where $\mathcal{M} := \{1, \dots, m\}$. Define

$$\mathcal{B}^{(\mathcal{P}, \mathcal{N}, z)} := \left\{ \beta \in \mathbb{R}^m \mid \sum_{i \in \mathcal{P}} \beta_i = 1 + z, \sum_{i \in \mathcal{N}} \beta_i = -z, \text{ and } \beta_i = 0 \text{ if } i \in (\mathcal{P} \cup \mathcal{N})^c \right\}.$$

Minimization of $Q(\beta)$ over the affinely restricted set $\mathcal{B}^{(\mathcal{P}, \mathcal{N}, z)}$ has the solution

$$\beta^{(\mathcal{P}, \mathcal{N}, z)} := \arg \min_{\beta \in \mathcal{B}^{(\mathcal{P}, \mathcal{N}, z)}} Q(\beta) = \begin{cases} \eta_i - \frac{(\sum_{j \in \mathcal{P}} \eta_j) - 1 - z}{p}, & i \in \mathcal{P}, \\ \eta_i - \frac{(\sum_{j \in \mathcal{N}} \eta_j) + z}{n}, & i \in \mathcal{N}, \\ 0, & i \in (\mathcal{P} \cup \mathcal{N})^c. \end{cases} \quad (2.3)$$

Recall that $\hat{\beta} \in \arg \min_{\beta \in \mathcal{T}} Q(\beta)$ and let $\hat{\mathcal{P}} := \{i \mid \hat{\beta}_i > 0\}$ and $\hat{\mathcal{N}} := \{i \mid \hat{\beta}_i < 0\}$. Furthermore, let $\mathcal{C}$ be the set of vectors β with elements that have the same signs as the elements of $\hat{\beta}$, then $\mathcal{A}_z \cap \mathcal{C} \subseteq \mathcal{B}^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}$. Notice that the difference between $\mathcal{A}_z \cap \mathcal{C}$ and $\mathcal{B}^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}$ is that there are no sign restrictions on elements $\beta_i \in \mathcal{B}^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}$, for which $i \in (\mathcal{P} \cup \mathcal{N})$. Consequently,

$$Q(\hat{\beta}) = \min_{\beta \in \mathcal{A}_z \cap \mathcal{C}} Q(\beta) \geq \min_{\beta \in \mathcal{B}^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}} Q(\beta) = Q(\beta^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}).$$

However, if $\hat{\beta} \neq \beta^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}$, then $\beta(\phi) := \phi \beta^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})} + (1 - \phi) \hat{\beta} \in \mathcal{A}_z \cap \mathcal{C}$ for sufficiently small $\phi > 0$. Furthermore, as $Q(\beta(\phi))$ is a parabola in ϕ with a minimum at $\lambda = 1$, we find that $Q(\beta(\phi)) < Q(\hat{\beta})$ for small $\phi > 0$. As $\hat{\beta} \in \arg \min_{\beta \in \mathcal{A}_z} Q(\beta)$, this is a contradiction. Hence, $\hat{\beta} = \beta^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}$, which is our first result.

Proposition 1. *If $\hat{\beta} \in \arg \min_{\beta \in \mathcal{T}} Q(\beta)$, then $\hat{\beta} = \beta^{(\hat{\mathcal{P}}, \hat{\mathcal{N}}, \hat{z})}$.*

So, the problem can be decomposed into finding the components of the triplet $(\mathcal{P}, \mathcal{N}, z)$ that minimizes $Q(\beta^{(\mathcal{P}, \mathcal{N}, z)})$. Next, we will study the properties of these components.

2.2.1 Properties of $Q(\beta^{(\mathcal{P}, \mathcal{N}, z)})$ as a function of $\mathcal{P}$ and $\mathcal{N}$

The sorting of η reveals an ordered structure in the sets $\mathcal{P}$ and $\mathcal{N}$ that minimize $Q(\beta^{(\mathcal{P}, \mathcal{N}, z)})$. This structure is described in the following result.

Proposition 2. *If $\beta^{(\mathcal{P}, \mathcal{N}, z)} \in \arg \min_{\beta \in \mathcal{A}_z} Q(\beta)$, then $\mathcal{P} = \{1, \dots, p\}$ and $\mathcal{N} = \{m - n + 1, \dots, m\}$ if $n \geq 1$, and $\mathcal{N} = \emptyset$ if $n = 0$.*

The proof is given in the Appendix. For sets such as $\mathcal{P} = \{1, \dots, p\}$ and $\mathcal{N} = \{m - n + 1, \dots, m\}$, we use the notation $\beta^{(p, n, z)} := \beta^{(\mathcal{P}, \mathcal{N}, z)}$, as in (2.3). The following result shows that p and n should be maximized such that $\beta^{(p, n, z)} \in \mathcal{A}_z$.

Lemma 1. *If $\beta^{(\tilde{p}, \tilde{n}, z)} \in \mathcal{A}_z$ and $\beta^{(p, n, z)} \in \mathcal{A}_z$, where $\tilde{p} \leq p$, $\tilde{n} \leq n$, $\tilde{p} + \tilde{n} < p + n$, then $Q(\beta^{(p, n, z)}) < Q(\beta^{(\tilde{p}, \tilde{n}, z)})$.*

The proof is given in the Appendix.

We will now consider the relationship between z and the pair (p, n) . With reference to (2.3), let us consider the sets

$$\mathcal{P}_z := \left\{ q \left| \eta_q - \frac{(\sum_{i=1}^q \eta_i) - 1 - z}{q} > 0 \right. \right\}, \quad (2.4)$$

$$\mathcal{N}_z := \left\{ q \left| \eta_{m-q+1} - \frac{(\sum_{i=1}^q \eta_{m-i+1}) + z}{q} < 0 \right. \right\}, \quad (2.5)$$

with cardinalities $|\mathcal{P}_z| = p_z$ and $|\mathcal{N}_z| = n_z$. As

$$\begin{aligned} \eta_q - \frac{(\sum_{i=1}^q \eta_i) - 1 - z}{q} &= \frac{q-1}{q} \left(\eta_q - \frac{(\sum_{i=1}^{q-1} \eta_i) - 1 - z}{q-1} \right) \\ &< \eta_{q-1} - \frac{(\sum_{i=1}^{q-1} \eta_i) - 1 - z}{q-1}, \end{aligned}$$

we find $\mathcal{P}_z = \{1, \dots, p_z\}$, and similarly $\mathcal{N}_z = \{m - n + 1, \dots, m\}$ if $z > 0$ and $\mathcal{N}_z = \emptyset$ if $z = 0$. Additionally, we find that p_z is increasing in z , and similarly that n_z is increasing in z . So, by Lemma 1 we have following result for $p_z + n_z \leq k$.

Proposition 3. If $\beta^{(\mathcal{P}, \mathcal{N}, z)} \in \arg \min_{\beta \in \mathcal{A}_z} Q(\beta)$ and $p_z + n_z \leq k$, then $\beta^{(\mathcal{P}, \mathcal{N}, z)} = \beta^{(p_z, n_z, z)}$.

We will now analyze how $Q(\beta^{(p_z, n_z, z)})$ varies with z if $p_z + n_z \leq k$, and use this to find a minimizer $\hat{\beta} \in \arg \min_{\beta \in \mathcal{T}} Q(\beta)$ if $p_s + n_s \leq k$. The case that $p_s + n_s > k$ is treated separately in Section 2.2.3.

2.2.2 Properties of $Q(\beta^{(p_z, n_z, z)})$ as a function of z for $p_z + n_z \leq k$

As p_z and n_z are integers, they increase discontinuously as z increases. In this subsection we show that $Q(\beta^{(p_z, n_z, z)})$ and its derivative are continuous in z despite these discontinuities in p_z and n_z . This will allow us to show that $\beta^{(p_s, n_s, s)} \in \arg \min_{\beta \in \mathcal{T}} Q(\beta)$ if $p_s + n_s \leq k$.

Let $z_1^+ = -1$ and $z_p^+ = z_{p-1}^+ + (p-1)(\eta_{p-1} - \eta_p) = \sum_{i=1}^{p-1} (\eta_i - \eta_p) - 1$, for $p = 2, \dots, m$. We then find the ordering $z_1^+ < z_2^+ < \dots < z_m^+$, and

$$\begin{aligned} \eta_p &= \frac{(\sum_{i=1}^p \eta_i) - 1 - z_p^+}{p}, \quad p = 1, \dots, m, \\ \eta_{p+1} &= \frac{(\sum_{i=1}^p \eta_i) - 1 - z_{p+1}^+}{p} = \frac{(\sum_{i=1}^{p+1} \eta_i) - 1 - z_{p+1}^+}{p+1}, \quad p = 1, \dots, m-1. \end{aligned} \quad (2.6)$$

Consequently, if $z_p^+ < z \leq z_{p+1}^+$, then

$$\eta_p > \frac{(\sum_{i=1}^p \eta_i) - 1 - z}{p} \geq \eta_{p+1}. \quad (2.7)$$

Similarly, let $z_m^- = 0$ and $z_{m-n+1}^- = z_{m-(n-1)+1}^- + (n-1)(\eta_{m-n+1} - \eta_{m-(n-1)+1}) = \sum_{i=1}^{n-1} (\eta_{m-n+1} - \eta_{m-i+1})$, for $n = 2, \dots, m$. Then $z_{m-1+1}^- < z_{m-2+1}^- < \dots < z_{m-m+1}^-$ and

$$\begin{aligned} \eta_{m-n+1} &= \frac{(\sum_{i=1}^n \eta_{m-i+1}) + z_{m-n+1}^-}{n}, \quad n = 1, \dots, m, \\ \eta_{m-n} &= \frac{(\sum_{i=1}^n \eta_{m-i+1}) + z_{m-n}^-}{n} = \frac{(\sum_{i=1}^{n+1} \eta_{m-i+1}) + z_{m-n}^-}{n+1}, \quad n = 1, \dots, m-1. \end{aligned} \quad (2.8)$$

Consequently, if $z_{m-n+1}^- < z \leq z_{m-n}^-$, then

$$\eta_{m-n+1} < \frac{(\sum_{i=1}^n \eta_{m-i+1}) + z}{n} \leq \eta_{m-n}. \quad (2.9)$$

Using the cardinalities p_z and n_z of the sets $\mathcal{P}_z$ and $\mathcal{N}_z$ in (2.4) and (2.5), let $z_m := \min\{z \mid p_z + n_z = m\}$. If $0 \leq z < z_m$, then $z_{p_z}^+ < z \leq z_{p_z+1}^+$. If $0 < z < z_m$, then $z_{m-n_z+1}^- < z \leq z_{m-n_z}^-$. The loss function

$$Q(\beta^{(p_z, n_z, z)}) = p_z \left\{ \frac{(\sum_{i=1}^{p_z} \eta_i) - 1 - z}{p_z} \right\}^2 + I_{\{n_z \geq 1\}} n_z \left\{ \frac{(\sum_{i=1}^{n_z} \eta_{m-i+1}) + z}{n_z} \right\}^2 + \sum_{i=p_z+1}^{m-n_z} \eta_i^2$$

is a continuous function of z for $0 \leq z \leq z_m$, with derivative

$$\frac{dQ(\beta^{(p_z, n_z, z)})}{dz} = -2 \left\{ \frac{(\sum_{i=1}^{p_z} \eta_i) - 1 - z}{p_z} \right\} + 2I_{\{n_z \geq 1\}} \left\{ \frac{(\sum_{i=1}^{n_z} \eta_{m-i+1}) + z}{n_z} \right\}, \quad (2.10)$$

which is continuous for $0 < z < z_m$. That is, using (2.6), if $z \uparrow z_{p_z+1}^+$, then

$$-\frac{(\sum_{i=1}^{p_z} \eta_i) - 1 - z}{p_z} \uparrow \eta_{p_z+1}$$

and if $z \downarrow z_{p_z+1}^+$, then

$$-\frac{(\sum_{i=1}^{p_z+1} \eta_i) - 1 - z}{p_z + 1} \downarrow \eta_{p_z+1}.$$

A similar continuity holds for the second term of (2.10) due to (2.8). The derivative (2.10) is increasing in z , but it is negative for $0 < z < z_m$ due to (2.7) and (2.9), which imply

$$\frac{dQ(\beta^{(p_z, n_z, z)})}{dz} \leq -2\eta_{p_z+1} + 2\eta_{m-n_z} \leq 0.$$

We summarize these results in a proposition.

Proposition 4. *The function $Q(\beta^{(p_z, n_z, z)})$ is continuous in z for $0 \leq z \leq z_m$, and the derivative with respect to z is negative for $0 < z < z_m$.*

As $Q(\beta^{(p_z, n_z, z)})$ is strictly decreasing in z over $0 < z \leq z_m$ if $p_z + n_z \leq k$, we conclude that $\beta^{(p_s, n_s, s)} \in \arg \min_{\beta \in \mathcal{T}} Q(\beta)$ if $p_s + n_s \leq k$.

2.2.3 The case that $p_s + n_s > k$

If $p_s + n_s > k$, then $\beta^{(p_s, n_s, s)} \notin \mathcal{T}$. So an alternative approach is required. By Lemma 1 and the fact that $p_s + n_s > k$, we should compare the objective values for all pairs (p, n) for which $p + n = k$, $p \leq p_s$ and $n \leq n_s$. In order to do so for a given pair (p, n) , we need to find the value of z that minimizes $Q(\beta^{(p, n, z)})$. This minimizing value, which we will denote by $\tilde{z}$, must satisfy $z_{pn}^* := \max\{z_p^+, z_{m-n+1}^-\} < \tilde{z} \leq s$. We will now show that $\tilde{z}$ is either equal to s or to $s_{pn} := \arg \min_z Q(\beta^{(p, n, z)})$.

We find

$$Q(\beta^{(p, n, z)}) = (p + n) \left(\frac{\sum_{i=1}^p \eta_i + \sum_{j=1}^n \eta_{m-j+1} - 1}{p + n} \right)^2 + \frac{p + n}{pn} (s_{pn} - z)^2 + \sum_{i=p+1}^{m-n} \eta_i,$$

where

$$s_{pn} = n \frac{(\sum_{i=1}^p \eta_i) - 1}{p + n} - p \frac{\sum_{i=1}^n \eta_{m-i+1}}{p + n}.$$

As $Q(\beta^{(p, n, z)})$ is quadratic in z with a minimum at s_{pn} , we find that if $z_{pn}^* < s \leq s_{pn}$, then $\tilde{z} = s$, and if $z_{pn}^* < s_{pn} < s$, then $\tilde{z} = s_{pn}$.

In the case that $s_{pn} \leq z_{pn}^*$, the minimum does not exist, since $Q(\beta^{(p, n, z)}) \downarrow Q(\beta^{(p, n, z_{pn}^*)})$ as $z \downarrow z_{pn}^*$. Furthermore, $\|\beta^{(p, n, z_{pn}^*)}\|_0 < k$. So if $p_{z_{pn}^*} + n_{z_{pn}^*} < k$ then $Q(\beta^{(p, n, z_{pn}^*)}) \geq Q(\beta^{(p_{z_{pn}^*}, n_{z_{pn}^*}, z_{pn}^*)})} > Q(\beta^{(p_z, n_z, z)})$ for some $z > z_{pn}^*$, by Proposition 3 and due to the negative gradient of $Q(\beta^{(p_z, n_z, z)})$. In the case that $p_{z_{pn}^*} + n_{z_{pn}^*} \geq k$, then $z_{pn}^* = z_p^+$ or $z_{pn}^* = z_n^-$, and $z_p^+ \neq z_n^-$. So if $z_{pn}^* = z_p^+$, then $Q(\beta^{(p, n, z_{pn}^*)}) = Q(\beta^{(p-1, n, z_{pn}^*)}) > Q(\beta^{(p-1, n+1, z_{pn}^*)})$ by Lemma

1. Similarly if $z_{pn}^* = z_n^-$, then $Q(\beta^{(p,n,z_{pn}^*)}) > Q(\beta^{(p+1,n-1,z_{pn}^*)})$. So if $s_{pn} \leq z_{pn}^*$, then $\beta^{(p,n,z)} \notin \arg \min_{\beta \in \mathcal{T}} Q(\beta)$ for all $z_{pn}^* < z \leq s$.

Hence, if $p_s + n_s > k$, we can compute $\tilde{z}$ for each pair (p, n) that satisfies $p + n = k$, $p \leq p_s$ and $n \leq n_s$ and use this to compute the objective value $Q(\beta^{(p,n,\tilde{z})})$. By comparing the objective values, we can find the triplet $(p, n, \tilde{z})$ for which $\beta^{(p,n,\tilde{z})} \in \arg \min_{\beta \in \mathcal{T}} Q(\beta)$.

Combining these findings with the findings from the previous sections, we can construct an algorithm to find an element of $\arg \min_{\beta \in \mathcal{T}} Q(\beta)$. This algorithm is presented in Algorithm 1.

Algorithm 1: Computing an element of $\arg \min_{\beta \in \mathcal{T}} \|\eta - \beta\|_2^2$

Input: Sorted m -vector η , parameters k and s .

Output: $\hat{\beta}$.

```

1  $\bar{p} = \arg \max_i (i \mid \sum_{j=1}^{i-1} (\eta_j - \eta_i) < 1 + s)$ 
2  $\bar{n} = \arg \max_i (i \mid \sum_{j=1}^{i-1} (\eta_{m-i+1} - \eta_{m-j+1}) < s)$ 
3 if  $\bar{p} + \bar{n} \leq k$  then
4    $\hat{\beta} = \beta^{(\bar{p}, \bar{n}, s)}$ 
5 else
6    $\mathcal{S} = \{(p, n) \mid p + n = k, p \leq \bar{p}, n \leq \bar{n}\}$ 
7   for  $(p, n) \in \mathcal{S}$  do
8      $s_{pn} = n \frac{(\sum_{i=1}^p \eta_i) - 1}{p+n} - p \frac{\sum_{i=1}^n \eta_{m-i+1}}{p+n}$ 
9     if  $s < s_{pn}$  then
10        $z_{pn} = s$ 
11     else
12        $z_{pn} = s_{pn}$ 
13      $Q_{pn} = \|\eta - \beta^{(p,n,z_{pn})}\|_2^2$ 
14    $(\hat{p}, \hat{n}) = \arg \min_{(p,n) \in \mathcal{S}} Q_{pn}$ 
15    $\hat{\beta} = \beta^{(\hat{p}, \hat{n}, z_{\hat{p}\hat{n}})}$ 

```

2.2.4 Extension

The case $k \leq m$ can be treated in a way similar to the case $k < m$, except that in the proof of Proposition 2 the assumption $k < m$ was needed. We therefore provide a proof for $k = m$.

Proposition 5. *Proposition 2 holds true when $k = m$.*

The proof is given in the Appendix.

2.3 Sparsity Under Orthogonality

In this section, we use the results from Section 2.2 to study the sparsity of the solution to (2.1) under orthogonality.

As both ℓ_0 - and ℓ_1 -regularization produce sparsity, we can analyze how the sparsity of the solution to (2.1) depends on the tuning parameters k and s . From Algorithm 1, it is straightforward to observe that the amount of non-zero elements in $\hat{\beta}$ is equal to $\min(k, \bar{p} + \bar{n})$, where $\bar{p} = \arg \max_i (i \mid \sum_{j=1}^{i-1} (\eta_j - \eta_i) < 1 + s)$ and $\bar{n} = \arg \max_i (i \mid \sum_{j=1}^{i-1} (\eta_{m-i+1} - \eta_{m-j+1}) < s)$. So the ℓ_1 -regularization component only produces additional sparsity if $k > \bar{p} + \bar{n}$.

In order to gain some insights into the sparsity produced by the ℓ_1 -regularization component, we consider the maximum sparsity produced by ℓ_1 -regularization if $k \geq \bar{p} + \bar{n}$. Notice that the sparsity is maximized if $\bar{p} + \bar{n}$ is minimized, which happens when $s = 0$. Furthermore, if $s = 0$, then $\bar{n} = 0$. So, the minimum number of non-zero elements is equal to

$$\begin{aligned}
 \min(\bar{p}, k) &= \bar{p} \\
 &= \arg \max_i \left(i \mid \sum_{j=1}^{i-1} (\eta_j - \eta_i) < 1 \right) \\
 &= \arg \max_i \left(i \mid \sum_{j=1}^{i-1} j \Delta_j < 1 \right), \tag{2.11}
 \end{aligned}$$

where $\Delta_j = \eta_j - \eta_{j+1}$. This shows that the maximum sparsity produced by ℓ_1 -regularization depends entirely on the size of the gaps between the $\bar{p} + 1$

largest elements of $\boldsymbol{\eta}$. So the maximum amount of sparsity does not change if the same constant is added to each element of $\boldsymbol{\eta}$.

To further analyze the maximum sparsity produced by the ℓ_1 -regularization component, we consider two special cases of $\boldsymbol{\eta}$: one case without noise and one case with noise.

Linear and Noiseless. Suppose that the $\bar{p} + 1$ largest elements of $\boldsymbol{\eta}$ are linearly spaced with distance $\Delta > 0$ (i.e. $\eta_i = a - (i - 1)\Delta$ for some a). Then, using (2.11), we can derive the following closed-form expression for the minimum number of non-zero elements:

$$\bar{p} = \left\lfloor \frac{1}{2} \left(\sqrt{\frac{\Delta + 8}{\Delta}} + 1 \right) \right\rfloor,$$

where $\lfloor \cdot \rfloor$ rounds down to the nearest integer. As this function is weakly decreasing in Δ , the maximum sparsity is increasing in Δ . So, we obtain the intuition that if the largest elements of $\boldsymbol{\eta}$ are more similar, then less sparsity can be produced by ℓ_1 -regularization.

Equal and Noisy. Let $\boldsymbol{\eta} = \boldsymbol{\beta}^* + \sigma \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon}$ has i.i.d. elements $\varepsilon_i \sim N(0, 1)$, $\sigma > 0$, and $\boldsymbol{\beta}^*$ is an m -vector with elements $\beta_i^* = \beta_j^*$ for all i, j . As all elements of $\boldsymbol{\beta}^*$ are equal, the gaps between the elements of $\boldsymbol{\eta}$ are equal to the gaps between the order statistics of $\boldsymbol{\varepsilon}$, scaled by the constant σ . So, the size of the gaps between the largest elements of $\boldsymbol{\eta}$ is increasing in σ . Therefore, according to (2.11), the maximum sparsity is increasing in σ . As an increase in σ represents an increase in noise, we can draw the intuitive conclusion that if $\boldsymbol{\beta}^*$ has elements of similar size, then the maximum amount of sparsity produced by ℓ_1 -regularization increases with noise.

2.4 General Case

In this section, we describe how a solution can be found for the general case, in which $\mathbf{X}$ is not required to be orthogonal. To do so, we adapt the framework laid out by Bertsimas et al. (2016) for standard linear regression. This

framework consists of two components. The first component is a Discrete First-Order (DFO) algorithm that uses an algorithm for the orthogonal problem as a subroutine in each iteration. The solution to this DFO algorithm is then used as an initial solution for the second component. The second component relies on reformulating (2.1) as an MIO problem, which can be solved to provable optimality by using an MIO solver.

2.4.1 Discrete First-Order Algorithm

In the construction of the DFO algorithm, we closely follow Bertsimas et al. (2016), but use a different constraint set that includes an additional ℓ_1 -norm restriction and unit-sum restriction.

Denote the objective function as

$$f(\boldsymbol{\beta}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2.$$

This function is Lipschitz continuously differentiable, as

$$\begin{aligned} \|\nabla f(\boldsymbol{\beta}) - \nabla f(\boldsymbol{\eta})\|_2^2 &= \|\mathbf{X}'\mathbf{X}(\boldsymbol{\beta} - \boldsymbol{\eta})\|_2^2 \\ &\leq \|\mathbf{X}'\mathbf{X}\|^2 \|\boldsymbol{\beta} - \boldsymbol{\eta}\|_2^2 \\ &= L^* \|\boldsymbol{\beta} - \boldsymbol{\eta}\|_2^2, \end{aligned}$$

where L^* is the largest absolute eigenvalue of $\mathbf{X}'\mathbf{X}$. So, we can apply the following result.

Proposition 6 (Nesterov, 2013; Bertsimas et al., 2016). *For a convex Lipschitz continuous function $f(\cdot)$, we have*

$$f(\boldsymbol{\eta}) \leq Q_L(\boldsymbol{\eta}, \boldsymbol{\beta}) := f(\boldsymbol{\beta}) + \frac{L}{2} \|\boldsymbol{\eta} - \boldsymbol{\beta}\|_2^2 + \nabla f(\boldsymbol{\beta})'(\boldsymbol{\eta} - \boldsymbol{\beta}), \quad (2.12)$$

for all $L \geq \bar{L}$, $\boldsymbol{\beta}$ and $\boldsymbol{\eta}$, where $\bar{L}$ is smallest constant such that $\|\nabla f(\boldsymbol{\beta}) - \nabla f(\boldsymbol{\eta})\|_2^2 \leq \bar{L} \|\boldsymbol{\beta} - \boldsymbol{\eta}\|_2^2$.

Given some fixed $\boldsymbol{\beta}$, we can minimize the bound in (2.12) with respect

to $\boldsymbol{\eta}$ under the constraint set $\mathcal{T}$, as given in (2.2). Following Bertsimas et al. (2016), we find

$$\begin{aligned}
 & \arg \min_{\boldsymbol{\eta} \in \mathcal{T}} Q_L(\boldsymbol{\eta}, \boldsymbol{\beta}) \\
 &= \arg \min_{\boldsymbol{\eta} \in \mathcal{T}} \left(f(\boldsymbol{\beta}) + \frac{L}{2} \|\boldsymbol{\eta} - \boldsymbol{\beta}\|_2^2 + \nabla f(\boldsymbol{\beta})'(\boldsymbol{\eta} - \boldsymbol{\beta}) + \frac{1}{2L} \|\nabla f(\boldsymbol{\beta})\|_2^2 - \frac{1}{2L} \|\nabla f(\boldsymbol{\beta})\|_2^2 \right) \\
 &= \arg \min_{\boldsymbol{\eta} \in \mathcal{T}} \left(f(\boldsymbol{\beta}) + \frac{L}{2} \|\boldsymbol{\eta} - (\boldsymbol{\beta} - \frac{1}{L} \nabla f(\boldsymbol{\beta}))\|_2^2 - \frac{1}{2L} \|\nabla f(\boldsymbol{\beta})\|_2^2 \right) \\
 &= \arg \min_{\boldsymbol{\eta} \in \mathcal{T}} \|\boldsymbol{\eta} - (\boldsymbol{\beta} - \frac{1}{L} \nabla f(\boldsymbol{\beta}))\|_2^2. \tag{2.13}
 \end{aligned}$$

Notice that (2.13) can be computed using Algorithm 1. Therefore, it is possible to use iterative updates in order to decrease the objective value. Specifically, let $\boldsymbol{\beta}_1 \in \mathcal{T}$ and recursively define $\boldsymbol{\beta}_{r+1} = \arg \min_{\boldsymbol{\eta} \in \mathcal{T}} Q_{L^*}(\boldsymbol{\eta}, \boldsymbol{\beta}_r)$, for all $r \in \mathbb{N}$. Then by Proposition 6,

$$f(\boldsymbol{\beta}_r) = Q_{L^*}(\boldsymbol{\beta}_r, \boldsymbol{\beta}_r) \geq Q_{L^*}(\boldsymbol{\beta}_{r+1}, \boldsymbol{\beta}_r) \geq f(\boldsymbol{\beta}_{r+1}).$$

In Algorithm 2, we present an algorithm that uses this updating step until some convergence criterion is reached.

Algorithm 2: First order algorithm

Input: Lipschitz constant L^* , convergence criterion ε , initial solution

$\boldsymbol{\beta}_1 \in \mathcal{T}$.

Output: $\hat{\boldsymbol{\beta}}$

1 $r = 1$

2 **repeat**

3 $\boldsymbol{\beta}_{r+1} \in \arg \min_{\boldsymbol{\eta} \in \mathcal{T}} \|\boldsymbol{\eta} - (\boldsymbol{\beta}_r - \frac{1}{L^*} \nabla f(\boldsymbol{\beta}_r))\|_2^2$ (using Algorithm 1).

4 $r = r + 1$

5 **until** $f(\boldsymbol{\beta}_r) - f(\boldsymbol{\beta}_{r-1}) < \varepsilon$;

6 $\hat{\boldsymbol{\beta}} = \boldsymbol{\beta}_r$

2.4.2 Mixed-Integer Optimization

In this section, an MIO formulation for problem (2.1) is presented. In order to formulate problem (2.1) as an MIO problem, we use three sets of auxiliary variables. The variables β_i^+ and β_i^- are used to specify the positive and negative parts of the arguments $\beta_i, i \in \{1, \dots, m\}$. The variable z_i serves as an indicator function for whether β_i is different from zero, $i \in \{1, \dots, m\}$. The MIO formulation is given as follows:

$$\begin{aligned}
 & \min_{\beta, z} \beta' X' X \beta - 2y' X \beta + \lambda \beta' \beta, \\
 & \text{s.t.} \\
 & \beta_i = \beta_i^+ - \beta_i^-, \quad i \in \{1, \dots, m\}, \\
 & \sum_{i=1}^m \beta_i = 1, \\
 & \sum_{i=1}^m \beta_i^+ \leq 1 + s, \\
 & \sum_{i=1}^m \beta_i^- \leq s, \\
 & \mathcal{M}_- z_i \leq \beta_i \leq \mathcal{M}_+ z_i, \quad i \in \{1, \dots, m\}, \\
 & \sum_{i=1}^m z_i \leq k, \\
 & \beta_i^+, \beta_i^- \geq 0, \quad i \in \{1, \dots, m\}, \\
 & z_i \in \{0, 1\}, \quad i \in \{1, \dots, m\},
 \end{aligned}$$

where β has elements β_i , and $\mathcal{M}_+$ and $\mathcal{M}_-$ are big-M parameters. These big-M parameters are used to enforce the sparsity constraint as follows: if $z_i = 1$ then $\beta_i \in [\mathcal{M}_-, \mathcal{M}_+]$, and if $z_i = 0$ then $\beta_i = 0$. Hence, $\mathcal{M}_-$ and $\mathcal{M}_+$ should be sufficiently large in absolute value to ensure that the solution to the MIO problem is the solution to (2.1). On the other hand, they should not be too large as tighter bounds decrease the size of the search space and improve the speed of the solver.

The ℓ_1 -restriction provides natural choices $\mathcal{M}_+ = 1 + s$ and $\mathcal{M}_- = -s$.

However, these bounds are conservative in practice. Mazumder et al. (2017) suggest the use of bounds based on the solution of the DFO algorithm. Similarly, we propose to use $\mathcal{M}_- = \max\{\frac{3}{2} \min(\min_i[\beta_i^{\text{DFO}}], -\frac{s}{k}), -s\}$ and $\mathcal{M}_+ = \min\{\frac{3}{2} \max(\max_i[\beta_i^{\text{DFO}}], \frac{1+s}{k}), 1+s\}$, where β_i^{DFO} is the i th element of the solution of the DFO algorithm.

2.5 Numerical Results

In this section we compare the performance of our $\ell_0\ell_1$ -regularized approach to ℓ_0 -regularization and ℓ_1 -regularization on simulated datasets, generated with multiple signal-to-noise ratios and values of β .

2.5.1 Setup Simulation Experiments

The setup of our simulation experiments largely follows the numerical experiments found in Mazumder et al. (2017) and Hastie et al. (2017). For a given set of parameters t (number of observations), m (number of variables), k^* (number of non-zero weights), p (number of positive weights), n (number of negative weights), s^* (sum of the negative weights), ρ (autocorrelation between the variables) and SNR (signal-to-noise ratio), the experiments are conducted as follows:

1. We randomly select k^* elements of β and set p of the elements equal to $(1 + s^*)/p$, and n of the elements equal to $-s^*/n$. The remaining elements are set equal to zero.
2. The rows of $t \times m$ matrix X are drawn i.i.d. from $N_m(\mathbf{0}, \Sigma)$, where Σ has elements $\sigma_{ij} = \rho^{|i-j|}$, $i, j \in \{1, \dots, p\}$.
3. The vector $\mathbf{y}$ is drawn from $N(X\beta, \sigma^2 I)$, where $\sigma^2 = \beta' \Sigma \beta / \text{SNR}$ in order to fix the signal-to-noise ratio.
4. We apply ℓ_0 -regularization, ℓ_1 -regularization and $\ell_0\ell_1$ -regularization to X and $\mathbf{y}$ for a range of tuning parameters. For both methods, we

select the tuning parameter(s) that minimize(s) the prediction error on a separate dataset $\tilde{X}, \tilde{y}$, generated in the same way as X and y .

5. We record several performance measures of the solutions that were found using the selected tuning parameters.

We repeat the above steps 100 times for each parameter setting. Throughout the experiments we use $t = 50$, $m = 100$, $k^* = 7$, $\rho = 0.2$. For each setting, we choose $s^* \in \{0.1, 2/3\}$ and $\text{SNR} \in \{2^{-1}, 2^0, 2^1\}$. This choice of s^* covers the case where the negative weights are small in comparison to the positive weights, as well as the case where the positive and negative weights are equal in magnitude. The tuning parameters corresponding to k^* and s^* are simultaneously selected over the grid $\{1, \dots, 20\} \times \{0, s^*/5, \dots, 2s^*\}$.

For each different combination of s^* and SNR, we record the following performance measures:

- **Relative risk.** As measure of predictive performance we use relative risk, defined for a solution $\hat{\beta}$ as

$$\text{RR}(\beta) = \frac{(\hat{\beta} - \beta)' \Sigma (\hat{\beta} - \beta)}{\beta' \Sigma \beta}.$$

This is one of the measures used in Hastie et al. (2017), and is similar to the predictive performance measures used in Bertsimas et al. (2016) and Mazumder et al. (2017). For this measure, a lower value is indicative of a better predictive performance and its minimum value is 0. The null score to beat is 1 (if $\hat{\beta} = 0$).

- **Number of non-zero elements.** As a second measure, we consider the number of non-zero elements in the estimated weights, in order to compare the sparsity obtained by both methods.
- **Sum of negative weights.** As a final measure, we consider the sum of the negative estimated weights. This allows us to compare the shrinkage produced by the ℓ_1 -regularization component of both methods.

2.5.2 Implementation and Stopping Criteria

In order to compute the solution to ℓ_1 -regularized unit-sum regression, we use an adaptation of the LARS method for ℓ_1 -regularization (Efron et al., 2004), based on the algorithm described by DeMiguel et al. (2009). The ℓ_0 -regularization solution is computed in the same way as the $\ell_0\ell_1$ -regularization solution by fixing the parameter s to some sufficiently large value.

For the $\ell_0\ell_1$ -regularization approach, we terminate the DFO algorithm if the improvement in the squared error is below some value ε , where we set $\varepsilon = 10^{-6}$. As the DFO algorithm can be sensitive to its initialization, we initialize it with the Forward-Stepwise Selection (FSS). We found that this typically yields a better performance than using the best solution out of 50 random initializations as used by Bertsimas et al. (2016). The FSS solution is implemented using successive applications of the adapted LARS algorithm.

The MIO formulation is implemented in the R-interface of Gurobi 7.1. Each instance is given 10 minutes of computation time. If the optimality of the solution is not confirmed within the allotted time, the solver is terminated and its best solution so far is used. This means that the combined maximum computation time is 44000 hours. However, in practice we find that the DFO algorithm often provides optimal or near-optimal solutions to the MIO solver. As a result, the MIO solver rarely uses the full 10 minutes and typically certifies optimality in seconds. The total computation time for the simulation experiments was approximately 600 hours on a standard university personal computer, including the computation of the initial solutions.

2.5.3 Results of Simulation Experiments

The results of the simulation experiments are displayed in Figure 2.1. We make the following observations.

Prediction. It can be observed that ℓ_0 -regularization typically performs worse than the other methods, especially when the SNR is low. Further-

more, $\ell_0\ell_1$ -regularization seems to outperform ℓ_1 -regularization for higher SNRs in terms of relative risk, while ℓ_1 -regularization fares similarly or even somewhat better for lower SNRs.¹

Sparsity. We find that ℓ_1 -regularization delivers considerably denser solutions than the other methods for all values of SNR and s^* . In addition, the number of non-zeros seems to move away from the true number of non-zeros as the SNR increases. On the other hand, ℓ_0 -regularization yields overly sparse solutions below the true value k^* , especially if the SNR is low. The number of non-zeros produced by $\ell_0\ell_1$ -regularization lies between the values other two methods, and is typically closer k^* than the number of non-zeros produced by ℓ_0 -regularization or ℓ_1 -regularization.

Shrinkage. In the third column of Figure 2.1, it can be seen that the sum of the negative weights of the solutions tends to be smaller than s^* . However, for the case that $s^* = 2/3$, there is a clear trend towards the true value of s^* as the SNR increases. Interestingly, both $\ell_0\ell_1$ -regularization and ℓ_1 -regularization have a similar sum of negative weights, despite the fact that the solutions of $\ell_0\ell_1$ -regularization are much sparser. This implies that the average magnitude of the weights of the ℓ_1 -regularization solution is much smaller than that of the $\ell_0\ell_1$ -regularization solution.

2.6 Application: Index Tracking

In order to demonstrate the use of our proposed methodology in practice, we consider an application to index tracking. Index tracking concerns the construction of a portfolio that replicates a stock index as closely as possible, while limiting the cost of holding the portfolio. Such a portfolio can be represented by a weight vector that sums to one, with positive elements

¹ These results differ slightly from the findings by Mazumder et al. (2017) for standard linear regression. They find that $\ell_0\ell_1$ -regularization performs as well as ℓ_1 -regularization if the SNR is low. We suspect that this difference could be caused by the fact that they do not consider an SNR below 1 and use a fixed-design setup where $\tilde{X} = X$.

that correspond to long positions and negative elements that correspond to short positions.

Two standard ways to limit the cost of holding the portfolio are to restrict the number of constituents in the portfolio and to avoid short positions. Using historical returns data, it is possible to find such a portfolio using $\ell_0\ell_1$ -regularized unit-sum regression. Specifically, let $\mathbf{y}$ represent the historical returns of a stock index and let each column of $\mathbf{X}$ represent the historical returns of one of its constituents. Then, using $s = 0$, problem (2.1) minimizes the squared error between the actual index returns and the returns of the portfolio, that consists of at most k constituents and has no short positions.

Notice that even if the ℓ_0 component is omitted, or equivalently $k = m$, then the remaining ℓ_1 -regularization may still produce a sparse portfolio (DeMiguel et al., 2009; Brodie et al., 2009). However, as the returns of an index are typically a dense linear combination of its constituent returns, with positive weights of similar magnitude, the intuitions from the orthogonal design case from Section 2.3 suggest that ℓ_1 -regularization may not be very effective in producing sparsity.

To compare the sparsity production and tracking performance of $\ell_0\ell_1$ -regularization and ℓ_1 -regularization, we use the index tracking datasets of the OR-library (Beasley et al., 2003; Canakgoz and Beasley, 2009). These datasets contain 290 weekly returns of 8 indexes varying from 31 to 2153 constituents.² Each dataset is split into two halves of 145 observations, where the first half is used to construct the portfolio and the second half is used to measure the performance of the portfolio.³ The performance is measured in out-of-sample R^2 , on the second half of the datasets. The results are presented in Table 2.1.

From the results we can make several observations regarding the sparsity of the solutions and the tracking performance. First it should be noted that

²Only the constituents that are part of the index for the entire period are included.

³From the second index (DAX) we removed two large consecutive outliers from the out-of-sample data. These two outliers were the largest two returns (in absolute value) and of opposing sign, suggesting a bookkeeping error in the index returns. This is supported by the fact that the outliers are not reflected in the returns of the constituents.

ℓ_1 -regularization by itself is not able to find a unique portfolio for the largest two stock indexes. In addition, even if ℓ_1 -regularization does have a unique solution, it is generally not able to produce a substantial amount of sparsity. In terms of out-of-sample tracking performance, lower values of k do generally result in worse performance. However, the difference is small, especially for the larger values of k .

2.7 Appendix

Proof of Proposition 2: If $\beta^{(\mathcal{P}, \mathcal{N}, z)} \in \arg \min_{\beta \in \mathcal{A}_z} Q(\beta)$, then $\mathcal{P} = \{1, \dots, p\}$ and $\mathcal{N} = \{m - n + 1, \dots, m\}$ if $n \geq 1$.

Proof. We show that if $\beta^{(\mathcal{P}, \mathcal{N}, z)} \in \arg \min_{\beta \in \mathcal{A}_z} Q(\beta)$, then two conditions hold true:

$$\mathcal{P}\text{-condition : } \max(\mathcal{P}) < \min\{(\mathcal{P} \cup \mathcal{N})^c\},$$

$$\mathcal{N}\text{-condition : } \min(\mathcal{N}) > \max\{(\mathcal{P} \cup \mathcal{N})^c\} \text{ if } n \geq 1.$$

We prove the $\mathcal{P}$ -condition. The proof of the $\mathcal{N}$ -condition is similar. Assume the $\mathcal{P}$ -condition is not true. In that case, we show that an index set $\overline{\mathcal{P}}$ exists, such that $\beta^{(\overline{\mathcal{P}}, \mathcal{N}, z)} \in \mathcal{A}_z$ and $Q(\beta^{(\overline{\mathcal{P}}, \mathcal{N}, z)}) < Q(\beta^{(\mathcal{P}, \mathcal{N}, z)})$, which is a contradiction, showing the validity of the $\mathcal{P}$ -condition.

Assuming the $\mathcal{P}$ -condition is not true, let $u := \max(\mathcal{P}) > v := \min\{(\mathcal{P} \cup \mathcal{N})^c\}$. As $\beta^{(\mathcal{P}, \mathcal{N}, z)} \in \mathcal{A}_z$, we have $\beta_u^{(\mathcal{P}, \mathcal{N}, z)} > 0$, which is equivalent to $z + 1 > \sum_{i \in \mathcal{P}} (\eta_i - \eta_u)$. Define

$$\mathcal{P}^* := \mathcal{P} \setminus u,$$

$$\tilde{\mathcal{P}}_j := \{i \mid i \in \mathcal{P}^*, i \leq j \in \mathcal{P}^*\},$$

$$\overline{\mathcal{P}}_j := \tilde{\mathcal{P}}_j \cup v.$$

Let $j_v := \max\{j \mid j \in \mathcal{P}^*, j < v\}$. As $\eta_i - \eta_u > 0$ if $i \in \mathcal{P}^*$, we find

$$z + 1 > \sum_{i \in \mathcal{P}} (\eta_i - \eta_u) = \sum_{i \in \mathcal{P}^*} (\eta_i - \eta_u) \geq \sum_{i \in \tilde{\mathcal{P}}_{j_v}} (\eta_i - \eta_u) > \sum_{i \in \tilde{\mathcal{P}}_{j_v}} (\eta_i - \eta_v) = \sum_{i \in \overline{\mathcal{P}}_{j_v}} (\eta_i - \eta_v).$$

Consequently, $\mathbf{x}^{(\overline{\mathcal{P}}_{j_v}, \mathcal{N}, z)} \in \mathcal{A}_z$. Therefore, the index set $\overline{\mathcal{P}}_{\bar{j}}$, where $\bar{j}$ is the maximum index such that $\mathbf{x}^{(\overline{\mathcal{P}}_{\bar{j}}, \mathcal{N}, z)} \in \mathcal{A}_z$, exists. Let $\overline{\mathcal{P}} = \overline{\mathcal{P}}_{\bar{j}}$.

We now show $Q(\beta^{(\overline{\mathcal{P}}, \mathcal{N}, z)}) < Q(\beta^{(\mathcal{P}, \mathcal{N}, z)})$. Let $\mathcal{R} = \mathcal{P}^* \setminus \tilde{\mathcal{P}}_{\bar{j}}$ with cardinality r . As $\mathcal{P}^* = \tilde{\mathcal{P}}_{\bar{j}} \cup \mathcal{R}$ has cardinality $p - 1$, the cardinality of $\overline{\mathcal{P}} =$

$\tilde{\mathcal{P}}_{\tilde{j}} \cup v$ equals $p - r$. Let

$$b := \frac{(\sum_{i \in \mathcal{P}} \eta_i) - 1 - z}{p}, \quad \text{and} \quad \bar{b} := \frac{(\sum_{i \in \bar{\mathcal{P}}} \eta_i) - 1 - z}{p - r}.$$

We find

$$\begin{aligned} Q(\beta^{(\bar{\mathcal{P}}, \mathcal{N}, z)}) - Q(\beta^{(\mathcal{P}, \mathcal{N}, z)}) &= (p - r)\bar{b}^2 - pb^2 + \eta_u^2 - \eta_v^2 + \sum_{i \in \mathcal{R}} \eta_i^2 \\ &= (p - r)(\bar{b} + b)(\bar{b} - b) + \eta_u^2 - \eta_v^2 + \sum_{i \in \mathcal{R}} (\eta_i^2 - b^2) \\ &= (\bar{b} + b) \left\{ \eta_v - \eta_u - \sum_{i \in \mathcal{R}} (\eta_i - b) \right\} + \eta_u^2 - \eta_v^2 + \sum_{i \in \mathcal{R}} (\eta_i^2 - b^2) \\ &= (\eta_u + \eta_v - \bar{b} - b)(\eta_u - \eta_v) + \sum_{i \in \mathcal{R}} (\eta_i - b)(\eta_i - \bar{b}) \\ &= \left(\beta_u^{(\mathcal{P}, \mathcal{N}, z)} + \beta_v^{(\bar{\mathcal{P}}, \mathcal{N}, z)} \right) (\eta_u - \eta_v) + \sum_{i \in \mathcal{R}} \beta_i^{(\mathcal{P}, \mathcal{N}, z)} (\eta_i - \bar{b}). \end{aligned}$$

As $\beta_u^{(\mathcal{P}, \mathcal{N}, z)} > 0$, $\beta_v^{(\bar{\mathcal{P}}, \mathcal{N}, z)} > 0$ and $\eta_u < \eta_v$, the first term is negative. As $\beta_i^{(\mathcal{P}, \mathcal{N}, z)} > 0$ and $\eta_i - \bar{b} < 0$ if $i \in \mathcal{R}$, and hence $i > \tilde{j}$, the second term is negative as well. Consequently, $Q(\beta^{(\bar{\mathcal{P}}, \mathcal{N}, z)}) < Q(\beta^{(\mathcal{P}, \mathcal{N}, z)})$, which is a contradiction. $\blacksquare$

Proof of Lemma 1: If $\beta^{(\tilde{p}, \tilde{n}, z)} \in \mathcal{A}_z$ and $\beta^{(p, n, z)} \in \mathcal{A}_z$, where $\tilde{p} \leq p$, $\tilde{n} \leq n$, $\tilde{p} + \tilde{n} < p + n$, then $Q(\beta^{(p, n, z)}) < Q(\beta^{(\tilde{p}, \tilde{n}, z)})$.

Proof. We can use the convexity of the quadratic function to show

$$\begin{aligned} Q(\beta^{(p, n, z)}) - Q(\beta^{(p+1, n, z)}) &= p \left(\frac{(\sum_{i=1}^p \eta_i) - 1 - z}{p} \right)^2 + \eta_{p+1}^2 - (p+1) \left(\frac{(\sum_{i=1}^{p+1} \eta_i) - 1 - z}{p+1} \right)^2 \\ &= (p+1) [\lambda y_1^2 + (1-\lambda)y_2^2 - \{\lambda y_1 + (1-\lambda)y_2\}^2] > 0, \end{aligned}$$

where $\lambda = p/(p+1)$, $y_1 = \{(\sum_{i=1}^p \eta_i) - 1 - z\} / p$ and $y_2 = \eta_{p+1}$. In a similar way we find $Q(\beta^{(p, n, z)}) > Q(\beta^{(p, n+1, z)})$. $\blacksquare$

Proof of Proposition 3: If $\beta^{(\mathcal{P}, \mathcal{N}, z)} \in \arg \min_{\beta \in \mathcal{A}_z} Q(\beta)$ and $p_z + n_z \leq k$, then $\beta^{(\mathcal{P}, \mathcal{N}, z)} = \beta^{(p_z, n_z, z)}$.

Proof of Proposition 5: Proposition 2 holds true when $k = m$.

Proof. Using Proposition 1, let $\beta^{(\mathcal{P}, \mathcal{N}, z)} \in \arg \min_{\beta \in \mathcal{A}_z} Q(\beta)$ where $\mathcal{P}$ and $\mathcal{N}$ have cardinalities p and n , respectively, where $p + n = m$. We will show that $u := \max(\mathcal{P}) < v := \min(\mathcal{N})$ if $n \geq 1$.

Suppose $u := \max(\mathcal{P}) > v := \min(\mathcal{N})$, then

$$\frac{(\sum_{i \in \mathcal{P}} \eta_i) - 1 - z}{p} < \eta_u < \eta_v < \frac{(\sum_{i \in \mathcal{N}} \eta_i) + z}{n}. \quad (\text{A.1})$$

Consequently,

$$Q\left(\beta^{(\mathcal{P}, \mathcal{N}, z)}\right) = p \left\{ \frac{(\sum_{i \in \mathcal{P}} \eta_i) - 1 - z}{p} \right\}^2 + n \left\{ \frac{(\sum_{i \in \mathcal{N}} \eta_{m-i+1}) + z}{n} \right\}^2$$

has a positive derivative

$$\begin{aligned} \frac{dQ\left(\beta^{(\mathcal{P}, \mathcal{N}, z)}\right)}{dz} &= -2 \left\{ \frac{(\sum_{i \in \mathcal{P}} \eta_i) - 1 - z}{p} \right\} + 2 \left\{ \frac{(\sum_{i \in \mathcal{N}} \eta_{m-i+1}) + z}{n} \right\} \\ &> 2(\eta_v - \eta_u) > 0. \end{aligned}$$

As a result, $\{Q\left(\beta^{(\mathcal{P}, \mathcal{N}, z)}\right) \mid z \text{ satisfies (A.1)}\}$ does not have a minimum. If

$$z^* = \min \left\{ z \mid \frac{(\sum_{i \in \mathcal{P}} \eta_i) - 1 - z}{p} \leq \eta_u < \eta_v \leq \frac{(\sum_{i \in \mathcal{N}} \eta_i) + z}{n} \right\},$$

then $\beta^{(\mathcal{P}^*, \mathcal{N}^*, z^*)}$ contains zeros, so that $Q(\beta^{(\mathcal{P}^*, \mathcal{N}^*, z^*)}) \geq Q(\beta^{(p_{z^*}, n_{z^*}, z^*)}) > Q(\beta^{(p_z, n_z, z)})$. ■

Figure 2.1. Performance of ℓ_0 -regularization, ℓ_1 -regularization and $\ell_0\ell_1$ -regularization in terms of relative risk, number of non-zero elements and the sum of the negative weights as a function of the SNR for two values of s^* . The vertical bars represent one standard error.

Index	k	#nz	R^2_{os}	Index	k	#nz	R^2_{os}
	5	5	0.909		20	20	0.922
Hang	15	15	0.982	Nikkei	60	60	0.957
Seng	25	25	0.991	(m = 225)	100	100	0.961
(m = 31)	31	25	0.991		225	127	0.961
	10	10	0.940		20	20	0.780
DAX	30	30	0.979	S&P	60	60	0.839
(m = 85)	50	50	0.981	500	100	100	0.857
	85	78	0.985	(m = 457)	457	122	0.855
	10	10	0.652		20	20	0.646
FTSE	30	30	0.948	Russel	60	60	0.679
(m = 89)	50	50	0.959	2000	100	100	0.691
	89	68	0.966	(m = 1319)	1319	-	-
	10	10	0.815		20	20	0.767
S&P	30	30	0.932	Russel	60	60	0.821
100	50	50	0.960	3000	100	100	0.836
(m = 98)	98	77	0.969	(m = 2152)	2152	-	-

Table 2.1. Out-of-sample R^2 (R^2_{os}) and number of non-zeros (#nz) of $\ell_0\ell_1$ -regularized unit-sum regression for all 8 index datasets and multiple values for the parameters k , using $s = 0$. The results for $k = m$ are equivalent to ℓ_1 -regularized unit-sum regression with $s = 0$. A hyphen (-) indicates that no unique solution was found.

Chapter 3

Exact Testing of Many Moment Inequalities against Multiple Violations

3.1 Introduction

As discussed by Chernozhukov et al. (2019a), 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*

This chapter is based on Koning and Bekker (2019).

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 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

¹ For 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).

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 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

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

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.

3.2 The model and two test statistics

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

$$X = \iota_n \mu' + E, \quad (\text{A.1})$$

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

Following CCK, we are interested in testing the hypothesis $H_0: \mu \leq 0$ against the alternative $H_1: \mu \not\leq 0$. That is, we want to test whether all elements of μ 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 $I_n = (e_1, \dots, e_n)$ be the $n \times n$ identity matrix, and define $P_{\iota_n} = \iota_n \iota_n' / \iota_n' \iota_n$. Define $\hat{\mu} = X' \iota_n / n$ with elements $\hat{\mu}_j$, and $\hat{\Sigma} = X'(I_n - P_{\iota_n})X / n$ with diagonal elements $\hat{\sigma}_j^2$. We assume $\hat{\mu} \not\leq 0$, otherwise we would not reject H_0 , and consider test statistics of the form

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

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

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

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, $t_j = n^{1/2} \widehat{\mu}_j / \widehat{\sigma}_j$, for $j = 1, \dots, p$. We will denote this statistic by $t_{\max} = t_{\max}(\mathbf{X}) = \max_{1 \leq j \leq p} t_j = T_{C^p}$, where $C^p = \{e_1, \dots, e_p\}$. 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 $\mathcal{U} = C^p \cup \{\mathbf{a}_1, \dots, \mathbf{a}_m\}$, 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 $\widehat{\Sigma} = \mathbf{1}_p \mathbf{1}_p'$, the $T_{\mathcal{U}}$ and $t_{\max}$ statistics coincide if $\mathcal{U} \subset \mathbb{R}_+^p$ is finite and $C^p \subset \mathcal{U}$. That is to say, $T_{\mathcal{U}} \geq t_{\max}$ by construction, and

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

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 $T_{\mathbb{R}^p}^2 = T^2 = n \widehat{\mu}' \widehat{\Sigma}^{-1} \widehat{\mu}$, which requires $\widehat{\Sigma}$ to be invertible. For the T_+ statistic $\widehat{\Sigma}$ need not be invert-

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 $\min_{\mu \leq 0} n(\widehat{\mu} - \mu)' \widehat{\Sigma}^{-1} (\widehat{\mu} - \mu)$ 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.

ible. Instead, a weaker condition suffices, which is a special case of Condition 1.

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

This condition is weaker than (strict) copositivity of $\hat{\Sigma}$, which requires that $\lambda' \hat{\Sigma} \lambda > 0$ for all $\lambda \geq 0$, $\lambda \neq 0$ (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 $\hat{\Sigma}$ are strictly positive (i.e. $\hat{\Sigma}$ is *positive*). Meinshausen et al. (2013) also provides an example where $\hat{\Sigma}$ is allowed to contain some negative entries. In particular, if $\mathcal{A}^c$ contains the indices of the the largest positive principal submatrix of $\hat{\Sigma}$ and $\mathcal{A}$ is its complement, then the condition only has to hold on the principal submatrix of $\hat{\Sigma}$ 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 $\hat{\mu}$. Condition 2 is equivalent to copositivity if $\hat{\mu} > 0$. At the other extreme, if $\hat{\mu} \leq 0$ the set $\{\lambda \mid \lambda \geq 0, \lambda \neq 0, \mu' \lambda \geq 0\}$ is empty, so that Condition 2 holds, whether or not copositivity holds. If the elements of $\hat{\mu}$ have comparable positive or negative sizes, then Condition 2 tends to become less restrictive the more negative elements there are. So, intuitively, even if $\hat{\Sigma}$ 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

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

This problem can be solved by the non-negative least squares algorithm of

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

Lawson and Hanson (1995).⁵ An interesting additional feature is that the resulting maximizing weight vector $\hat{\lambda}$ may contain zeros, resulting in a selection of moment inequalities that are ‘suspected’ to violate the null.

3.3 Symmetry based inference

The randomization tests that we consider are based on the symmetry assumption. For completeness, we include this section to 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 $\mathcal{R}_{E^*} = \{E^*, \mathbf{R}_2 E^*, \dots, \mathbf{R}_N E^*\}$, where $E^* \in \mathbb{R}^{n \times p}$ acts as a representative of the class. For simplicity, we only consider orbits $\mathcal{R}_{E^*} \in \mathcal{E}$, for which the cardinality satisfies $|\mathcal{R}_{E^*}| = |\mathcal{R}| = N$. This is equivalent to assuming that we only consider the subset of orbits $\mathcal{E}^0 \subset \mathcal{E}$ where the elements of $\mathcal{R}_{E^*} \in \mathcal{E}^0$ 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 $\mathbf{R}E$ and E have the same distribution, for all $\mathbf{R} \in \mathcal{R}$. Equivalently, we assume the conditional distribution of E , given an orbit $E \in \mathcal{R}_{E^*}$, is uniform for all orbits $\mathcal{R}_{E^*} \in \mathcal{E}^0$.

Let $g : \mathcal{R}_{E^*} \rightarrow \mathbb{R} \cup \{\infty\}$ be a mapping. Notice that if g is injective on all orbits $\mathcal{R}_{E^*} \in \mathcal{E}^0$, then the symmetry assumption implies that the conditional distribution of $g(E)$, given an orbit $\mathcal{R}_{E^*} \in \mathcal{E}^0$, is uniform just as well. Furthermore, the function $p(E) = |\{\mathbf{R} \in \mathcal{R} \mid g(\mathbf{R}E) \geq g(E)\}|/N$ has a con-

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

ditional distribution, given an orbit, that is uniform over $\{1/N, 2/N, \dots, 1\}$. As this holds for all $\mathcal{R}_{E^*} \in \mathcal{E}^0$, the function $p(E)$ has an unconditional uniform distribution over $\{1/N, 2/N, \dots, 1\}$. This leads to the following results.

Proposition 7. $P(p(E) \leq \alpha) \leq \alpha$, for $\alpha \in [0, 1]$. If g is injective on all $\mathcal{R}_{E^*} \in \mathcal{E}^0$, then $P(p(E) \leq \alpha) = \alpha$, 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 $p(E)$ is intractable even in small samples. Therefore, we provide the following result to allow for sampling from $\mathcal{R}$.⁶

Proposition 8. 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 $p_M(E) = |\{\mathbf{R} \in \mathcal{R}^M \mid g(\mathbf{R}E) \geq g(E)\}|/M$, then $P(p_M(E) \leq \alpha) \leq \alpha$, for $\alpha \in [0, 1]$ and, if g is injective, then $P(p_M(E) \leq \alpha) = \alpha$ for $\alpha \in \{1/M, 2/M, \dots, 1\}$.

3.4 Randomization tests for moment inequalities

To test the moment inequalities of model (A.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 $\hat{\mu} \leq \mathbf{0}$ 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 $g(E) = T(E)$. So, Proposition 8 can be used based on $p_M(E) = p_M(\mathbf{X}) = |\{\mathbf{R} \in \mathcal{R}^M \mid T(\mathbf{R}(\mathbf{X})) \geq T(\mathbf{X})\}|/M$. 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.

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

2. Generate the transformed dataset $\mathbf{RX}$ and compute $T(\mathbf{RX})$, 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 $T(\mathbf{RX}) \geq T(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{RE})$, then the test is conservative.⁷ In that case we find, that $T(\mathbf{X}) = T(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{E})$ is an inferential function, $g(\mathbf{E}) = T(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{E})$, so that

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

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.

First, note that in the definition of $T_{\mathcal{U}}$, the statistic is invariant to the scale of $\boldsymbol{\lambda} \in \mathcal{U} \subseteq \mathbb{R}_+^p$. Therefore, we can rescale the elements of $\mathcal{U}$ to have length 1 without loss of generality, so that the resulting $\mathcal{U}$ is bounded. For the remainder, we make the further assumption that $\mathcal{U}$ is closed, so that it is compact.

Condition 3. *As $n \rightarrow \infty$ and p is fixed, $\boldsymbol{\lambda}' \mathbf{E}' \boldsymbol{\iota}_n / n$ and $\boldsymbol{\lambda}' (\mathbf{E}' \mathbf{E} / n - \boldsymbol{\Sigma}) \boldsymbol{\lambda}$ converge almost surely to 0, uniformly in $\boldsymbol{\lambda} \in \mathcal{U}$, where $\boldsymbol{\Sigma} > \mathbf{0}$ and $\mathcal{U} \subseteq \mathbb{R}_+^p$ is closed. That is, $\max_{\boldsymbol{\lambda} \in \mathcal{U}} |\boldsymbol{\lambda}' \mathbf{E}' \boldsymbol{\iota}_n / n| \xrightarrow{a.s.} 0$ and $\max_{\boldsymbol{\lambda} \in \mathcal{U}} |\boldsymbol{\lambda}' (\mathbf{E}' \mathbf{E} / n - \boldsymbol{\Sigma}) \boldsymbol{\lambda}| \xrightarrow{a.s.} 0$.*

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

To verify $T(\mathbf{RX}) \geq T(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{RE})$, 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 $\boldsymbol{\lambda} \geq 0$. Furthermore, notice that

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

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 $T^*(\mathbf{RX})$,

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

as $\boldsymbol{\iota}'_n \mathbf{R} \boldsymbol{\iota}_n < n$. Due to Condition 3, as $\mathbf{E}$ and $\mathbf{RE}$ follow the same distribution,

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

uniformly in $\boldsymbol{\lambda}$. Furthermore, $T^*(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{RE})$ maximizes

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

as a function of $\boldsymbol{\lambda} \in \mathcal{U}$. So that

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

uniformly in $\boldsymbol{\lambda}$. Let $\boldsymbol{\lambda}^*$ be the maximizer of $T^*(\boldsymbol{\iota}_n \boldsymbol{\mu}' + \mathbf{RE})$, then due to the uniform convergence we have convergence of the maximum, so that

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

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

3.5 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 $\mu = \mathbf{0}$, 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 $\mu = \mathbf{0}$, then $T(\mathbf{X}_{\mathcal{J}(\mathbf{X})}) = T(\mathbf{E}_{\mathcal{J}(\mathbf{E})}) = g(\mathbf{E})$ is an inferential function and

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

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 8. 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 I_n and $M - 1$ matrices drawn from $\mathcal{R} \setminus \{I_n\}$ without replacement.
2. Compute the selected inequalities $\mathcal{J}(RX)$, for each $R \in \mathcal{R}_M$.
3. Compute $T(RX_{\mathcal{J}(RX)})$ for all $R \in \mathcal{R}_M$.
4. Compute $p_{\text{sel}}^T(X)$, by counting the proportion that exceeds $T(X_{\mathcal{J}(X)})$.
5. Reject the null if $p_{\text{sel}}^T(X) \leq \alpha$.

If $\mu \leq \mathbf{0}$ and $\mu \neq \mathbf{0}$, the test is conservative if

$$T(RX_{\mathcal{J}(RX)}) \geq T\left((\iota\mu' + RE)_{\mathcal{J}(\iota\mu' + RE)}\right).$$

The power of the test varies with the value of μ and the selection method. We follow CCK and use $\mathcal{J}(X) = \{j \in \{1, \dots, p\} \mid t_j > c\}$, where the selection constant is chosen using their empirical bootstrap procedure. The simulations show that the size of the tests is controlled when $\mu \leq \mathbf{0}$.

3.6 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 $X = \iota_n\mu' + EA$, where ι_n is an n -vector of ones, $n \times p$ matrix E has i.i.d. elements drawn from a distribution F , and A is defined such that $\Sigma = A'A$, where Σ has elements $\sigma_{ij} = \rho^{|i-j|}$.

⁸ We 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.

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 μ , we consider four different designs. The signs of the elements of μ 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 μ 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 3.1. For the sake of brevity, the asymmetric error distributions were only 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 μ 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 $t_{\max}^t = T_{C_p \cup \{t_p\}}$ 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 $t_{\max}^t$

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

statistics.

For each setting, 1000 realizations of 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 $t'_{\max}$ statistics this is approximately 5-7 seconds.

3.6.1 Results

The results of the simulation experiments can be found in Tables 3.1 to 3.5, and are discussed separately for each design.

Design 1: $\mu = 0$.

In Design 1, all elements of μ 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 3.1 for the symmetric error distribution, and Table 3.5 for the asymmetric error distributions.

In Table 3.1, for the case that the number of observations is large ($n = 400$), we see that the EB tests for both $t_{\max}$ and $t'_{\max}$ reject with probability close to 0.05, which is in line with Remark 1. For the symmetry randomization (SR) tests, Proposition 8 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 $T_+(X)$ and critical values, defined as the $1 - \alpha$ quantile of the set $\{T_+(RX) \mid R \in \mathcal{R}_M\}$, 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 to Proposition 8. 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 3.5 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 3.2.

From Table 3.2, 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 3.4 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 3.3.

In Table 3.3, 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 $t'_{\max}$ and $t_{\max}$ statistics is quite remarkable, as the $t'_{\max}$ 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 $t'_{\max}$ statistic generally performs as well as or better than the $t_{\max}$ statistic, especially if $\rho = 0$.

Design 4: $\mu \not\leq 0, \mu \not\geq 0$.

In Design 4, for $n = 400$, the first $0.1p$ elements of μ are equal to 0.02 and the remaining elements are equal to -0.75 . For $n = 30$, the first 10 elements of μ 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 3.4.

The results in Table 3.4 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 $t'_{\max}$ 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 $t'_{\max}$ 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 ι_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 3.2 which suggest that the condition may still be satisfied if many of the moment inequalities are strictly satisfied. The $t'_{\max}$ statistic performs best overall for $n = 30$.

Figure 3.1. 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_+	
			EB	SR	EB sel	SR sel	EB	SR	EB sel	SR sel	SR	SR sel
200		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		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
200		0	.094	.055	-	-	.111	.061	-	-	.000	-
		0.5	.122	.056	-	-	.117	.048	-	-	.000	-
		0.9	.128	.050	-	-	.153	.045	-	-	.048	-
30	500	0	.114	.040	-	-	.124	.045	-	-	.000	-
		0.5	.142	.044	-	-	.141	.049	-	-	.000	-
		0.9	.167	.052	-	-	.173	.058	-	-	.000	-
1000		0	.137	.049	-	-	.155	.045	-	-	.000	-
		0.5	.170	.063	-	-	.174	.051	-	-	.000	-
		0.9	.210	.049	-	-	.204	.060	-	-	.000	-

Table 3.1. 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 $t_{\max}$, $t'_{\max}$ and T_+ statistics, both with pre-selection (sel) and without pre-selection.

n	p	ρ	$t_{\max}$				$t'_{\max}$				T_+	
			EB	SR	EB sel	SR sel	EB	SR	EB sel	SR sel	SR	SR sel
400	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
	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
30	200	0	.004	.034	-	-	.006	.044	-	-	.009	-
		.5	.006	.020	-	-	.013	.046	-	-	.012	-
		.9	.010	.023	-	-	.014	.033	-	-	.014	-
	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 3.2. Monte Carlo rejection probabilities with 1000 repetitions for Design 2: $\mu \leq 0$, 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 $t_{\max}$, $t'_{\max}$ 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	.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 3.3. Monte Carlo rejection probabilities with 1000 repetitions for Design 3: $\mu > \mathbf{0}$, 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 $t_{\max}$, $t_{\max}^t$ and T_+ statistics, both with pre-selection (sel) and without pre-selection.

n	p	ρ	$t_{\max}$				$t'_{\max}$				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 3.4. Monte Carlo rejection probabilities with 1000 repetitions for Design 4: $\mu \not\leq \mathbf{0}, \mu \not\geq \mathbf{0}$, 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 $t_{\max}$, $t'_{\max}$ and T_+ statistics, both with pre-selection (sel) and without pre-selection.

γ	p	ρ	$t_{\max}$				$t'_{\max}$				T_+	
			EB	SR	EB sel	SR sel	EB	SR	EB sel	SR sel	SR	SR sel
-.667	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
	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
.667	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
	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 3.5. Monte Carlo rejection probabilities with 1000 repetitions for Design 1: $\mu = \mathbf{0}$, $n = 400$ 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 $t_{\max}$, $t'_{\max}$ and T_+ statistics, both with pre-selection (sel) and without pre-selection.

Chapter 4

Testing Against Specific Alternatives with Conic Statistics

4.1 Introduction

This paper concerns testing hypotheses on the location of a high-dimensional parameter μ , whose dimension p may exceed the number of observations n . Traditionally, This paper concerns testing hypotheses on the location of a high-dimensional parameter μ , whose dimension p may exceed the number of observations n . Traditionally, such tests are often based on a quadratic statistic like the Wald statistic. However, tests based on quadratic statistics are known to have low power against subsets of the parameter space (henceforth, parameter subspaces) in high dimensions. This has provoked a recent interest in constructing tests that direct power towards parameter subspaces of interest (see e.g. Fan et al., 2015; Kock and Preinerstorfer, 2019).

Such parameter subspaces arise if the econometrician has some prior information about the nature of violations of the null hypothesis. For example,

This chapter is based on Koning (2019).

suppose that we want to test the null hypothesis $H_0 : \boldsymbol{\mu} = \mathbf{0}$, describing a collection of p equalities. Then one may have prior information that if H_0 is false, only some number $k < p$ of these equalities are violated. That is, the parameter $\boldsymbol{\mu}$ is expected to be k -sparse under the alternative (i.e. containing $p - k$ zeros). So, we may want to test against this specific alternative by directing power towards the parameter subspace consisting of k -sparse vectors. The following example describes how such situations can arise in practice.

Example 1. *Multi-factor pricing models describe the excess returns of assets as a linear combination of a limited number of factors. In particular, for a collection of p assets at time n , let $\mathbf{y}_n$ be the p -vector of excess returns, let $\mathbf{f}_n$ be the m -vector containing the m observable factors and let $\boldsymbol{\varepsilon}_n$ be a p -vector of errors. Let $\mathbf{B}$ denote the $p \times m$ matrix of factor loadings, and $\boldsymbol{\mu}$ the p -vector of asset-specific returns that are not captured by the factors (also known as ‘alpha’ in the finance literature). Then, a multi-factor pricing model can be written as*

$$\mathbf{y}_n = \boldsymbol{\mu} + \mathbf{B}\mathbf{f}_n + \boldsymbol{\varepsilon}_n.$$

If the factor model is well specified, efficient pricing theory suggests that the excess returns should be fully captured by the factors so that $\boldsymbol{\mu} = \mathbf{0}$. To assess this theory, one may want to test the null hypothesis $H_0 : \boldsymbol{\mu} = \mathbf{0}$. As such models are well founded in theory, one may expect that if this hypothesis is false then it would only be violated by a small number of exceptional assets. So, if the alternative hypothesis holds, then the underlying vector of excess returns is likely to be sparse. Therefore, one may want to test against the specific alternative that this vector is sparse.

Besides low power against parameter subspaces, another problem with quadratic statistics is that they typically involve an inverse covariance matrix, which is usually not known in practice and must be estimated. This is problematic in high-dimensional settings, as standard estimators of such a covariance matrix are not invertible if $p > n$. Without additional assumptions, this means that the Wald statistic does not exist. The solution provided in the literature is to estimate the covariance matrix by imposing restrictions

on its structure, through regularization or other means (see e.g. Srivastava and Du, 2008; Chen et al., 2011; Fan et al., 2015). However, if no information about the covariance matrix is available, then it is not clear which restrictions should be used. Imposing the wrong restrictions may lead to poor estimates, which in turn may lead to a loss of power.

I simultaneously address these two issues by proposing a novel test statistic that is large in a conic parameter subspace that is specified by the econometrician.¹ This test statistic is a generalization of the Wald statistic. It nests many well-known test statistics that can be recovered by choosing a particular cone. In reference to Example 1, the econometrician would choose a cone that corresponds to a hypothesized level of sparsity under the alternative.

For the test statistic to exist, a restricted eigenvalue condition must be satisfied. The strength of this condition depends directly on the cone of interest. The condition is weaker than the condition that the sample covariance matrix is positive-definite. Unlike the standard Wald statistic, the test statistic can therefore be used even in high-dimensional settings, as long as the cone is sufficiently ‘small’.

In addition, I show that the computation of the test statistic can be formulated as a quadratic minimization problem, where the arguments over which is minimized are restricted to the cone of interest. If the standard sample covariance matrix and sample mean are used as estimators, this problem reduces to regularized linear regression with a constant dependent variable. I provide an additional result for the special case of diagonal estimators of the covariance matrix and cones defined by sign or sparsity restrictions. In that case, the computation reduces to solving a minimum distance problem which can lead to a closed-form test statistic.

For the standard location model, I construct a critical value using reflection-based randomization (see e.g. Lehmann and Romano, 2006). The construction relies on a symmetry assumption on the distribution of the error term, and guarantees size control for a simple null hypothesis in small samples,

¹ A cone $\mathcal{C}$ is a collection of vectors that is closed under positive scalar multiplication. That is, if $\lambda \in \mathcal{C}$, then $\gamma\lambda \in \mathcal{C}$, for all $\gamma > 0$. See e.g. Vock (2007) for an overview of work on testing cone hypotheses, and Wei et al. (2019) for recent work

regardless of p . In addition, I provide a condition that guarantees asymptotic size control by generalizing a result by Koning and Bekker (2019). This condition allows p to be infinite, but instead restricts the dimensionality of the problem through the aforementioned cone of interest.

The test is illustrated by testing against the specific alternative that the parameter is a k -sparse vector. Here, the existence criterion reduces to the requirement that the covariance matrix estimator is at least of rank k . This condition is typically satisfied if $k < n$. The computation of the test statistic is a special case of best-subset selection, also known as ℓ_0 -regularization. Best-subset selection has seen a recent surge of interest after Bertsimas et al. (2016) showed that this problem is easier to solve than was commonly assumed, by exploiting gradient descent methods and commercially available mixed-integer optimization solvers.² If one uses a diagonal covariance matrix estimator, then the test statistic has a closed form and is related to threshold-type statistics (see e.g., Fan, 1996; Zhong et al., 2013).

In addition, I recover the cone for which the computation of the test statistic coincides with Lasso, which is another popular technique used to obtain sparse solutions (Tibshirani, 1996). This cone is centered around sparse vectors, and its width depends on the Lasso regularization parameter. Therefore, the parameter subspace corresponding to Lasso can be interpreted as the collection of vectors that are ‘nearly’ sparse. Intuitively, these may be used if one suspects only a limited number of large violations of the null hypothesis $H_0 : \mu = \mathbf{0}$ might occur.

The power properties of the tests proposed in this paper are studied using Monte Carlo simulation. The results suggest that the tests perform favorably compared to the Wald test if $p < n$ and retains power if $p > n$. In addition, the test typically outperforms the power enhancement technique of Fan et al. (2015).

²See e.g. Hastie et al. (2017); Mazumder et al. (2017); Hazimeh and Mazumder (2018); Koning and Bekker (2018) for more recent work.

4.1.1 Comparison to the Power Enhancement Technique

The power enhancement technique is a recently proposed method that increases power towards a specific alternative of interest. This technique works as follows: given some initial test, one chooses an enhancement test which has asymptotic size zero and is consistent against a parameter subspace, against which the initial test is not consistent. By rejecting if at least one of these tests rejects, the resulting test has asymptotic size equal to the initial test, and may be consistent against a strictly larger set of alternatives. The existence of power enhancement tests has recently been discussed by Kock and Preinerstorfer (2019).

The main advantage of the methodology proposed in this paper over the power enhancement technique is that it is constructive: when the econometrician specifies a specific alternative of interest, this simultaneously defines a test statistic. For the power enhancement technique, both an initial and enhancement test must be chosen by the econometrician. A typical choice for the initial test, as used by Fan et al. (2015), would be based on a Wald-type statistic, which suffers from the issues in high-dimensional settings that were previously described. Finding an appropriate enhancement test for a given specific alternative may also not be straightforward in practice.

The advantage of the power enhancement technique is its attractive asymptotic properties, which promise a strictly larger subset of the parameter space against which the test is consistent while maintaining asymptotic size. It remains unclear how this translates to finite samples. Fan et al. (2015) conduct several numerical experiments that suggest small size distortions may occur in finite samples. Larger size distortions are found in the Monte Carlo simulations presented in Section 4.6.

4.1.2 Notation

For the remainder of the paper, the following notation is used. For a set $\mathcal{A}$, the set $\mathcal{A}^0 = \mathcal{A} \cup \{0\}$ denotes the union of the set and the origin. The set $\mathcal{A}^\circ = \mathcal{A} \setminus \{0\}$ excludes the origin. For a cone $\mathcal{C}$ and matrix $\mathbf{M}$, let $\mathbf{MC} =$

$\{x \mid x = M\lambda, \exists \lambda \in \mathcal{C}\}$. Let $I_n = (e_1, e_2, \dots, e_n)$ denote the $n \times n$ identity matrix, and ι_n an n -vector of ones. For a matrix X with full column rank, the matrix $P_X = X(X'X)^{-1}X'$ denotes the projection matrix of X . For a vector x with elements x_j , $j = 1, \dots, p$, the subscripted vector $x_{\mathcal{J}}$ has j th element equal to x_j if $j \in \mathcal{J}$ and 0, otherwise. Finally, the following standard notation is used

$$\|x\|_q = \begin{cases} \sum_j 1_{\{x_j \neq 0\}}, & q = 0, \\ (\sum_j |x_j|^q)^{\frac{1}{q}}, & 0 < q < \infty. \end{cases}$$

4.2 From Quadratic to Conic statistics

Let μ be the p -vector of interest from the parameter space $\mathbb{R}^p$, and let m be an estimator for μ with covariance matrix $\Sigma > 0$. The matrix $S \geq 0$ denotes a generic estimator of Σ . If S is positive definite, a quadratic or Wald-type statistic that is typically used for testing the hypothesis $H_0 : \mu = \mathbf{0}$ is of the form

$$T = \sqrt{m'S^{-1}m},$$

which is a weighted Euclidean norm of m . Let us slightly extend its definition so that $T = 0$ if $m = \mathbf{0}$ and S is singular. Then, T can be equivalently written as a maximization problem. Specifically, let $\mathcal{S} = \{\lambda \in \mathbb{R}^p \mid \lambda'S\lambda = 1\}$ denote the unit ellipsoid induced by S , then

$$T = \max_{\theta' \theta = 1} m'S^{-\frac{1}{2}}\theta = \max_{\lambda'S\lambda = 1} m'S^{-\frac{1}{2}}S^{\frac{1}{2}}\lambda = \max_{\lambda \in \mathcal{S}} m'\lambda = \max_{\lambda \in \mathcal{S}^0} m'\lambda,$$

Table 4.1. Several test statistics that are nested by the $T_{\mathcal{C}}$ statistic.

$\mathcal{C}$	$T_{\mathcal{C}}$ statistic
$\mathbb{R}^p$	$\sqrt{\mathbf{m}'\mathbf{S}^{-1}\mathbf{m}}$ (Wald statistic)
$\mathcal{C}_+ = \{\boldsymbol{\lambda} \mid \boldsymbol{\lambda} \geq \mathbf{0}\}$	T_+ (Koning and Bekker, 2019)
$\mathcal{C}_{(1)} = \{\boldsymbol{\lambda} \mid \lambda_j = 0, j \neq 1\}$	$ m_1 /s_1$ (two-sided t -statistic)
$\mathcal{C}_{(1)} \cap \mathcal{C}_+$	$(m_1/s_1)^+$ (one-sided t -statistic)
$\mathcal{C}_1 = \{\boldsymbol{\lambda} \mid \ \boldsymbol{\lambda}\ _0 \leq 1\}$	$\max_j m_j /s_j$ (Bugni et al., 2016)
$\mathcal{C}_1 \cap \mathcal{C}_+$	$\max_j (m_j/s_j)^+$ (Chernozhukov et al., 2019a)
$\{(-1, 1, \mathbf{0}), (1, -1, \mathbf{0})\}$	$ m_1 - m_2 /\sqrt{s_1^2 + s_2^2}$, if $\mathbf{S}$ is diagonal (Welch, 1947)

This table contains some examples of statistics that are nested by the $T_{\mathcal{C}}$ statistic. The left column describes the cone $\mathcal{C}$ and the right column contains the corresponding $T_{\mathcal{C}}$ statistic. The notation $(\cdot)^+ = \max\{\cdot, 0\}$ is used for the positive-part operator.

where the first step follows from the Cauchy-Schwarz inequality.³ Notice that the right-hand side reformulation does not use the inverse of $\mathbf{S}$, but the equality holds as the maximization is unbounded if $\mathbf{S}$ is singular and $\mathbf{m} \neq \mathbf{0}$.

The idea to arrive at a conic statistic, is to extend the definition of T by restricting the weights $\boldsymbol{\lambda}$ to a section of $\mathcal{S}$ that represents a direction of interest. The resulting statistic may exist even if $\mathbf{S}$ is singular. Specifically, let $\mathcal{C} \subseteq \mathbb{R}^p$ be a (not necessarily convex) cone and assume that $\mathcal{C}$ is closed in $\mathbb{R}^p$. Then, I propose the following ‘conic’ statistic

$$T_{\mathcal{C}} = \max_{\boldsymbol{\lambda} \in \mathcal{S}^0 \cap \mathcal{C}} \mathbf{m}'\boldsymbol{\lambda}.$$

The dependence on $\mathbf{m}$ and $\mathbf{S}$ is suppressed for notational convenience. Table 4.1 contains an overview of several test statistics that are nested by $T_{\mathcal{C}}$ and can be obtained for different choices of $\mathcal{C}$. Notice that $T = T_{\mathbb{R}^p}$.

³By the Cauchy-Schwarz inequality with $\boldsymbol{\theta} \neq \mathbf{0}$, we have $\mathbf{x}'\boldsymbol{\theta} \leq \sqrt{\mathbf{x}'\mathbf{x}}\sqrt{\boldsymbol{\theta}'\boldsymbol{\theta}}$, so that $\mathbf{x}'\boldsymbol{\theta}/\sqrt{\boldsymbol{\theta}'\boldsymbol{\theta}} \leq \sqrt{\mathbf{x}'\mathbf{x}}$, which implies

$$\sqrt{\mathbf{x}'\mathbf{x}} = \max_{\boldsymbol{\theta}} \frac{\mathbf{x}'\boldsymbol{\theta}}{\sqrt{\boldsymbol{\theta}'\boldsymbol{\theta}}} = \max_{\substack{\boldsymbol{\theta} \\ \boldsymbol{\theta}'\boldsymbol{\theta}=1}} \mathbf{x}'\boldsymbol{\theta}.$$

4.2.1 Existence of T_C

While the quadratic statistic $T_{\mathbb{R}^p}$ requires S to be positive definite in order to exist, T_C typically requires a weaker condition.

Let $\mathcal{H}_x = \{\lambda \mid x'\lambda > 0\}$. If $x \neq \mathbf{0}$, then $\mathcal{H}_x$ describes the open halfspace of $\mathbb{R}^p$ in which x lies, that is supported by the hyperplane on which $x'\lambda = 0$. The following restricted eigenvalue-type condition guarantees that T_C exists.⁴

Condition 4. $\lambda'S\lambda > 0$, for all $\lambda \in \mathcal{C}^\circ \cap \mathcal{H}_m$.

Notice this condition reduces to $\lambda'S\lambda > 0$ for all $\lambda \in \mathbb{R}^p$, if $\mathcal{C} = \mathbb{R}^p$ (and $m \neq \mathbf{0}$). This is equivalent to $S > 0$, which indeed coincides with the condition for the typical Wald-type statistic $T_{\mathbb{R}^p}$ to exist.

Condition 4 is of a weaker form than the typical restricted eigenvalue conditions, due its dependence on m through the inclusion of the restriction $\lambda \in \mathcal{H}_m$. For example, if this restriction is omitted, the cone $\mathcal{C}(\mathcal{A}, L) = \{\lambda \mid \|\lambda_{\mathcal{A}^c}\|_1 \leq L\|\lambda_{\mathcal{A}}\|_1\}$, $L > 0$, $\mathcal{A} \subseteq \{1, \dots, p\}$ produces the well-known compatibility condition used to prove oracle properties of the Lasso (Van de Geer, 2007; Van De Geer et al., 2009).

4.2.2 Geometric interpretation

Condition 4 has the following geometric interpretation. Let a *ray* be defined as half of a line proceeding from the origin. Notice that if $\lambda'S\lambda = 1$ for some λ , then $x'Sx > 0$ for all positive scalar multiples x of λ . Hence, in order to verify Condition 4 we only have to verify whether every ray in $\mathcal{C} \cap \mathcal{H}_m$ intersects with the unit ellipsoid S .

A visual illustration of this geometric interpretation in two dimensions is provided in Figure 4.1. In Panel I, we can see that S is a proper ellipsoid,

⁴First, observe that $T_C = 0$ if $\lambda = \mathbf{0}$. Then for $\lambda \neq \mathbf{0}$, we can write

$$T_{C^\circ} = \max_{\lambda \in C^\circ} \frac{m'\lambda}{\sqrt{\lambda'S\lambda}} = \max_{\substack{\lambda \in C \\ m'\lambda = \gamma}} \frac{\gamma}{\sqrt{\lambda'S\lambda}} = \frac{\gamma}{\min_{\substack{\lambda \in C \\ m'\lambda = \gamma}} \sqrt{\lambda'S\lambda}}, \text{ for all } \gamma > 0.$$

so that every ray intersects $\mathcal{S}$. So, $S > 0$, which implies that Condition 4 holds.

In Panel II, $\mathcal{S}$ is a defect ellipsoid (two parallel lines), so $S \not> 0$. However, all rays in $\mathcal{C}$ intersect $\mathcal{S}$ so that $\lambda' S \lambda > 0$ for all $\lambda \in \mathcal{C}$, and therefore also for all $\lambda \in \mathcal{C} \cap \mathcal{H}_m$. So, Condition 4 does hold.

In Panel III, $\mathcal{S}$ is again a defect ellipsoid. However, there now exists a single ray in $\mathcal{C}$ that does not intersect with $\mathcal{S}$: the ray consisting of the scalar multiples of the vector $(-1, 1)$. So, there exists a $\lambda \in \mathcal{C}$ such that $\lambda' S \lambda = 0$. However, in Panel IV we see that this ray is not contained in $\mathcal{H}_m$. So, Condition 4 holds.

4.2.3 Computation

The following result shows that $T_{\mathcal{C}}$ can be computed using restricted quadratic minimization. This is convenient as some of such problems have been well studied in the optimization literature. A proof of the result can be found in Appendix A.

Proposition 9. *Let $\lambda' S \lambda > 0$ for all $\lambda \in \mathcal{C}^\circ$. Let $\hat{\beta} = \arg \min_{\beta \in \mathcal{C}} \beta' S \beta - 2m' \beta$ and $\hat{\lambda} = \arg \max_{\lambda \in \mathcal{S}^0 \cap \mathcal{C}} m' \lambda$ be unique optimizers. Then*

$$\hat{\lambda} = \frac{\hat{\beta}}{\sqrt{\hat{\beta}' S \hat{\beta}}},$$

if $\hat{\beta} \neq \mathbf{0}$ and $\hat{\lambda} = \mathbf{0}$, otherwise.

In Section 4.5, I discuss a case where the computation reduces to ℓ_0 - and ℓ_1 -regularized regression. Section 4.3.3 covers the special case if S is diagonal, where the computation simplifies for cones that are described by sign and sparsity restriction.

4.3 Hypotheses

In this section, I describe the hypotheses for which I propose $T_{\mathcal{C}}$ to be used as a test statistic.

Figure 4.1. Visual illustration of Condition 4 for the cone $\mathcal{C}(\{1\}, .5)$ (Panels I-II) and the top-left orthant cone (Panels III-IV) in red, a proper ellipsoid S (Panel I) and defect ellipsoid S (Panels II-IV), and a vector m . The half space $\mathcal{H}_m$ is the same for all four panels, but it is only displayed in Panel IV in blue to avoid cluttering. Condition 4 holds if all rays in $\mathcal{C} \cap \mathcal{H}_m$ intersect S .

Figure 4.2. Two illustrations of a cone $\mathcal{C}$ and its polar cone $\mathcal{C}^\circ$. Note that the cones $\mathcal{C}$ in both panels share the same polar cone.

4.3.1 Hypotheses

Let $\mathcal{C}^\circ$ denote the polar cone of $\mathcal{C}$, defined by $\mathcal{C}^\circ = \{\boldsymbol{\mu} \mid \boldsymbol{\mu}'\boldsymbol{\lambda} \leq 0, \forall \boldsymbol{\lambda} \in \mathcal{C}\}$. The polar cone of $\mathcal{C}$ consists of the vectors that have a right angle or larger angle with any vector in $\mathcal{C}$. Intuitively, $\mathcal{C}^\circ$ can be interpreted as the cone that ‘points away’ from $\mathcal{C}$. Illustrations of cones and their polar cone can be found in Figure 4.2.

Notice that the polar cone can equivalently be defined as

$$\mathcal{C}^\circ = \{\boldsymbol{m} \mid \max_{\boldsymbol{\lambda} \in \mathcal{S}^0 \cap \mathcal{C}} \boldsymbol{m}'\boldsymbol{\lambda} = 0\}.$$

So, the polar cone $\mathcal{C}^\circ$ corresponds exactly to the $\boldsymbol{m}$ for which $T_{\mathcal{C}}$ equals zero, and its complement to the $\boldsymbol{m}$ where $T_{\mathcal{C}}$ is positive. Therefore, it seems intuitive to use $T_{\mathcal{C}}$ as a test statistic for the null hypothesis, and its complement as the alternative hypothesis

$$H_0 : \boldsymbol{\mu} \in \mathcal{C}^\circ, \tag{A.1}$$

$$H_1 : \boldsymbol{\mu} \notin \mathcal{C}^\circ, \tag{A.2}$$

where we reject H_0 if T_C is large.

The two following examples demonstrate several hypotheses that can be constructed using different cones $\mathcal{C}$.

Example 2. The cone $\mathcal{C}_+ = \{\lambda \mid \lambda \geq \mathbf{0}\}$ has polar cone $\mathcal{C}_+^\circ = \{\mu \mid \mu \leq \mathbf{0}\}$, producing the multi-dimensional ‘one-sided’ hypotheses $H_0 : \mu \leq \mathbf{0}$, $H_1 : \mu \not\leq \mathbf{0}$. In Table 4.1, we can see that corresponding test statistic is T_+ , which is indeed proposed to test these hypotheses by Koning and Bekker (2019).

Example 3. The cone $\mathcal{C}_{(1)} = \{\lambda \mid \lambda_j = 0, j \neq 1\}$ has polar cone $\mathcal{C}_{(1)}^\circ = \{\mu \mid \mu_1 = 0\}$ corresponding to the one-dimensional simple hypotheses $H_0 : \mu_1 = 0$, $H_1 : \mu_1 \neq 0$. The intersection $\mathcal{C}_{(1)+} = \mathcal{C}_{(1)} \cap \mathcal{C}_+$ has polar cone $\mathcal{C}_{(1)+}^\circ$, which constructs the one-dimensional one-sided hypotheses $H_0 : \mu_1 \leq 0$, $H_1 : \mu_1 > 0$. From Table 4.1 we can see that the corresponding test statistics are the two-sided and one-sided t -statistics, which are indeed used to test these hypotheses.

4.3.2 Specific alternatives

While quadratic statistic are known to perform poorly against parameter subspaces as the dimensionality p increases, conic statistics offer an interesting opportunity to counteract this problem.

Here, the key observation is that different cones $\mathcal{C}$ can share the same polar cone, and therefore produce the same hypotheses (A.1) and (A.2). As different cones can produce the same hypotheses, we have some freedom to change the cone $\mathcal{C}$ without affecting the hypotheses. This allows us to pick the cone $\mathcal{C}$ from a collection of cones that share the same polar cone. This is illustrated in the following example, and can be seen visually in the two panels of Figure 4.2.

Example 4. Let us consider $\mathcal{C}_1 = \{\lambda \mid \|\lambda\|_0 \leq 1\}$, containing the cones that coincide with the axes, and the cone $\mathbb{R}^p$ that contains the entire parameter space. Notice that the polar cone corresponding to both $\mathcal{C}_1$ and $\mathbb{R}^p$ is $\{\mathbf{0}\}$. So, both $\mathcal{C}_1$ and $\mathbb{R}^p$ produce the hypotheses $H_0 : \mu = \mathbf{0}$ and $H_1 : \mu \neq \mathbf{0}$. However, $T_{\mathcal{C}_1}$ and $T_{\mathbb{R}^p}$ are clearly different for almost all $\mathbf{m}$ and $\mathbf{S}$: they only coincide if $\mathbf{m} \in \mathcal{C}_1$.

In addition to specifying a null and alternative hypothesis, this effectively forces the user to designate a conic parameter subspace of interest $\mathcal{C}$ as a *specific alternative*. For notational convenience, I will denote this specific alternative as $H_s : \boldsymbol{\mu} \in \mathcal{C}$.

Remark 3. *One may argue that $H_1 : \boldsymbol{\mu} \in \mathcal{C}$ should be considered the alternative hypothesis, and that the concept of specific alternative is unnecessary. However, in light of Example 3 this would link the one-sided t -statistic to the hypotheses $H_0 : \mu_1 \leq 0$ against $H_1 : \mu_1 > 0, \mu_j = 0, j \neq 1$, which is clearly not desirable.*

Depending on the covariance structure, there may be a distortion so that $T_{\mathcal{C}}$ is not necessarily largest if $\boldsymbol{m} \in \mathcal{C}$. In particular, notice that $T_{\mathbb{R}^p} \geq T_{\mathcal{C}}$, by construction. In order to establish for which cone $T_{\mathcal{C}}$ is large, we can find the $\boldsymbol{m}$ for which $T_{\mathbb{R}^p} = T_{\mathcal{C}}$. From Proposition 9, we find that if $\boldsymbol{m} \neq \mathbf{0}$ then $T_{\mathbb{R}^p}$ is maximized by some $\hat{\boldsymbol{\lambda}} \neq \mathbf{0}$ such that $\gamma S \hat{\boldsymbol{\lambda}} = \boldsymbol{m}$. So, if $\boldsymbol{m} \in SC^{\odot}$, then $T_{\mathcal{C}} = T_{\mathbb{R}^p}$, so that $T_{\mathcal{C}}$ is ‘large’. While this distortion could lead to a loss of power, the resulting test may still be substantially more powerful than a test based on a quadratic statistic.

In addition, there are important examples where $\mathcal{C}$ and SC coincide. If we have some information about the structure of $\boldsymbol{\Sigma}$ we can make use of matrix cone invariance. In particular, there exist cones $\mathcal{C}$ and matrices S such that $SC = \mathcal{C}$. For this to hold, it suffices to guarantee that the equality holds for all elements of the parameter space of $\boldsymbol{\Sigma}$. An example of this is illustrated in Section 4.3.3. A weaker condition is subset invariance $SC \subseteq \mathcal{C}$. An example of this is the case that it is known $\boldsymbol{\Sigma}$ has positive elements and $\mathcal{C}$ is the non-negative orthant (see Koning and Bekker, 2019).

4.3.3 Diagonal covariance matrices and scones

Suppose that it is known that the elements of $\boldsymbol{m}$ are uncorrelated, so that $\boldsymbol{\Sigma}$ is diagonal with positive diagonal elements. Then, there exist cones $\mathcal{C}$ for which $\mathcal{C} = \boldsymbol{\Sigma}\mathcal{C}$. In particular, let $\mathcal{D}$ be a set of vectors that is closed under multiplication by positive definite diagonal matrices. That is, if $\boldsymbol{\lambda} \in \mathcal{D}$, then

$D\lambda \in \mathcal{D}$, for all diagonal matrices $D > 0$. Notice here that $\mathcal{D}$ is a cone.⁵ So, we have indeed found a type of cone and covariance matrix for which $\mathcal{D} = \Sigma\mathcal{D}$.

I have been unable to find a name of these types of cones and will henceforth refer to them as *sign cones*. A sign cone can be viewed as a union of sub-orthants (i.e. a union of possibly lower dimensional orthants embedded in a higher dimensional space). As pre-multiplication by a diagonal matrix $D > 0$ is a sign-preserving operation, the sign-conic parameter subspaces relate directly to hypotheses concerning sign or sparsity restrictions. To provide some more intuitions for sign cones, the following examples are given.

Example 5. *The sign cones in $\mathbb{R}^2$ are the origin, the four half-axes, the four quadrants, and any union of any these objects.*

Example 6. *The set of k -sparse vectors, defined by $\mathcal{D}_k = \{\mu \mid \|\mu\|_0 \leq k\}$ is a sign cone. To see this, suppose $\mu = (\mu_1, \dots, \mu_k, 0, \dots, 0)'$, without loss of generality. So $\mu \in \mathcal{D}_k$. Let $D > 0$ have diagonal elements $d_1, \dots, d_p$. Then $\|D\mu\|_0 = \|(d_1\mu_1, \dots, d_k\mu_k, 0, \dots, 0)'\|_0 = \|\mu\|_0$. So $D\mu \in \mathcal{D}_k$.*

If Σ is assumed to be diagonal, it would be sensible to use a diagonal estimator S , which ensures $\mathcal{D} = S\mathcal{D}$. In this case, the computation of $T_{\mathcal{C}}$ can be simplified substantially. In particular, Proposition 10 shows that computing $T_{\mathcal{C}}$ requires solving a minimum distance problem. Such problems are typically easier to solve than linear regression. For example, the following section considers testing against sparse subspaces, for which this problem has a closed-form solution.

Proposition 10. *Let $S > 0$ be a diagonal matrix and let $\mathcal{D}$ be a scone. Suppose that $\hat{\lambda} = \arg \max_{\lambda \in S^0 \cap \mathcal{D}} m' \lambda$ and $\hat{\beta} = \arg \min_{\beta \in \mathcal{D}} \|S^{-\frac{1}{2}}m - \beta\|_2^2$ are unique optimizers. Then*

$$\hat{\lambda} = S^{-\frac{1}{2}}\hat{\beta} / \sqrt{\hat{\beta}'\hat{\beta}},$$

⁵ This follows from the observation that multiplication by a scalar $\gamma > 0$ is equivalent to pre-multiplication with by diagonal matrix γI , where I is the identity matrix. So, $\mathcal{D}$ is also closed under positive scalar multiplication, which makes it a cone.

if $\hat{\beta} \neq \mathbf{0}$ and $\hat{\lambda} = \mathbf{0}$, otherwise.

4.4 Randomization Test in Location Model

Let us consider testing H_0 against H_1 with specific alternative $H_s : \mu \in \mathcal{C}$ in the standard location model, where $n \times p$ data matrix X can be decomposed as

$$X = \iota_n \mu' + E,$$

where E is a random matrix. In light of this data generating process, define $m = \hat{\mu} := \frac{1}{n} X' \iota_n$ to be the sample mean and $S = \hat{\Sigma} := \frac{1}{n} X' (I_n - P_{\iota_n}) X$ to be the sample covariance matrix.

Deriving null distributions of $T_{\mathcal{C}}$ for general cones is challenging.⁶ Therefore, I construct a critical value based on a randomization assumption. Let $\mathcal{R}$ contain all $n \times n$ reflection matrices, which are diagonal matrices with diagonal elements in $\{-1, 1\}$. So $\mathcal{R}$ constitutes the reflection group of order $N = 2^n$ under matrix multiplication. I then make the following reflection symmetry assumption on the error term.

Condition 5. $E \stackrel{(d)}{=} RE$, for all $R \in \mathcal{R}$.

Based on Condition 5, a randomization test can then be conducted as follows (see e.g. Ch. 15 of Lehmann and Romano (2006) for a more general discussion). Let $\mathcal{R}^M := \{I_n, R_2, \dots, R_M\} \subseteq \mathcal{R}$, where R_i is drawn without replacement from $\mathcal{R} \setminus \{I_n\}$, $i = 2, \dots, M$. Here, $M \leq N$ is the number of re-samples. Denote the set containing the reflection transformed data by $\mathcal{R}_X^M = \{X, R_2 X, \dots, R_M X\}$. Let

$$T_{\mathcal{C}}(\mathcal{R}_X^M) = \{T_{\mathcal{C}}(X), T_{\mathcal{C}}(R_2 X), \dots, T_{\mathcal{C}}(R_M X)\},$$

be the set containing test statistic of interest computed on each of the reflection transformed data sets. To ensure $T_{\mathcal{C}}(\mathcal{R}_X^M)$ consists of M unique ele-

⁶See e.g. Chernozhukov et al. (2019a,b), for special case of cones consisting of a finite collection of rays.

ments, I assume random infinitesimal perturbations are added so that any ties are resolved with equal probability in either direction.

Let $\alpha \in \{1/M, 2/M, \dots, 1\}$ and suppose that $\mathcal{C}^\circ = \{\mathbf{0}\}$, so that $H_0 : \mu = \mathbf{0}$. Then, we have $\mathbf{X} = \mathbf{E}$ under null, so that each of the elements of $T_{\mathcal{C}}(\mathcal{R}_X^M)$ follow the same distribution. So, under H_0 , the probability that $T_{\mathcal{C}}(\mathbf{X})$ is larger than the $1 - \alpha$ quantile c_α among the elements of $T_{\mathcal{C}}(\mathcal{R}_X^M)$ is α . Therefore, we arrive at the well known result that the test $T_{\mathcal{C}}(\mathbf{X}) > c_\alpha$ has size α . In order to present the upcoming results, let us also define $p^{Tc}(\mathbf{X})$ to be the proportion of elements in $T_{\mathcal{C}}(\mathcal{R}_X^M)$ that are larger than or equal to $T_{\mathcal{C}}(\mathbf{X})$.

If $\mathcal{C}^\circ \neq \{\mathbf{0}\}$, exactness can no longer be guaranteed. However, I provide the following condition that guarantees size control in large samples by generalizing a result of Koning and Bekker (2019).

Condition 6. As $n \rightarrow \infty$, $\max_{\lambda \in \mathcal{S}^0 \cap \mathcal{C}} |\lambda' \mathbf{E}' \mathbf{t}_n / n| \xrightarrow{a.s.} 0$ and $\max_{\lambda \in \mathcal{S}^0 \cap \mathcal{C}} |\lambda' (\mathbf{S} - \Sigma) \lambda| \xrightarrow{a.s.} 0$.

Note that this condition permits p to grow arbitrarily in n , but instead restricts the ‘dimensionality’ through $\mathcal{C}$. For example, for a given $\mathbf{E}$, Condition 6 may hold for $\mathcal{C} = \{e_1\}$, but not for $\mathcal{C} = \mathbb{R}^p$. Condition 6 is of a weaker form than the condition required for the T_+ statistic of Koning and Bekker (2019), who require fixed p , while Condition 6 only requires $\mathcal{C}$ to be fixed. The reason for this difference is that the cone they implicitly consider, $\mathcal{C}_+ = \{\lambda \mid \lambda \geq \mathbf{0}\}$, grows in p .

Proposition 11. If $\mu \in \mathcal{C}^\circ$, and Condition 5 and 6 hold then $P(\lim_{n \rightarrow \infty} p^{Tc}(\mathbf{X}) \leq \alpha) \leq \alpha$.

A proof can be found in Appendix 4.8. Interestingly, similar reasoning can be used to show that the test has asymptotic power of at least α over the dual cone $\mathcal{C}^* = \{\mu \mid \mu' \lambda \geq 0, \forall \lambda \in \mathcal{C}\}$ of $\mathcal{C}$. Its proof is also provided in Appendix 4.8.

Proposition 12. If $\mu \in \mathcal{C}^*$ and Conditions 5 and 6 hold, then $P(\lim_{n \rightarrow \infty} p^{Tc}(\mathbf{X}) \leq \alpha) \geq \alpha$.

4.5 Sparse specific alternatives

In this section, I illustrate the methodology described in the previous sections by testing against *sparse* specific alternatives, where the cone $\mathcal{C}$ corresponds to the parameter subspace consisting of vectors that contain zeros.

Such a sparse parameter subspace can be described as follows. For a given level of sparsity k , the k -sparse parameter subspace $\mathcal{D}_k = \{\boldsymbol{\mu} \in \mathbb{R}^p \mid \|\boldsymbol{\mu}\|_0 \leq k\}$ consists of all k -sparse vectors with p elements. Recall here that $\|\mathbf{x}\|_0 = \sum_{j=1}^p 1_{\{x_j \neq 0\}}$ is the ℓ_0 -norm of a vector $\mathbf{x}$, which counts the number of non-zero elements in the vector. As detailed in Example 6, $\mathcal{D}_k$ is a sign cone, because $\|\mathbf{x}\|_0 = \|\mathbf{D}\mathbf{x}\|_0$ for all diagonal matrices $\mathbf{D} > 0$.

Note that $\mathcal{D}_k^\circ = \{\mathbf{0}\}$. So, following Section 4.3 the hypotheses are

$$H_0 : \boldsymbol{\mu} = \mathbf{0},$$

$$H_1 : \boldsymbol{\mu} \neq \mathbf{0},$$

$$H_s : \boldsymbol{\mu} \in \mathcal{D}_k.$$

The interpretation of this specific alternative is that (at most) k elements of the equalities specified by the null hypothesis are violated. The corresponding test statistic is $T_{\mathcal{D}_k}$, which will be abbreviated as T_k for notational convenience.

Following Condition 4, a sufficient existence condition for T_k is then given by

Condition 7. $\text{rank}(\widehat{\boldsymbol{\Sigma}}) \geq k$.

This condition is typically satisfied if the number of observations n exceeds the hypothesized sparsity level k . So, if the hypothesized sparsity level is sufficiently small, is a much weaker condition than the requirement that $\widehat{\boldsymbol{\Sigma}} > 0$ which requires that the number of observations exceeds the total number of parameters p .

Although the maximizing vector $\widehat{\boldsymbol{\lambda}}$ in the computation of T_k can be viewed as a by-product, it may also be of interest in its own right. In some applications it may be interesting to not just know *whether* $H_0 : \boldsymbol{\mu} = \mathbf{0}$ is violated,

but also *which* of the p equalities specified by the null hypothesis are violated. A collection of k potentially violated equalities is provided by the k non-zero elements of $\hat{\lambda}$. The computation of $\hat{\lambda}$ is discussed in the following two sections.

4.5.1 Computation: full covariance matrix

For the case that $\mathbf{m} = \hat{\boldsymbol{\mu}}$ and $\mathbf{S} = \hat{\boldsymbol{\Sigma}}$, Proposition 13 shows that the computation of $T_{\mathcal{C}}$ reduces to regularized linear regression where the regularization is determined by $\mathcal{C}$. The result is proven in Appendix 4.8.

Proposition 13. *Let $\boldsymbol{\lambda}'\hat{\boldsymbol{\Sigma}}\boldsymbol{\lambda} > 0$ for all $\boldsymbol{\lambda} \in \mathcal{C}^\circ$. Let $\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta} \in \mathcal{C}} \frac{1}{n} \|\boldsymbol{\iota}_n - \mathbf{X}\boldsymbol{\beta}\|_2^2$ and $\hat{\boldsymbol{\lambda}} = \arg \max_{\boldsymbol{\lambda} \in (\mathcal{S} \cap \mathcal{C}) \cup \{\mathbf{0}\}} \boldsymbol{\mu}'\boldsymbol{\lambda}$ be unique optimizers. Then*

$$\hat{\boldsymbol{\lambda}} = \hat{\boldsymbol{\beta}} / \sqrt{\hat{\boldsymbol{\beta}}'\hat{\boldsymbol{\Sigma}}\hat{\boldsymbol{\beta}}},$$

if $\hat{\boldsymbol{\beta}} \neq \mathbf{0}$ and $\hat{\boldsymbol{\lambda}} = \mathbf{0}$, otherwise.

If $\mathcal{C} = \mathcal{D}_k$, this result implies that T_k can be computed using sparse linear regression or ‘best subset selection’ (BSS), which is defined as follows

$$\hat{\boldsymbol{\beta}} = \arg \min_{\substack{\boldsymbol{\beta} \\ \|\boldsymbol{\beta}\|_0 \leq k}} \|\boldsymbol{\iota}_n - \mathbf{X}\boldsymbol{\beta}\|_2^2. \quad (\text{A.3})$$

While this optimization problem was long deemed computationally infeasible for $p \gtrsim 40$, Bertsimas et al. (2016) have recently shown that it can be solved for problems of practical size within reasonable time by using gradient descent methods and mixed-integer optimization solvers. In particular, they solve BSS with n of order 10^3 and p of order 10^2 in minutes. If p and n are very large or if re-sampling is used to construct a critical value, then this may still be costly in terms of computation to conduct a test. However, promising work by Hazimeh and Mazumder (2018) shows that problems of order $p = 10^6$ can be approximated in seconds. An alternative approach is to approximate the solution of (A.3) using greedy methods such as forward stepwise selection (see e.g. Hastie et al. 2009, 2017).

4.5.2 Computation: diagonal covariance matrix

Let us now also consider the computation of the statistic for the case that S is a positive definite diagonal estimator. Popular examples of diagonal estimators are $\text{diag}(\widehat{\Sigma})$ or the pooled variance estimator $\text{tr}(\widehat{\Sigma}/p)I_p$. In order to distinguish the statistic for the diagonal covariance estimator, I use the superscripted notation T_k^d . Proposition 10 then shows that if $S > 0$, T_k^d can be computed by solving

$$\min_{\substack{\beta \\ \|\beta\|_0 \leq k}} \|S^{-\frac{1}{2}}\widehat{\mu} - \beta\|_2^2.$$

It is straightforward to show that $\widehat{\beta} = S^{-1}\widehat{\mu}_{\mathcal{J}_k}$, where $\mathcal{J}_k$ contains the largest k elements of $S^{-1}\widehat{\mu}$.⁷ This leads to the closed form

$$T_k^d = \sqrt{\widehat{\mu}' S^{-1} \widehat{\mu}_{\mathcal{J}_k}}.$$

If $S = \text{diag}(\widehat{\Sigma})$, the T_k^d statistic is closely related to threshold-type statistics (see e.g. Fan, 1996; Zhong et al., 2013), which include the screening statistic proposed by Fan et al. (2015). They define the screening statistic as

$$J_0 = \sqrt{p}\widehat{\mu}' \text{diag}(\widehat{\Sigma})^{-1} \widehat{\mu}_{\mathcal{J}},$$

where $\mathcal{J} = \{j \in \{1, \dots, p\} \mid |\widehat{\mu}_j| > \widehat{\sigma}_j \delta\}$, for a given threshold value δ . Here, Fan et al. (2015) choose δ to grow sufficiently fast so that using 0 as critical value results in a test with asymptotic size 0. Notice that if δ is instead chosen such that $|\mathcal{J}| = k$, then J_0 is a monotonic function of T_k^d : $J_0 = \sqrt{p}(T_k^d)^2$. Hence, the difference is that a threshold-type statistic implicitly defines the sparsity level, while the T_k^d explicitly specifies the sparsity level of the alternative. The latter option seems more intuitive, though there may be applications for which a threshold-type specification of the sparsity level is desirable.

⁷ See e.g. Proposition 3 of Bertsimas et al. (2016) for a formal proof.

4.5.3 Lasso

In this section, I demonstrate that testing with a conic statistic may also be useful if the parameter subspace of interest is not a cone. To do so, I consider the popular sparsifying tool Lasso (Tibshirani, 1996), which is linear regression under an ℓ_1 -norm restriction.

Unlike $\mathcal{D}_k$, the Lasso constraint set $\mathcal{B}_t = \{\lambda \mid \|\lambda\|_1 \leq t\}$, where $t \geq 0$, is not a sign cone nor even a cone. However, we can generate cones using this set, as follows. Let $\Lambda_t = \mathcal{S} \cap \mathcal{B}_t$ be the intersection of the Lasso constraint set and $\mathcal{S}$, for some $t \geq 0$. Then we can construct the cone $\mathcal{C}_{\Lambda_t} = \{\lambda \mid \gamma\lambda \in \Lambda_t, \exists \gamma > 0\} \cup \{\mathbf{0}\}$ that consists of all the scalar multiples of Λ_t and the origin.

Notice that if $\mathcal{S} = s^2 \mathbf{I}_p$, where $s > 0$, then $\mathcal{C}_{\Lambda_t}$ coincides with $\mathcal{D}_1$ when $t = 1/s$, and $\mathbb{R}^p$ if $t \geq \frac{\sqrt{p}}{s}$. Furthermore, if $t \geq 1/s$, $\mathcal{C}_{\Lambda_t}^\circ = \{\mathbf{0}\}$. So, according to the construction in Section 4.3, this yields the hypotheses $H_0 : \mu = \mathbf{0}$, $H_1 : \mu \neq \mathbf{0}$, and $H_s : \mathcal{C}_{\Lambda_t}$.

A visual illustration of the two-dimensional case is given in Figure 4.3, for $t = 1.2$ and $s = 1$. This figure shows that the specific alternative corresponding to the Lasso constraint is set of convex cones that are centered around the axes, and whose widths are regulated by the parameter t . Therefore, the specific alternative induced by Lasso can be viewed as directing power towards *nearly-sparse* violations. This interpretation of a nearly-sparse specific alternative extends to higher dimensions.

By Proposition 13, the computation of the test statistic $T_{\mathcal{C}_{\Lambda_t}}$ reduces to Lasso regression

$$\hat{\beta} = \arg \min_{\substack{\beta \\ \|\beta\|_1 \leq t / \sqrt{\hat{\beta}' \hat{\Sigma} \hat{\beta}}}} \|\iota_n - X\beta\|_2^2,$$

so that $\hat{\lambda} = \hat{\beta} / \sqrt{\hat{\beta}' \hat{\Sigma} \hat{\beta}}$.⁸

⁸Unfortunately, as of writing this manuscript, all existing efficient implementations of Lasso that I could find fail for a constant dependent variable. This has prohibited the inclusion of the Lasso based test in the Monte Carlo experiments. The authors of several

Figure 4.3. The construction of a Lasso specific alternative in two dimensions for $S = I_2$. The area inside the diamond corresponds to the Lasso constraint set $\mathcal{B}_{1,2}$, the circle is the unit sphere $\mathcal{U}$, the thicker sections of the circle correspond to the intersection $\Lambda_{1,2} = \mathcal{U} \cap \mathcal{B}_{1,2}$, and the cone corresponding to the intersection $\mathcal{C}_{\Lambda_{1,2}}$ is shown as the gray dashed areas.

4.6 Simulation experiments

In this section, the Monte Carlo simulation experiments are presented. In particular, I consider testing against sparse specific alternatives, where the level of sparsity is varied. The goal of the experiments is to analyze the size and power properties of the T_C test, and to compare it to the power enhancement approach of Fan et al. (2015). For this purpose, I will vary the sample size, the number of parameters, the underlying covariance matrix and the level of sparsity in the parameter vector. Code to replicate the experiments will be made available at <https://github.com/nickwkoning>.

4.6.1 Data generation

I generate the $n \times p$ data matrix $X = \iota_n \mu' + EA$, where the elements of E are independently drawn from a standard normal distribution, and $A'A = \Sigma$, where Σ has off-diagonal elements $\rho \in \{0, .5, .7\}$, $i \neq j$ and diagonal elements equal to 1. For the number of observations and parameters, I consider $n \in \{30, 250\}$ and $p \in \{100, 300, 500\}$, respectively.

The parameter vector of interest μ has elements $\mu_j = b_{n,s} > 0$ for $1 \leq j \leq s$ and $\mu_j = 0$, if $s < j \leq p$. So, μ is s -sparse. The values of $b_{n,s}$ are chosen to ensure that all tests that control size have rejection rates smaller than 1 at the significance level $\alpha = 0.05$. In particular, I use

$$b_{n,s} = \begin{cases} 0.75, & \text{if } n = 30, \quad s = 1, \\ 0.25, & \text{if } n = 30, \quad s = 20, \\ 0.25, & \text{if } n = 250, \quad s = 1, \\ 0.07, & \text{if } n = 250, \quad s = 20. \end{cases}$$

prominent Lasso implementations have been made aware of this issue.

4.6.2 Tests

On the data described above, I test the null hypothesis $H_0 : \mu = \mathbf{0}$ against the alternative $H_1 : \mu \neq \mathbf{0}$, with the s -sparse specific alternative $H_1^s : \mu \in \mathcal{D}_s$. As this sparsity level may not be known in practice, I also consider a possible over- and underspecification of s by testing against the specific alternative $H_1^1 : \mu \in \mathcal{D}_1$, if $s = 20$, and against $H_1^{20} : \mu \in \mathcal{D}_{20}$, if $s = 1$.

In particular, I use the T_k statistic where $k \in \{1, 20\}$. As estimators for μ and Σ , I choose the sample mean $\mathbf{m} = \mathbf{X}'\mathbf{t}_n/n$, and sample covariance $\mathbf{S} = \widehat{\Sigma} := \mathbf{X}'(\mathbf{I}_n - \mathbf{P}_{\mathbf{t}_n})\mathbf{X}/n$. In addition, I also consider the T_k^d statistic with $k \in \{1, 20\}$, where $\mathbf{S} = \text{diag}(\widehat{\Sigma})$. Note here that the T_1^d statistic and T_1 statistic coincide.

The T_k statistic is computed using the methodology proposed by Hazi-meh and Mazumder (2018). To construct a critical value for T_k and T_k^d , I use reflection based randomization as described in Section 4.4. Note that this will yield an exact test, as the normal distribution is reflection symmetric. The number of reflection re-samples used is 1000. Performing a single test with the T_{20} statistic on the largest setting with $n = 250$ and $p = 500$ takes approximately 45 seconds on a standard 2017 edition 13 inch Macbook Pro.

For the comparison to the power enhancement technique, an initial test and power enhancement test must be selected. As initial test I use the standard Wald test, and as statistic for the enhancement test I use the same screening statistic as used by Fan et al. (2015).⁹ This screening statistic is given by

$$J_0 = \sqrt{p}\mathbf{m}' \text{diag}(\widehat{\Sigma})^{-1}\mathbf{m}_{\mathcal{J}},$$

with $\mathcal{J} = \{j \in \{1, \dots, p\} \mid |m_j| > \widehat{\sigma}_j \delta_{n,p}\}$, where $\widehat{\sigma}_j^2$ are the diagonal elements of $\widehat{\Sigma}$, and $\delta_{n,p} = \log \log n \sqrt{\log p}$.¹⁰ As critical value for the Wald

⁹One may argue that a comparison could instead be made to the ‘feasible Wald test’ proposed as initial test by Fan et al. (2015). However, this test relies on an additional sparsity assumption on Σ which would cloud the comparison.

¹⁰The code provided the online supplementary material of Fan et al. (2015) suggests that $1.06\delta_{n,p}$, $\sqrt{1.5}\delta_{n,p}$ and $.9\delta_{n,p}$ were used in their numerical experiments, instead of $\delta_{n,p}$. It is unclear where these constants come from, as they are not mentioned in the paper. I therefore choose to simply use $\delta_{n,p}$, instead.

statistic I use $pn/(n - p)$ times the $1 - \alpha$ quantile of the F distribution with p and $n - p$ degrees of freedom (also known as a Hotelling T^2 test). If $p > n$, the initial test is not defined and I reject with probability α . As critical value for the enhancement test I use 0. The combined test then rejects if either the initial test or enhancement test rejects.¹¹ The simulation experiments are also conducted on the initial test to infer the contribution of the constituent tests to the performance of the power enhancement technique.

For each combination of setting and test, I perform 1000 repetitions and record the proportion of rejections. The results are reported in Tables 4.2, 4.3 and 4.4.

4.6.3 Results

Table 4.2 reports the rejection rates of the tests on data that was generated under the null hypothesis. As these rejection rates are used to infer the size of the tests, they should ideally be close to the nominal significance level $\alpha = 0.05$. As expected, the table shows that the T_C tests and the Wald test all have good size control under H_0 .

In contrast, the power enhancement test suffers from a small size distortion in the large sample ($n = 250$). Furthermore, as the power enhancement test is based on an asymptotic critical value it fails to control size in the small sample ($n = 30$). The rejection rates for the power enhanced test for $n = 30$ will therefore be ignored in the remainder without mention.

Table 4.3 contains the rejection rates under the alternative where $s = 1$, so that only a single element of μ is unequal to zero. As this data is generated under the alternative, the rejection rates are used to infer the power of the test, which should be as large as possible. Overall, the T_1 test outperforms all other tests. For the large sample ($n = 250$), the power enhanced test performs second best and delivers a large power improvement compared to the Wald test. Despite the mis-specification of the sparsity level, the T_{20}

¹¹ The formulation of the power enhancement technique here differs slightly from the one described by Fan et al. (2015) and is of the form analyzed by Kock and Preinerstorfer (2019). This formulation permits the use of the power enhancement approach even if the test statistic for the initial test does not exist.

test performs reasonably well if $n = 250$, substantially outperforming the Wald test. The T_{20} test performs especially well under large correlations. As expected, the T_{20}^d test performs better than the T_{20} test under no correlations, but loses power as the correlations increase.

Table 4.4 presents the results for $s = 20$, which are more diverse than the results in Table 4.2 and 4.3. The T_{20}^d test outperforms the other tests if $\rho = 0$ as the sparsity level is now correctly specified. If $\rho \neq 0$ and $n = 30$, the T_1 test outperforms the T_{20} test even though $s = 20$. This may be a consequence of the low rank of $\hat{\Sigma}$, so that Condition 7 is barely satisfied. For $n = 250$ and $\rho \neq 0$, the power enhancement test performs well in the low-dimensional setting $p = 100$. This performance seems almost entirely carried by the initial Wald test, which performs well as s is large compared to p . In the settings where $p > n$, the T_{20} test substantially outperforms the other tests.

From the results described above, the following conclusions are drawn.

- The T_k^d test performs well if Σ is diagonal and k is close to s . It loses power as k is further from s or Σ is further from diagonal.
- If Σ is not diagonal, the T_k test performs well if k is close to the true sparsity level s , and n is sufficiently large compared to k .
- The power enhancement test trades a small size distortion for a possibly large power gain in large samples.
- The power enhancement test is competitive with the T_s test if s is small, but loses power compared to the T_s test if s is somewhat larger.

4.7 Discussion

The goal of this paper is to discuss an alternative to quadratic statistics in hypothesis testing situations where they suffer from existence and power loss issues, with a special focus on $p > n$ settings.

A potential solution treated here is to test using a conic statistic, which relies on the choice of a cone. By selecting a sufficiently small cone, the existence issue is immediately addressed through the existence Condition 4. At the same time, I proposed to set the cone equal to the specific alternative of interest, in order to divert power towards it. Together, this addresses both the issues encountered with quadratic statistics.

However, there are some limitations. First of all the choice of cone. This choice is mostly done for practical reasons, as no translation has to be done from alternative to cone. However, it is not obvious whether this choice is in any sense ‘optimal’ from a power perspective: there may exist another cone that yields a test with higher power. In particular, this may depend on the covariance matrix. Yet, the simulation results seem promising in comparison to the power enhancement method.

Further research is also warranted in exploring the trade-off in the choice of cone, in the way that it influences both the existence of the statistic and the specific alternative. This link causes a curious balance: depending on the setting, changing the cone may be beneficial from an existence perspective, but undesirable in the selection of specific alternative. Here, both the existence and choice of specific alternative have consequences for the power of the test.

4.8 Appendix A

In order to present the proofs of Proposition 9, 10 and 13, I first provide the following simple lemma, where the notation $\gamma\mathcal{A}$ is used to mean the set containing the elements of $\mathcal{A} \subseteq \mathbb{R}^p$ scalar multiplied by γ .

Lemma 2. *Let $\mathcal{C}$ be a cone and $\gamma > 0$. If Condition 4 holds, then*

$$\arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda' S \lambda = \gamma^2}} m' \lambda = \gamma \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda' S \lambda = 1}} m' \lambda.$$

Proof. As Condition 4 holds, a maximizing argument exists. I find

$$\arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda' S \lambda = \gamma^2}} m' \lambda = \arg \max_{\substack{\lambda \in \mathcal{C} \\ \frac{1}{\gamma^2} \lambda' S \lambda = 1}} m' \lambda = \arg \max_{\substack{(\frac{1}{\gamma} \lambda) \in \mathcal{C} \\ (\frac{1}{\gamma} \lambda)' S (\frac{1}{\gamma} \lambda) = 1}} m' \lambda = \gamma \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda' S \lambda = 1}} m' \lambda.$$

■

Proposition 9. Let $\lambda' S \lambda > 0$ for all $\lambda \in \mathcal{C}^\circ$. Let $\hat{\beta} = \arg \min_{\beta \in \mathcal{C}} \beta' S \beta - 2m' \beta$ and $\hat{\lambda} = \arg \max_{\lambda \in \mathcal{S}^0 \cap \mathcal{C}} m' \lambda$ be unique optimizers. Then

$$\hat{\lambda} = \frac{\hat{\beta}}{\sqrt{\hat{\beta}' S \hat{\beta}}},$$

if $\hat{\beta} \neq \mathbf{0}$ and $\hat{\lambda} = \mathbf{0}$, otherwise.

Proof. I will consider two cases: $\hat{\beta} \neq \mathbf{0}$ and $\hat{\beta} = \mathbf{0}$.

Case $\hat{\beta} \neq \mathbf{0}$

I find

$$\begin{aligned} \hat{\lambda} &= \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda' S \lambda = 1}} m' \lambda = \frac{1}{\sqrt{\hat{\beta}' S \hat{\beta}}} \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda' S \lambda = \hat{\beta}' S \hat{\beta}}} m' \lambda \\ &= \frac{1}{\sqrt{\hat{\beta}' S \hat{\beta}}} \arg \min_{\substack{\lambda \in \mathcal{C} \\ \lambda' S \lambda = \hat{\beta}' S \hat{\beta}}} \lambda' S \lambda - 2m' \lambda = \frac{1}{\sqrt{\hat{\beta}' S \hat{\beta}}} \hat{\beta}, \end{aligned}$$

where the second equality follows from Lemma 2, the third equality from the fact that $\lambda' S \lambda$ is constant due to the constraint, and the final equality from the definition of $\hat{\beta}$.

Case $\hat{\beta} = \mathbf{0}$

I will prove that $\hat{\lambda} = \mathbf{0}$ by contradiction. Suppose that $\hat{\lambda} \in \mathcal{C}^\circ$ is arbitrarily given. Then $\hat{\lambda}' S \hat{\lambda} - 2m' \hat{\lambda} > 0$, as $\hat{\beta} = \mathbf{0}$ is a unique minimizer. So $\frac{1}{2} > m' \hat{\lambda} / \hat{\lambda}' S \hat{\lambda}$, because we assumed that $\lambda' S \lambda > 0$, for all $\lambda \in \mathcal{C}^\circ$. As $\hat{\lambda} \in \mathcal{C}^\circ$,

we have that $\gamma\hat{\lambda} \in C^\circ$ for all $\gamma > 0$. Then $\lim_{\gamma \rightarrow 0^+} \mathbf{m}'(\gamma\hat{\lambda})/(\gamma\hat{\lambda})'\mathbf{S}(\gamma\hat{\lambda}) = \lim_{\gamma \rightarrow 0^+} \frac{1}{\gamma} \mathbf{m}'\hat{\lambda}/\hat{\lambda}'\mathbf{S}\hat{\lambda} < 1/2$, which implies $\mathbf{m}'\hat{\lambda} \leq 0$. This contradicts the assumption that $\hat{\lambda}$ is a unique minimizer. Hence, $\hat{\lambda} = \mathbf{0}$. $\blacksquare$

Proposition 10. Let $\mathbf{S} > 0$ be a diagonal matrix and let $\mathcal{D}$ be a scone. Suppose that $\hat{\lambda} = \arg \max_{\lambda \in \mathcal{S}^0 \cap \mathcal{D}} \mathbf{m}'\lambda$ and $\hat{\beta} = \arg \min_{\beta \in \mathcal{D}} \|\mathbf{S}^{-\frac{1}{2}}\mathbf{m} - \beta\|_2^2$ are unique optimizers. Then

$$\hat{\lambda} = \mathbf{S}^{-\frac{1}{2}}\hat{\beta}/\sqrt{\hat{\beta}'\hat{\beta}},$$

if $\hat{\beta} \neq \mathbf{0}$ and $\hat{\lambda} = \mathbf{0}$, otherwise.

Proof. I will consider two cases: $\hat{\beta} \neq \mathbf{0}$ and $\hat{\beta} = \mathbf{0}$.

Case $\hat{\beta} \neq \mathbf{0}$

Using the substitution $\beta = \mathbf{S}^{\frac{1}{2}}\lambda$ and $\hat{\mathbf{m}} = \mathbf{S}^{-\frac{1}{2}}\mathbf{m}$ yields

$$\hat{\lambda} = \arg \max_{\substack{\lambda \in \mathcal{D}_\Lambda \\ \lambda'\mathbf{S}\lambda=1}} \hat{\mu}'\lambda = \mathbf{S}^{-\frac{1}{2}} \arg \max_{\substack{\beta \in \mathcal{D}_\Lambda \\ \beta'\beta=1}} \hat{\mathbf{m}}'\beta,$$

where $\beta \in \mathcal{D}_\Lambda$ if $\lambda \in \mathcal{D}_\Lambda$, as $\mathbf{S}$ is diagonal and positive-definite. Define $\hat{\beta} = \arg \min_{\beta \in \mathcal{D}_\Lambda} \|\beta - \hat{\mathbf{m}}\|_2^2$. From Lemma 2 it follows that

$$\hat{\lambda} = \frac{1}{\sqrt{\hat{\beta}'\hat{\beta}}} \mathbf{S}^{-\frac{1}{2}} \arg \max_{\substack{\beta \in \mathcal{D}_\Lambda \\ \beta'\hat{\beta}=\hat{\beta}'\hat{\beta}}} \hat{\mathbf{m}}'\beta = \frac{1}{\sqrt{\hat{\beta}'\hat{\beta}}} \mathbf{S}^{-\frac{1}{2}} \arg \min_{\beta \in \mathcal{D}_\Lambda} \|\beta - \hat{\mathbf{m}}\|_2^2 = \frac{1}{\sqrt{\hat{\beta}'\hat{\beta}}} \mathbf{S}^{-\frac{1}{2}} \hat{\beta}.$$

Case $\hat{\beta} = \mathbf{0}$.

This case is analogous to the case that $\hat{\beta} = \mathbf{0}$ in the proof of Proposition 9. $\blacksquare$

Proposition 11. If $\mu \in C^\circ$, and Condition 5 and 6 hold then $P(\lim_{n \rightarrow \infty} p^{Tc}(\mathbf{X}) \leq \alpha) \leq \alpha$.

Proof. As much of this proof employs similar reasoning as used by Koning

and Bekker (2018), I employ similar notation and only include the steps up to the point that the reasoning is the same.

Let $\mu \in \mathcal{C}^\circ$ and draw $\mathcal{R}_n^M$ with equal probability from the subsets of $\mathcal{R}_n$ of size M that include I_n . Define

$$p_M^{T_C}(\mathbf{X}_n) = |\{\mathbf{R}_n \in \mathcal{R}_n^M \mid T_C(\mathbf{R}_n \mathbf{X}_n) \geq T_C(\mathbf{X}_n)\}|/M.$$

Here, I assume $T_C(\mathbf{R}_n^i \mathbf{X}_n) \neq T_C(\mathbf{R}_n^j \mathbf{X}_n)$, for all $\mathbf{R}_n^i, \mathbf{R}_n^j \in \mathcal{R}_n^M, i \neq j$, so that there exists a strict ordering. This can be done without loss of generality if we assume any ties are resolved with equal probability in either direction.

The idea of the proof is to show

$$P(\lim_{n \rightarrow \infty} p_M^{T_C}(\mathbf{X}_n) \geq \lim_{n \rightarrow \infty} p_M^g(E_n)) = 1, \quad (\text{A.4})$$

where $g(E_n) = T_C(\iota_n \mu' + E_n)$. As g is a function of E_n , $P(p_M^g(E_n) \leq \alpha) = \alpha$ for all n , by Condition 5. Therefore, by (A.4), $P(\lim_{n \rightarrow \infty} p_M^{T_C}(E_n) \leq \alpha) \leq \alpha$.

Writing out (A.4) yields

$$\begin{aligned} P\left(\lim_{n \rightarrow \infty} |\{\mathbf{R}_n \in \mathcal{R}_n^M \mid T_C(\mathbf{R}_n \mathbf{X}_n) \geq T_C(\mathbf{X}_n)\}|/M \geq \lim_{n \rightarrow \infty} |\{\mathbf{R}_n \in \mathcal{R}_n^M \mid g(\mathbf{R}_n E_n) \geq g(E_n)\}|/M\right) \\ = 1. \end{aligned}$$

Notice that $g(E_n) = T_C(\mathbf{X}_n)$, so in order to show (A.4) it remains to verify that $T_C(\mathbf{R}_n \mathbf{X}_n) \geq T_C(\iota_n \mu' + \mathbf{R}_n E_n)$ for large n .

First, notice that $T_C(\mathbf{R}_n \mathbf{X}_n) = T_C(\iota_n \mu' + E_n)$ if $\mathbf{R}_n = I_n$ or $\mu' \lambda = 0$. Hence, we can limit ourselves to the case that $\mathbf{R}_n \neq I_n$ and $\mu' \lambda < 0$, as $\mu' \lambda \leq 0$ for all $\mu \in \mathcal{C}^\circ$ and $\lambda \in \mathcal{C}$.

As $\mu' \lambda < 0$ and $\iota_n' \mathbf{R}_n \iota_n < n$, it then follows from Condition 6 and the same reasoning as used by Koning and Bekker (2019) that $n^{1/2} T_C(\mathbf{R}_n \mathbf{X}_n) > n^{1/2} T_C(\iota_n \mu' + \mathbf{R}_n E_n) + o_{a.s.}(0)$. This completes the proof. ■

Proposition 12. *If $\mu \in \mathcal{C}^*$ and Conditions 5 and 6 hold, then $P(\lim_{n \rightarrow \infty} p^{T_C}(\mathbf{X}) \leq \alpha) \geq \alpha$.*

Proof. This proof is analogous to the proof of Proposition 11, using the fact that $\mu' \lambda \geq 0$ for $\mu \in \mathcal{C}^*$ and $\lambda \in \mathcal{C}$, and flipping several inequalities.

Let $\mu \in \mathcal{C}^*$ and draw $\mathcal{R}_n^M$ with equal probability from the subsets of $\mathcal{R}_n$ of size M that include $\mathbf{I}_n$. The idea is now to instead show

$$P(\lim_{n \rightarrow \infty} p_M^{T_C}(\mathbf{X}_n) \leq \lim_{n \rightarrow \infty} p_M^g(E_n)) = 1, \quad (\text{A.5})$$

where $g(E_n) = T_C(\iota_n \mu' + E_n)$. As g is a function of E_n and g is injective, $P(p_M^g(E_n) \leq \alpha) = \alpha$ for all n , by Condition 5. Therefore, by (A.4), $P(\lim_{n \rightarrow \infty} p_M^{T_C}(E_n) \leq \alpha) \geq \alpha$.

Writing out (A.5) yields

$$P\left(\lim_{n \rightarrow \infty} |\{\mathbf{R}_n \in \mathcal{R}_n^M \mid T_C(\mathbf{R}_n \mathbf{X}_n) \geq T_C(\mathbf{X}_n)\}|/M \leq \lim_{n \rightarrow \infty} |\{\mathbf{R}_n \in \mathcal{R}_n^M \mid g(\mathbf{R}_n E_n) \geq g(E_n)\}|/M\right) = 1.$$

Notice that $g(E_n) = T_C(\mathbf{X}_n)$, so in order to show (A.4) it remains to verify that $T_C(\mathbf{R}_n \mathbf{X}_n) \leq T_C(\iota_n \mu' + \mathbf{R}_n E_n)$ for large n .

First, notice that $T_C(\mathbf{R}_n \mathbf{X}_n) = T_C(\iota_n \mu' + E_n)$ if $\mathbf{R}_n = \mathbf{I}_n$ or $\mu' \lambda = 0$. Hence, we can limit ourselves to the case that $\mathbf{R}_n \neq \mathbf{I}_n$ and $\mu' \lambda > 0$, as $\mu' \lambda \geq 0$ for all $\mu \in \mathcal{C}^*$ and $\lambda \in \mathcal{C}$.

As $\mu' \lambda > 0$ and $\iota_n' \mathbf{R}_n \iota_n < n$, it then follows from Condition 6 and the same reasoning as used by Koning and Bekker (2019) that $n^{1/2} T_C(\mathbf{R}_n \mathbf{X}_n) < n^{1/2} T_C(\iota_n \mu' + \mathbf{R}_n E_n) + o_{a.s.}(0)$. This completes the proof. $\blacksquare$

Proposition 13. *Let $\lambda' \hat{\Sigma} \lambda > 0$ for all $\lambda \in \mathcal{C}^\circ$. Let $\hat{\beta} = \arg \min_{\beta \in \mathcal{C}} \frac{1}{n} \|\iota_n - \mathbf{X} \beta\|_2^2$ and $\hat{\lambda} = \arg \max_{\lambda \in (\mathcal{S} \cap \mathcal{C}) \cup \{0\}} \mu' \lambda$ be unique optimizers. Then*

$$\hat{\lambda} = \hat{\beta} / \sqrt{\hat{\beta}' \hat{\Sigma} \hat{\beta}},$$

if $\hat{\beta} \neq 0$ and $\hat{\lambda} = 0$, otherwise.

Proof. Let $\widehat{G} = X'X/n$ denote the Gramian matrix. I will consider two cases.

Case $\widehat{\beta} \neq 0$

Writing out the ℓ_2 -norm and using Proposition 9 yields

$$\widehat{\beta} = \arg \min_{\beta \in \mathcal{C}} 1 - 2\widehat{\mu}'\beta + \widehat{\beta}'\widehat{G}\widehat{\beta} = \sqrt{\widehat{\beta}'\widehat{G}\widehat{\beta}} \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda'\widehat{G}\lambda=1}} \widehat{\mu}'\lambda,$$

where I use the fact that $\lambda'\widehat{\Sigma}\lambda := \lambda'\widehat{G}\lambda - (\widehat{\mu}'\lambda)^2 > 0$ for all $\lambda \in \mathcal{C}^\circ$, implies that $\lambda'\widehat{G}\lambda > 0$ for all $\lambda \in \mathcal{C}^\circ$. By Lemma 2

$$\widehat{\lambda} = \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda'\widehat{G}\lambda - (\widehat{\mu}'\lambda)^2=1}} \widehat{\mu}'\lambda = \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda'\widehat{G}\lambda=1+(\widehat{\mu}'\widehat{\lambda})^2}} \widehat{\mu}'\lambda = \sqrt{1 + (\widehat{\mu}'\widehat{\lambda})^2} \arg \max_{\substack{\lambda \in \mathcal{C} \\ \lambda'\widehat{G}\lambda=1}} \widehat{\mu}'\lambda.$$

Combining this and applying some straightforward algebra yields

$$\widehat{\lambda} = \sqrt{\frac{1+(\widehat{\mu}'\widehat{\lambda})^2}{\widehat{\beta}'\widehat{G}\widehat{\beta}}} \widehat{\beta} = \frac{1}{\sqrt{\widehat{\beta}'\widehat{\Sigma}\widehat{\beta}}} \widehat{\beta}.$$

Case $\widehat{\beta} = 0$

This case is analogous to the case that $\widehat{\beta} = 0$ in the proof of Proposition 9. ■

Table 4.2. Monte Carlo rejection rates if $H_0 : \mu = \mathbf{0}$ is true.

n	p	ρ	T_C			PE	Wald
			T_1/T_1^d	T_{20}	T_{20}^d		
30	100	0	.042	.038	.055	.786	.056
		.5	.043	.047	.052	.450	.045
		.7	.065	.048	.048	.270	.059
	300	0	.042	.043	.064	.895	.043
		.5	.037	.055	.061	.434	.049
		.7	.056	.041	.058	.274	.047
	500	0	.041	.052	.043	.950	.055
		.5	.053	.045	.047	.459	.062
		.7	.048	.043	.051	.261	.049
250	100	0	.054	.054	.066	.080	.058
		.5	.056	.043	.048	.058	.044
		.7	.044	.046	.045	.074	.056
	300	0	.041	.045	.049	.080	.039
		.5	.056	.053	.053	.077	.052
		.7	.046	.045	.048	.053	.053
	500	0	.054	.046	.050	.054	.055
		.5	.063	.041	.047	.060	.052
		.7	.048	.049	.050	.060	.049

The Monte Carlo rejection rates (between 0 and 1) for data generated under the null hypothesis $H_0 : \mu = \mathbf{0}$. Columns three to six contain the results for the reflection based T_C tests. The column headed by PE contains the rejection rates for the power enhancement test, and the final column contains the results for the Wald test that functions as the initial test in the power enhancement test.

Table 4.3. Monte Carlo rejection rates if $H_1^1 : \mu \in \mathcal{D}_1$ is true.

n	p	ρ	T_C			PE	Wald
			T_1/T_1^d	T_{20}	T_{20}^d		
100		0	.592	.068	.235	.985	.053
		.5	.697	.128	.071	.972	.059
		.7	.798	.176	.067	.959	.049
30	300	0	.485	.055	.209	.980	.053
		.5	.566	.107	.066	.955	.057
		.7	.672	.147	.052	.938	.045
	500	0	.404	.064	.158	.991	.056
		.5	.506	.091	.062	.949	.052
		.7	.616	.147	.066	.917	.051
100		0	.708	.324	.384	.665	.195
		.5	.737	.644	.076	.752	.455
		.7	.837	.921	.075	.876	.755
250	300	0	.577	.257	.296	.506	.054
		.5	.677	.453	.073	.473	.056
		.7	.767	.726	.059	.515	.053
	500	0	.535	.205	.276	.429	.061
		.5	.621	.384	.061	.434	.057
		.7	.722	.610	.074	.398	.046

The Monte Carlo rejection rates (between 0 and 1) for data generated under the alternative hypothesis with sparsity level $\|\mu\|_0 = 1$. Columns three to six contain the results for the reflection based T_C tests. The column headed by PE contains the rejection rates for the power enhancement test, and the final column contains the results for the Wald test that functions as the initial test in the power enhancement test.

Table 4.4. Monte Carlo rejection rates if $H_1^{20} : \mu \in \mathcal{D}_{20}$ is true.

n	p	ρ	T_C			PE	Wald
			T_1/T_1^d	T_{20}	T_{20}^d		
30	100	0	.278	.102	.675	.982	.045
		.5	.213	.118	.163	.796	.053
		.7	.197	.162	.124	.567	.048
	300	0	.130	.060	.351	.979	.051
		.5	.147	.072	.104	.701	.045
		.7	.141	.106	.092	.510	.053
	500	0	.155	.064	.237	.980	.060
		.5	.113	.085	.085	.698	.053
		.7	.136	.088	.090	.466	.050
250	100	0	.194	.337	.440	.387	.315
		.5	.164	.537	.109	.622	.607
		.7	.151	.779	.118	.880	.898
	300	0	.117	.152	.231	.094	.052
		.5	.108	.288	.072	.080	.051
		.7	.115	.473	.094	.089	.038
	500	0	.104	.107	.137	.088	.041
		.5	.087	.206	.083	.085	.036
		.7	.121	.325	.073	.051	.057

The Monte Carlo rejection rates (between 0 and 1) for data generated under the alternative hypothesis with sparsity level $\|\mu\|_0 = 20$. Columns three to six contain the results for the reflection based T_C tests. The column headed by PE contains the rejection rates for the power enhancement test, and the final column contains the results for the Wald test that functions as the initial test in the power enhancement test.

Chapter 5

Conclusion and Discussion

This dissertation contains a collection of several articles that are broadly linked through their focus on high-dimensional applications. The general context, as well as the context of each of the chapters is discussed in the introduction chapter. In this chapter, I briefly summarize the main conclusions of each chapter and mention some limitations and potential ideas for future research.

In Chapter 2 I discuss a new approach for obtaining sparse solutions in unit-sum regression by combining ℓ_0 - and ℓ_1 -regularization. This solves the issue that ℓ_1 -regularization may not be able to guarantee a unique solution in practical settings nor produce any sparsity, as the unit-sum restriction prevents shrinkage. Furthermore, compared to ℓ_0 -regularization, this combination offers the better predictive performance typically found with ℓ_1 -regularization. In addition, it offers sparser solutions than typically offered by ℓ_1 .

The main focus of the chapter is to compute such a regularized solution to unit-sum regression. The starting point of this approach is the construction of a first-order algorithm that finds a ‘good’ local solution, which is then fed as a warm start to a mixed-integer optimization solver. In simulation results we find that our approach performs favorably compared to separate ℓ_0 - or ℓ_1 -regularization. This is in line with the results found by Mazumder et al. (2017) for $\ell_0 + \ell_1$ -regularization in standard linear regression.

While the use of a warm start may give a considerable boost to a mixed-integer optimization solver, the solver remains the bottleneck. However, with further advances in computing power and solvers, this bottleneck may widen considerably in the future. In addition, Hazimeh and Mazumder (2018) have demonstrated that the MIO step may be removed altogether, by combining a first-order algorithm with local combinatorial optimization. This may be extended to unit-sum regression.

In Chapter 3 I discuss a new test for moment inequalities. The main focus of the paper is to increase the power of a test based on the maximum t statistic. The central idea is to expand the set t statistics over which a maximum is taken, with t statistics that are designed to be large if multiple moment inequalities are violated. By continuing to add more t statistics, this produces a one-sided generalization of Hotelling's T^2 statistic.

In the paper we propose inference based on a symmetry assumption on the error term, and provide a condition that guarantees size control in large samples. In numerical experiments it is shown to also control size in small samples. In addition, the numerical experiments demonstrate that the expansion of the set of t statistics can lead to considerable power gains against alternatives where multiple moment inequalities are violated. This suggests that tests based on the proposed statistic may indeed be better suitable for testing against such violations than a test based on a maximum t statistic.

While these results are promising, there are some limitations and points for future research. First, the numerical results suggest that it may be possible to find a condition for size control in small samples, or if the number of moment inequalities is allowed to grow with the sample size. In addition, it would be interesting to further explore the consequences of asymmetry in the error term. In particular: what are the consequences if the distribution is only approximately symmetric? It is found numerically that left skewness and right skewness distort the size in different directions. This suggests that the effect may be 'canceled out' in some sense if half the variables are left skewed and the other half right skewed. Finally, the number of elements in the used reflection group typically far exceeds the amount

necessary (or even computationally feasible) to use for inference. This has lead people to study subgroups, to replace the reflection group and thereby weaken the symmetry assumption. However, I conjecture that for the type of statistic considered in this chapter, there may exist subgroups that yield a more powerful test in finite samples.

In Chapter 4, I extend the statistic of Chapter 3 to arbitrary (not necessarily convex) cones. I analyze properties of this conic statistic, with a focus on its existence and computation, and apply it to hypothesis testing.

The main focus of the chapter is to find a substitute for quadratic statistics, which suffer from power loss against specific alternatives and may be poorly defined in high-dimensional settings. The conic statistic is argued to be a good candidate, as its existence condition is weaker and it is ‘large’ under a specific alternative of interest.

In the chapter, its use is further explored by constructing a randomization test, where a condition from Chapter 3 is generalized to guarantee asymptotic size control. Furthermore, numerical results show that the test performs well compared to other existing tests. Together, this demonstrates that tests based on a conic statistic may indeed be a good substitute in situations where quadratic statistics fail.

The main limitation of this chapter is that setting the cone equal to the specific alternative interest is mostly a practical choice. In particular, there may exist other cones that yield a test with higher power against the specific alternative. This warrants further research on the link between the choice of cone and specific alternative. Furthermore, it may be possible to weaken the condition that guarantees size control to somehow permit (growing) sequences of cones which could be used to demonstrate size control under a growing dimensionality.

Bibliography

- Abadie, A., Diamond, A., and Hainmueller, J., 2010. Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association*, 105 (490):493–505.
- Allen, R., 2018. Testing moment inequalities: Selection versus recentering. *Economics Letters*, 162:124–126.
- Andrews, D. W. and Barwick, P. J., 2012. Inference for parameters defined by moment inequalities: A recommended moment selection procedure. *Econometrica*, 80(6):2805–2826.
- Andrews, D. W. and Guggenberger, P., 2009. Validity of subsampling and plug-in asymptotic inference for parameters defined by moment inequalities. *Econometric Theory*, 25(3):669–709.
- Andrews, D. W. and Shi, X., 2013. Inference based on conditional moment inequalities. *Econometrica*, 81(2):609–666.
- Armstrong, T. B., 2014. Weighted ks statistics for inference on conditional moment inequalities. *Journal of Econometrics*, 181(2):92–116.
- Armstrong, T. B., 2015. Asymptotically exact inference in conditional moment inequality models. *Journal of Econometrics*, 186(1):51–65.
- Armstrong, T. B. and Chan, H. P., 2016. Multiscale adaptive inference on conditional moment inequalities. *Journal of Econometrics*, 194(1):24–43.
- Beasley, J. E., Meade, N., and Chang, T.-J., 2003. An evolutionary heuristic for the index tracking problem. *European Journal of Operational Research*, 148(3):621–643.
- Bekker, P. A., 2002. Exact inference for the linear model with groupwise heteroscedastic spherical disturbances. *Journal of Econometrics*, 111(2):285 – 302. ISSN 0304-4076.
- Bekker, P. A. and Lawford, S., 2008. Symmetry-based inference in an instrumental variable setting. *Journal of econometrics*, 142(1):28–49.

- Bertsimas, D., King, A., Mazumder, R., et al., 2016. Best subset selection via a modern optimization lens. *The Annals of Statistics*, 44(2):813–852.
- Brodie, J., Daubechies, I., De Mol, C., Giannone, D., and Loris, I., 2009. Sparse and stable markowitz portfolios. *Proceedings of the National Academy of Sciences*, 106(30):12267–12272.
- Bugni, F. A., Caner, M., Kock, A. B., and Lahiri, S., 2016. Inference in partially identified models with many moment inequalities using lasso. *arXiv preprint arXiv:1604.02309*.
- Canakgoz, N. A. and Beasley, J. E., 2009. Mixed-integer programming approaches for index tracking and enhanced indexation. *European Journal of Operational Research*, 196(1):384–399.
- Canay, I. A., 2010. El inference for partially identified models: Large deviations optimality and bootstrap validity. *Journal of Econometrics*, 156(2): 408–425.
- Chen, L. S., Paul, D., Prentice, R. L., and Wang, P., 2011. A regularized hotellings t^2 test for pathway analysis in proteomic studies. *Journal of the American Statistical Association*, 106(496):1345–1360.
- Chernozhukov, V., Hong, H., and Tamer, E., 2007. Estimation and confidence regions for parameter sets in econometric models 1. *Econometrica*, 75(5): 1243–1284.
- Chernozhukov, V., Lee, S., and Rosen, A. M., 2013. Intersection bounds: Estimation and inference. *Econometrica*, 81(2):667–737.
- Chernozhukov, V., Chetverikov, D., and Kato, K., 2019a. Inference on causal and structural parameters using many moment inequalities. *The Review of Economic Studies*, 86(5):1867–1900.
- Chernozhukov, V., Chetverikov, D., Kato, K., and Koike, Y., 2019b. Improved central limit theorem and bootstrap approximations in high dimensions. *arXiv preprint arXiv:1912.10529*.
- Ciliberto, F. and Tamer, E., 2009. Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791–1828.

- Coleman, T. F., Li, Y., and Henniger, J., 2006. Minimizing tracking error while restricting the number of assets. *Journal of Risk*, 8(4):33.
- Cutler, A. and Breiman, L., 1994. Archetypal analysis. *Technometrics*, 36(4): 338–347.
- DeMiguel, V., Garlappi, L., Nogales, F. J., and Uppal, R., 2009. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55(5):798–812.
- Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al., 2004. Least angle regression. *The Annals of Statistics*, 32(2):407–499.
- Fan, J., 1996. Test of significance based on wavelet thresholding and neyman’s truncation. *Journal of the American Statistical Association*, 91(434): 674–688.
- Fan, J., Liao, Y., and Yao, J., 2015. Power enhancement in high-dimensional cross-sectional tests. *Econometrica*, 83(4):1497–1541.
- Fastrich, B., Paterlini, S., and Winker, P., 2014. Cardinality versus q-norm constraints for index tracking. *Quantitative Finance*, 14(11):2019–2032.
- Fisher, R. A., 1935. *The Design of Experiments*. Oliver and Boyd, London.
- Hastie, T., Tibshirani, R., and Friedman, J., 2009. *The elements of statistical learning: data mining, inference and prediction*. Springer, 2 edition. URL <http://www-stat.stanford.edu/~tibs/ElemStatLearn/>.
- Hastie, T., Tibshirani, R., and Tibshirani, R. J., 2017. Extended comparisons of best subset selection, forward stepwise selection, and the lasso. *arXiv preprint arXiv:1707.08692*.
- Hazimeh, H. and Mazumder, R., 2018. Fast best subset selection: Coordinate descent and local combinatorial optimization algorithms. *arXiv preprint arXiv:1803.01454*.
- Hiriart-Urruty, J.-B. and Seeger, A., 2010. A variational approach to copositive matrices. *SIAM review*, 52(4):593–629.

- Kock, A. B. and Preinerstorfer, D., 2019. Power in high-dimensional testing problems. *Econometrica*, 87(3):1055–1069.
- Koning, N. W., 2019. Testing against specific alternatives with conic statistics. *arXiv preprint arXiv:1907.05077*.
- Koning, N. W. and Bekker, P. A., 2018. Sparse unit-sum regression. *arXiv preprint arXiv:1907.04620*.
- Koning, N. W. and Bekker, P. A., 2019. Exact testing of many moment inequalities against multiple violations. *arXiv preprint arXiv:1904.12775*.
- Lawson, C. L. and Hanson, R. J., 1995. *Solving least squares problems*, volume 15. Siam.
- Lee, S., Song, K., and Whang, Y.-J., 2013. Testing functional inequalities. *Journal of Econometrics*, 172(1):14–32.
- Lee, S., Song, K., and Whang, Y.-J., 2018. Testing for a general class of functional inequalities. *Econometric Theory*, 34(5):1018–1064.
- Lehmann, E. L. and Romano, J. P., 2006. *Testing statistical hypotheses*. Springer Science & Business Media.
- Maritz, J. S., 1995. *Distribution-free statistical methods*, volume 17. CRC Press.
- Mazumder, R., Radchenko, P., and Dedieu, A., 2017. Subset selection with shrinkage: Sparse linear modeling when the snr is low. *arXiv preprint arXiv:1708.03288*.
- Meinshausen, N. et al., 2013. Sign-constrained least squares estimation for high-dimensional regression. *Electronic Journal of Statistics*, 7:1607–1631.
- Natarajan, B. K., 1995. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234.
- Nesterov, Y., 2013. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media.

- Pakes, A., 2010. Alternative models for moment inequalities. *Econometrica*, 78(6):1783–1822. doi: 10.3982/ECTA7944. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA7944>.
- Romano, J. P., 1990. On the behavior of randomization tests without a group invariance assumption. *Journal of the American Statistical Association*, 85(411):686–692.
- Romano, J. P. and Shaikh, A. M., 2008. Inference for identifiable parameters in partially identified econometric models. *Journal of Statistical Planning and Inference*, 138(9):2786–2807.
- Rosen, A. M., 2008. Confidence sets for partially identified parameters that satisfy a finite number of moment inequalities. *Journal of Econometrics*, 146(1):107–117.
- Slawski, M., Hein, M., et al., 2013. Non-negative least squares for high-dimensional linear models: Consistency and sparse recovery without regularization. *Electronic Journal of Statistics*, 7:3004–3056.
- Srivastava, M. S. and Du, M., 2008. A test for the mean vector with fewer observations than the dimension. *Journal of Multivariate Analysis*, 99(3): 386–402.
- Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.
- Timmermann, A., 2006. Forecast combinations. *Handbook of economic forecasting*, 1:135–196.
- Van de Geer, S. The deterministic lasso. JSM Proceedings, 140. Am. Statistic. Assoc., Alexandria, VA., 2007.
- Van De Geer, S. A., Bühlmann, P., et al., 2009. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392.
- Vock, M. “One-sided” Statistical Inference for a Multivariate Location Parameter. PhD thesis, University of Bern, 2007.

- Wei, Y., Wainwright, M. J., Guntuboyina, A., et al., 2019. The geometry of hypothesis testing over convex cones: Generalized likelihood ratio tests and minimax radii. *The Annals of Statistics*, 47(2):994–1024.
- Welch, B. L., 1947. The generalization of Student's problem when several different population variances are involved. *Biometrika*, 34(1-2):28–35.
- Zhong, P.-S., Chen, S. X., Xu, M., et al., 2013. Tests alternative to higher criticism for high-dimensional means under sparsity and column-wise dependence. *The Annals of Statistics*, 41(6):2820–2851.

Samenvatting

In het afgelopen decennium heeft er een enorme groei plaatsgevonden in de hoeveelheid data die beschikbaar is voor economisch onderzoek. Dit betreft in veel gevallen data op een gedetailleerder niveau dan voorheen beschikbaar was. Een bijzondere eigenschap van dergelijke data is de aanwezigheid van zeer veel variabelen, wat leidt tot modellen waarin het aantal parameters mogelijk substantieel groter is dan het aantal observaties ('hoog-dimensioneel'). Klassieke econometrische technieken zijn hiervoor vaak ongeschikt, omdat ze over het algemeen juist ontwikkeld zijn voor situaties waarin het aantal observaties groot is ten opzichte van het aantal parameters. De overkoepelende bijdrage van dit proefschrift betreft de ontwikkeling van nieuwe technieken die geschikt zijn voor hoog-dimensionele data.

Het onderwerp van Hoofdstuk 2 is eenheidssomregressie. Eenheidssomregressie betreft lineaire regressie onder de restrictie dat de parameters optellen tot 1. Dit heeft vele toepassingen binnen de economische wetenschap. Prominente voorbeelden hiervan zijn te vinden in portfolio-optimalisatie, synthetisch gecontroleerde omstandigheden, en voorspellingscombinatie.

In het hoofdstuk ligt de focus op de ℓ_1 -ijlte van de parametervector. Dat wil zeggen dat de parametervector voornamelijk bevolkt is door nullen. In standaard lineaire regressie is het gebruikelijk om, onder verdenking van ℓ_1 -ijlte, zogenoemde ℓ_1 -regularisatie toe te passen, wegens praktische alsook prestatiegerichte overwegingen. Hierbij wordt de som van de absolute waarde van de parameters beperkt. In eenheidssomregressie hindert de eenheidssomrestrictie echter de mate waarin de parameters kunnen worden beperkt. De som van de absolute waarde kan immers niet gedwongen worden om kleiner te zijn dan 1, als de som 1 moet zijn. Hierdoor wordt de effectivi-

teit van ℓ_1 -regularisatie belemmerd. In het hoofdstuk stellen wij voor om naast ℓ_1 -regularisatie ook ℓ_0 -regularisatie toe te voegen, waar het aantal parameters dat niet gelijk is aan nul expliciet wordt gerespecteerd. Vervolgens demonstreren we hoe een oplossing kan worden berekend en bestuderen we de eigenschappen middels simulaties.

In Hoofdstuk 3 wordt het toetsen van momentongelijkheden besproken. Dergelijke momentongelijkheden komen veel voor in de economische wetenschap. Ze ontstaan bijvoorbeeld bij het maken van discrete keuzes, waarbij de onderzoeker slechts de kwalitatieve observatie kan maken dat de keuze door een beslisser werd verkozen boven een selectie aan andere mogelijkheden. Deze ongelijkheden worden gebruikt om onderliggende eigenschappen van de beslisser te schatten, die gerelateerd zijn aan het keuzeproces. Vervolgens kunnen de ongelijkheden worden getoetst om betrouwbaarheidsintervallen te verkrijgen voor de schattingen.

In sommige toepassingen kan het aantal ongelijkheden erg groot zijn. Hiervoor hebben Chernozhukov et al. (2019a) toetsen ontwikkeld op basis van het grootste (gestudentiseerde) steekproefmoment. In het hoofdstuk ontwikkelen wij een toets met een groter onderscheidend vermogen tegen alternatieven waaronder meerdere momentongelijkheden zijn overschreden. Daarvoor generaliseren wij de toets van Chernozhukov et al. (2019a) door niet slechts de steekproefmomenten te gebruiken die overeenkomen met de momentongelijkheden, maar ook sommige (of alle) steekproefmomenten die overeenkomen met momentongelijkheden die worden gempliceerd door de originele verzameling van momentongelijkheden. Verder creëren wij een kritieke waarde op basis van een symmetrieaanname, en geven wij een conditie waaronder dit leidt tot een juiste toetsingsgrootte. In simulaties vinden wij dat dit inderdaad leidt tot een toets met een groter onderscheidend vermogen tegen alternatieven waaronder meerdere momentongelijkheden zijn overschreden.

In het laatste hoofdstuk worden de ideeën uit Hoofdstuk 3 gegeneraliseerd naar generieke kegels, wat leidt tot zogenoemde kegelstatistieken. Deze kegelstatistieken worden vergeleken met kwadratische statistieken (zoals een Wald-type statistiek), waarvan bekend is dat ze een relatief laag onderscheidend vermogen hebben tegen specifieke alternatieven. Een ander

probleem met kwadratische statistieken is dat ze vaak de inverse van een steekproefcovariantiematrix bevatten, die niet bestaat in hoogdimensionele toepassingen. De voordelen van kegelstatistieken zijn dat ze zwakkere bestaanscriteria hebben, en met behulp van de gekozen kegel het onderscheidend vermogen kunnen richten op specifieke alternatieven. Voor het standaard locatiemodel wordt op een vergelijkbare manier als in Hoofdstuk 3 een conditie afgeleid waaronder een juiste toetsingsgrootte kan worden gegarandeerd. Uit simulaties blijkt dat deze toets gunstig presteert in vergelijking met toetsen die gebaseerd zijn op een kwadratische statistiek, en de onderscheidendvermogenverhogende techniek van Fan et al. (2015).