Published in final edited form as:

SIAM Rev Soc Ind Appl Math. 2015 ; 57(2): 167–197. doi:10.1137/130921179.

Linear Models Based on Noisy Data and the Frisch Scheme*

Lipeng Ning[†], Tryphon T. Georgiou[‡], Allen Tannenbaum[§], and Stephen P. Boyd[¶]

Lipeng Ning: lning@bwh.harvard.edu; Tryphon T. Georgiou: tryphon@umn.edu; Allen Tannenbaum: allen.tannenbaum@stonybrook.edu; Stephen P. Boyd: boyd@stanford.edu

[†]Brigham and Women's Hospital, Harvard Medical School, Boston, MA 02115

[‡]Department of Electrical & Computer Engineering, University of Minnesota, Minneapolis, MN 55455

[§]Departments of Computer Science and Applied Mathematics, Stony Brook University, Stony Brook, NY 11794

[¶]Department of Electrical Engineering, Stanford University, Stanford, CA 94305

Abstract

We address the problem of identifying linear relations among variables based on noisy measurements. This is a central question in the search for structure in large data sets. Often a key assumption is that measurement errors in each variable are independent. This basic formulation has its roots in the work of Charles Spearman in 1904 and of Ragnar Frisch in the 1930s. Various topics such as errors-in-variables, factor analysis, and instrumental variables all refer to alternative viewpoints on this problem and on ways to account for the anticipated way that noise enters the data. In the present paper we begin by describing certain fundamental contributions by the founders of the field and provide alternative modern proofs to certain key results. We then go on to consider a modern viewpoint and novel numerical techniques to the problem. The central theme is expressed by the Frisch–Kalman dictum, which calls for identifying a noise contribution that allows a maximal number of simultaneous linear relations among the noise-free variables—a rank minimization problem. In the years since Frisch's original formulation, there have been several insights, including trace minimization as a convenient heuristic to replace rank minimization. We discuss convex relaxations and theoretical bounds on the rank that, when met, provide guarantees for global optimality. A complementary point of view to this minimum-rank dictum is presented in which models are sought leading to a uniformly optimal quadratic estimation error for the error-free variables. Points of contact between these formalisms are discussed, and alternative regularization schemes are presented.

Keywords

linear models; factor analysis; identification

*This work was supported in part by grants from NSF, NIH, AFOSR, ONR, and MDA. This work is part of the National Alliance for Medical Image Computing (NA-MIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, grant U54 EB005149. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>. Finally, this project was supported by grants from the National Center for Research Resources (P41-RR-013218) and the National Institute of Biomedical Imaging and Bioengineering (P41-EB-015902) of the National Institutes of Health.

1. Introduction

The standard paradigm in modeling is to postulate that measured quantities contain a contribution of “accidental deviation” [72] from the otherwise “uniformities” that characterize an underlying law. Therefore, a key issue when identifying dependencies between variables is to decide on how to account for the contribution of noise in the data. Various assumptions on the structure of noise and of the possible dependencies lead to a number of corresponding methodologies.

The purpose of the present paper is to present from a modern computational perspective the basic paradigm of identifying linear relations under the assumption that noise components are independent and to discuss variants and generalizations. The independence assumption underlies the errors-in-variables model [25, 52, 67, 21] and factor analysis [6, 5, 55, 39, 44, 68] and has a century-old history [35, 64, 53]; see also [45, 46, 58, 78, 36, 71, 4, 30, 12]. Accordingly, given the long history of this problem, the paper aims to provide a tutorial flavor.

The precise formulation of the central question has its roots in the work of Ragnar Frisch in the 1930s. The key assumption is that the noise components are independent of the underlying variables and that they are mutually independent as well, while the “noise-free” variables satisfy a set of linear relations. The data are typically abstracted in the form of an (estimated) covariance matrix. The basic modeling premise dictates a decomposition of a covariance matrix obtained from the data into the sum of a covariance matrix corresponding to the “noise free” variables and a diagonal one corresponding to noise. However, in addition, since there is no uniqueness in such a decomposition and therefore several alternative sets of linear relations can be consistent, a maximal set of simultaneous dependencies is sought as a means to limit uncertainty and to provide canonical models [45, 46]. These hypotheses define the Frisch–Kalman scheme and lead to a (nonconvex) rank-minimization problem where a given covariance matrix is sought to be decomposed into a sum of two covariance matrices, a singular one of minimal rank and another diagonal corresponding to noise.

In spite of extensive work over half a century by many researchers, exact solutions for such a decomposition have not been forthcoming, with the exception of a result due to Reiersøl [64]; see also [45, 52]. This result pertains to the special case where the number of simultaneous linear relations is equal to 1 and is no greater. Interestingly, this case as well as its complement are generic in that small changes in the data may leave them unaffected. For the case where there are no two simultaneous linear relations possible, necessary and sufficient conditions on the data covariance matrix were provided by Reiersøl over half a century ago. Further, for this case, a detailed parametrization of all solutions has been obtained. To date no other case is known that admits such precise necessary and sufficient conditions and closed-form parametrization of solutions.

In view of the fundamental nature of this problem and the scarcity of exact solutions, several computational approaches have emerged over the years [68, 69, 22, 8, 79, 9, 59]. A key trend

has been to seek numerical decompositions that employ the trace as a surrogate for the rank [28]. In parallel, attempts have been made to determine computable a priori lower bounds on the rank of the “noise-free” component [37, 68, 69]. The rank of the “noise-free” covariance matrix represents the number of “factors” in factor analysis, and such bounds on the minimum rank, when met, serve as “certificates” that guarantee optimality. That is, when numerical decomposition produces a covariance for the noise-free component that attains a lower bound of the rank, this rank represents the minimum number of independent “factors” while the co-rank (nullity) represents the maximal number of linear relations consistent with the assumptions of the Frisch scheme.

Relaxations of modeling assumptions and constraints have been considered throughout the development of the subject so as to account for statistical errors that are inherent in empirical covariance estimates and, also, in order to provide insights into potential underlying algebraic structure in the data [68]. Related in spirit are recent ideas on minimax robust analysis [48, 77, 47, 51, 27] and the use of error loss functions together with regularization terms that promote low-rank decompositions [28, 14, 59, 65, 16, 1]. Our development will draw on some of these ideas and provide points of contact.

Our exposition begins by motivating Frisch’s error-in-variables paradigm while explaining the origin of the modeling assumptions (section 3). In section 4 we precisely define the Frisch problem. We then analyze the special case where the maximal number of linear dependencies is exactly equal to one (the Reiersøl case). In the same section (section 4), we also define and develop in parallel a problem due to Shapiro [68]. Shapiro’s problem differs from Frisch’s in that the requirement for the diagonal summand to be positive semidefinite is relaxed in order to gain insight into potential algebraic relations imposed by the off-diagonal elements of the data covariance matrix. The Frisch and the Shapiro problems are quite similar, and their parallel presentation helps underscore insights into the structure of the respective solutions. Thus, section 4 details the theory for both problems in parallel. More specifically, in this section we discuss the well-posedness of the two problems, we present a new geometric principle that is used to derive and explain the analytic conditions for the Reiersøl case (section 4.1), we develop a certain dual viewpoint for both problems (section 4.2), we derive necessary and sufficient conditions as well as a parametrization of solutions for the special Reiersøl case (sections 4.3, 4.4, and 4.5), and lastly we discuss a generalization of Frisch’s problem for the case of complex-valued covariances and provide partial results for an analogue of the Reiersøl case (section 4.6).

In section 5 we turn to convex relaxation methods for the problems of Frisch and Shapiro. We revisit the trace minimization heuristic for rank-minimization problems. Since the rank remains invariant under scaling of rows and columns, it is natural to consider weighted-trace minimization; this is explained in section 5.1. In section 6 we provide lower bounds for the minimum rank which provide potential certificates for optimality of solutions which can be computed by convex optimization.

As already indicated, the Frisch scheme centers around a certain decomposition of estimated covariance statistics. In section 7 we consider what such a decomposition implies about the data and, in particular, a conformable decomposition of the data into signal plus noise

directly in the time domain. More specifically, we show that to any Frisch decomposition there corresponds a family of consistent decompositions of the data into signal plus noise. This family corresponds to a matrix ball and, under a Gaussian assumption, the “uncertainty matrix radius” coincides with the variance of the conditional expectation of the “noise-free” variables (section 8). This analysis is brought in as a prelude to certain alternative relaxations of the Frisch paradigm; it motivates the related problem of minimizing uniformly the variance of conditional expectation as an alternative to the Frisch rationale (section 8.1). Thus, instead of a maximal number of simultaneous linear relations, for certain types of applications one may seek a uniformly optimal estimator for the unobserved data under the independence assumption of the Frisch scheme. Thereby, the optimal estimator can be obtained as a solution to a minmax optimization problem. Points of contact with certain robust techniques in signal analysis [48, 77, 47, 51, 27] and with trace regularization are discussed in section 8.2.

In section 9 we consider ways that one may account for the effect of statistical estimation errors in the sample covariance matrices. To this end, it is natural to relax the requirement for an exact decomposition of sample statistics. Among the many possible alternatives, we chose to draw a connection between approximating Gaussian probability density functions in the 2-Wasserstein metric and a distance between the corresponding covariance matrices. This can be conveniently expressed as a semidefinite program. Thus, trace minimization as well as accounting for statistical uncertainty and quadratic performance can all be expressed in a similar format employing semidefinite programming. An example is discussed in section 9 that helps underscore the scope and potential limitations of the Frisch scheme and of the formalism in this paper. The aim of the paper has been to thread together historical as well as modern ideas and techniques which have grown around basic ideas in modeling that have roots in the work of Charles Spearman and Ragnar Frisch. Thus, the main theme of the Frisch scheme and alternative viewpoints are reiterated in the concluding remarks (section 10).

2. Notation

$\mathcal{R}(\cdot), \mathcal{N}(\cdot)$	range space, null space
$\Pi_{\mathcal{X}}$	orthogonal projection onto $\mathcal{X}$
> 0 (≥ 0)	positive definite (resp., positive semidefinite)
$\mathbf{S}_n$	$= \{M \mid M \in \mathbb{R}^{n \times n}, M = M^T\}$
$\mathbf{S}_{n,+}$	$= \{M \mid M \in \mathbf{S}_n, M \geq 0\}$
$\mathbf{H}_n$	$= \{M \mid M \in \mathbb{C}^{n \times n}, M = M^H\}$
$\mathbf{H}_{n,+}$	$= \{M \mid M \in \mathbf{H}_n, M \geq 0\}$
$[\cdot]_{k\ell}, ([\cdot]_k)$	(k, ℓ) th entry (resp., k th entry)
$ M $	determinant of $M \in \mathbb{R}^{n \times n}$
$n_+(\cdot)$	number of positive eigenvalues
$\text{diag} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^n : M \mapsto d$	where $[d]_i = [M]_{ii}$ for $i = 1, \dots, n$
$\text{diag}^* : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n} : d \mapsto D$	where D is diagonal and $[D]_{ii} = [d]_i$ for $i = 1, \dots, n$
$M \succ_e 0$ ($\succeq_e 0, \prec_e 0, \preceq_e 0$)	the off-diagonal entries are > 0 (resp., $\geq 0, < 0, \leq 0$), or can be made so by changing the signs of selected rows and corresponding columns

3. Data and Basic Assumptions

Consider a Gaussian vector $\mathbf{x}$ taking values in $\mathbb{R}^{n \times 1}$ having zero mean and covariance matrix Σ . We assume that it represents an additive mixture of a Gaussian “noise-free” vector $\hat{\mathbf{x}}$ and a “noise component” $\tilde{\mathbf{x}}$; thus

$$\mathbf{x} = \hat{\mathbf{x}} + \tilde{\mathbf{x}}. \quad (3.1)$$

The entries of $\tilde{\mathbf{x}}$ are assumed independent of one another and independent of the entries of $\hat{\mathbf{x}}$ with both vectors having zero mean and covariance matrices $\hat{\Sigma}$ and $\tilde{\Sigma}$, respectively. Thus,

$$\mathcal{E}(\tilde{\mathbf{x}}\tilde{\mathbf{x}}') =: \tilde{\Sigma} \text{ is diagonal}, \quad (3.2a)$$

$$\mathcal{E}(\hat{\mathbf{x}}\tilde{\mathbf{x}}') = 0, \quad (3.2b)$$

while

$$\Sigma = \hat{\Sigma} + \tilde{\Sigma}. \quad (3.2c)$$

Throughout $\mathcal{E}(\cdot)$ denotes the expectation operation and 0 denotes the zero vector/matrix of appropriate size. The entries of $\hat{\mathbf{x}}$, which are thought of as “noise-free” variables, are assumed to satisfy a set of q simultaneous linear relations. The coefficients of these relations form the columns of a matrix $M \in \mathbb{R}^{n \times q}$ having $\text{rank}(M) = q > 0$ and satisfying $M\hat{\mathbf{x}} = 0$. From an applications standpoint, *the modeling problem is to infer q as well as the parameters specifying such linear relations between the entries of $\hat{\mathbf{x}}$ from samples of the random vector $\mathbf{x}$.*

An equivalent statement to the existence of q independent relations between the entries of $\hat{\mathbf{x}}$ is that the covariance matrix $\mathcal{E}(\hat{\mathbf{x}}\hat{\mathbf{x}}') =: \hat{\Sigma}$ has

$$\text{rank}(\hat{\Sigma}) = n - q, \quad (3.2d)$$

in which case $\hat{\Sigma}M = 0$. The “noise-free” component may then be expressed as a linear combination of $n - q$ latent variables referred to as “factors,” which are the entries of a random vector $\mathbf{v}$, and written in the form $\hat{\mathbf{x}} = F\mathbf{v}$; the matrix $F \in \mathbb{R}^{n \times (n-q)}$ is often referred to as the *factor loading*. The broader subject of factor analysis is an especially timely subject due to the growing interest in explaining high-dimensional data [54, 18, 12]. Comprehensive accounts on the foundation of factor analysis and variations on this basic scheme can be found in [6, 5, 8, 9], while recent trends to use factor analysis in time series are seen in [40, 67, 30, 20, 32, 31].

Statistics of a physical variable are estimated based on finite observation records. To this end, consider a sequence

$$x_t \in \mathbb{R}^{n \times 1}, \quad t=1, \dots, T,$$

of independent measurements (realizations) of $\mathbf{x}$ and, likewise, let $\hat{x}_t$ and $\tilde{x}_t$ represent the corresponding values of the noise-free variable and noise components, respectively. For simplicity, we assume that the mean of all variables is zero. Denote by

$$X = [x_1 \ x_2 \ \dots \ x_T] \in \mathbb{R}^{n \times T}$$

the matrix of observations of $\mathbf{x}$ and similarly denote by $\hat{X}$ and $\tilde{X}$ the corresponding matrices of the noise-free and noise entries, respectively. Data for identifying relations among the noise-free variables are typically limited to the observation matrix X and, neglecting a scaling factor of $1/T$, the data are typically abstracted in the form of a sample covariance XX' . For the most part we will assume that sample covariances are accurate approximations of true covariance matrices, and hence the modeling assumptions amount to

$$\tilde{X}\tilde{X}' \simeq \text{diagonal}, \quad (3.3a)$$

$$\hat{X}\tilde{X}' \simeq 0, \quad (3.3b)$$

$$\text{rank}(\hat{X}) = n - q \quad (3.3c)$$

since $\tilde{M}\hat{X} = 0$.

The number of possible linear relations among the noise-free variables and the corresponding coefficient matrix are to be determined from either Σ or X . If no prior assumption on the structure of the noise is available, then singular value decomposition of Σ , also X , often provides a satisfactory decomposition $X = \hat{X} + \tilde{X}$ grouped into components that correspond to the top and bottom singular values of X (see e.g., [42]). In this, however, the “noise” component fails to satisfy the independence assumption (3.3a). It is often the case that more is known about the noise and the independence assumption in the Frisch model represents one of the earliest and most appealing paradigms.

Thus, the need to decompose data into signal and noise, relying on a structural prior on the noise statistics, motivates two basic mathematical problems that are formulated in section 4—the Frisch and Shapiro problems. These address decompositions of the covariance matrix Σ as in (3.2c) where the summands abide by the structural assumptions (3.2a)–(3.2b).

An alternative line of reasoning, rooted in optimal estimation, is presented in section 8. This rationale motivates a complementary viewpoint aimed at addressing estimation problems under modeling uncertainty; we will return to this in section 8.

4. The Problems of Frisch and Shapiro

We begin by formulating the Frisch problem. This pertains to the decomposition of a covariance matrix Σ in a way that is consistent with the assumptions of section 3. These assumptions are somewhat stringent in that, in practice, Σ is an empirical sample covariance. This fact motivates relaxing assumptions (3.2a)–(3.2d) in various ways. In particular, relaxation of the constraint $\tilde{\Sigma} \succeq 0$ leads to a problem that was studied by Shapiro and others and is formulated below as well. There is a strong similarity between the two problems, and the parallel treatment below underscores insights that can be gained.

Problem 1 (the Frisch problem)—Given $\Sigma \in \mathbf{S}_{n,+}$, determine

$$\text{mr}_+(\Sigma) := \min \{ \text{rank}(\hat{\Sigma}) \mid \Sigma = \tilde{\Sigma} + \hat{\Sigma}, \hat{\Sigma} \succeq 0, \tilde{\Sigma} \text{ is diagonal} \}. \quad (4.1)$$

Problem 2 (the Shapiro problem)—Given $\Sigma \in \mathbf{S}_{n,+}$, determine

$$\text{mr}(\Sigma) := \min \{ \text{rank}(\hat{\Sigma}) \mid \Sigma = \tilde{\Sigma} + \hat{\Sigma}, \hat{\Sigma} \succeq 0, \tilde{\Sigma} \text{ is diagonal} \}. \quad (4.2)$$

Frisch's problem has been studied by several researchers for over fifty years (see [46, 58, 79, 67] and the references therein). On the other hand, Shapiro [68] and others [56, 38] investigated the relaxed version stated as Problem 2, by removing the requirement that $\tilde{\Sigma} \succeq 0$, in an attempt to gain understanding of the algebraic constraints imposed by the off-diagonal elements of Σ on the decomposition. We will develop *Shapiro's problem* in parallel with *Frisch's problem* because of the strong parallels between the two and the insights this analysis provides.

Either problem requires computing the minimum rank of a partially specified positive semidefinite matrix. While the computation of the minimum rank for such partially specified matrices has been recently investigated in the literature (see, e.g., [34, 63]), the positive semidefinite requirement brings in an additional element that impacts the structure of solutions. Below, we will develop in parallel the theory of the problems of Frisch and Shapiro. This parallel treatment allows insight into the effect of positivity of the summands and underscores important differences between the two. Furthermore, it gives us an opportunity to provide a modern exposition of classical results of the subject.

We first discuss some properties of $\text{mr}_+(\cdot)$ and $\text{mr}(\cdot)$. We refer to $\text{mr}_+(\cdot)$ as the *Frisch minimum rank* and $\text{mr}(\cdot)$ as the *Shapiro minimum rank*. The former is lower semicontinuous, whereas the latter is not, as stated next. Naturally, this difference may turn out to be important in certain cases when dealing with experimental data.

Proposition 1— $\text{mr}_+(\cdot)$ is lower semicontinuous, whereas $\text{mr}(\cdot)$ is not.

Proof: Assume that for a given $\Sigma \succ 0$ there exists a sequence $\Sigma_1, \Sigma_2, \dots$ of positive definite matrices such that $\Sigma_j \rightarrow \Sigma$, while

$$\text{mr}_+(\sum_i) < \text{mr}_+(\sum) = r \quad \forall i=1, 2, \dots$$

Decompose $\Sigma_i = \hat{\Sigma}_i + D_i$ with $\text{rank}(\hat{\Sigma}_i) < r$, $\Sigma_i - D_i = 0$, and D_i diagonal. Then there exist convergent subsequences $\hat{\Sigma}_{i_k} \rightarrow \hat{\Sigma}$ and $D_{i_k} \rightarrow D$, as $k \rightarrow \infty$. Since $\Sigma_{i_k} \rightarrow \hat{\Sigma} + D = \Sigma$, by the lower semicontinuity of the rank,

$$\text{rank}(\hat{\Sigma}) \leq \liminf_{k \rightarrow \infty} \text{rank}(\hat{\Sigma}_{i_k}) < r = \text{mr}_+(\Sigma).$$

This is a contradiction. On the other hand, to see that $\text{mr}(\cdot)$ is not lower semicontinuous, consider

$$\Sigma = \begin{bmatrix} 3 & -1 & -1 \\ -1 & 3 & 0 \\ -1 & 0 & 3 \end{bmatrix} \quad \text{and} \quad \Sigma_\varepsilon = \begin{bmatrix} 3 & -1 & -1 \\ -1 & 3 & \varepsilon \\ -1 & \varepsilon & 3 \end{bmatrix}, \quad \hat{\Sigma}_\varepsilon = \begin{bmatrix} 1/\varepsilon & -1 & -1 \\ -1 & \varepsilon & \varepsilon \\ -1 & \varepsilon & \varepsilon \end{bmatrix}$$

for $\varepsilon > 0$. Clearly $\text{mr}(\Sigma) = 2$. Also $\lim_{\varepsilon \rightarrow 0} \Sigma_\varepsilon = \Sigma$. Yet $\Sigma_\varepsilon = \hat{\Sigma}_\varepsilon + D_\varepsilon$, while Σ_ε has rank 1 and D_ε is diagonal ($\neq 0$). Hence $\text{mr}(\Sigma_\varepsilon) = 1$.

Assuming that the off-diagonal entries of $\Sigma > 0$ of size $n \times n$ are known with absolute certainty, any “minimum rank” ($\text{mr}_+(\cdot)$ and $\text{mr}(\cdot)$) is bounded below by the so-called Ledermann bound, i.e.,

$$\frac{2n+1-\sqrt{8n+1}}{2} \leq \text{mr}(\Sigma) \leq \text{mr}_+(\Sigma), \quad (4.3)$$

which holds on a generic set of positive definite matrices Σ , that is, on a (Zariski open) subset of positive definite matrices. Equivalently, the set of matrices Σ for which $\text{mr}(\Sigma)$ is lower than the Ledermann bound is nongeneric—their entries satisfy algebraic equations which fail under small perturbation. To see this, consider any factorization

$$\Sigma = FF',$$

with $F \in \mathbb{R}^{n \times r}$. There are $(n-r)r + \frac{r(r+1)}{2}$ independent entries in F (when accounting for the action of a unitary transformation of F on the right), whereas the value of the off-diagonal entries of Σ impose $\frac{n(n-1)}{2}$ constraints. Thus, the number of independent entries in F exceeds the number of constraints when $(n-r)^2 - n + r$, which then leads to the inequality $\frac{2n+1-\sqrt{8n+1}}{2} \leq r$. The bound was first noted in [55], while the independence of the constraints has been detailed in [10]. In general, the computation of the exact value for $\text{mr}_+(\Sigma)$ and $\text{mr}(\Sigma)$ is a nontrivial matter. Thus, it is rather surprising that an exact analytic result is

available for both in the special case when $r = n - 1$. We review this next in the form of two theorems.

Theorem 2 (Reiersøl's theorem [64])—Let $\Sigma \in \mathbf{S}_{n,+}$ and $\Sigma > 0$. Then

$$\text{mr}_+(\Sigma) = n - 1 \iff \Sigma^{-1} \succ_e 0.$$

Theorem 3 (Shapiro's theorem [69])—Let $\Sigma \in \mathbf{S}_{n,+}$ and be irreducible. Then

$$\text{mr}(\Sigma) = n - 1 \iff \Sigma \preceq_e 0.$$

The characterization of covariance/covariance matrices Σ for which $\text{mr}_+(\Sigma) = n - 1$ was first recognized by T. C. Koopmans in 1937 [53] and proven by Reiersøl [64], who used the Perron–Frobenius theory to improve on Koopmans' analysis. Later on, R. E. Kalman streamlined and completed the steps in [45] relying again on the Perron–Frobenius theorem (see also Klepper and Leamer [52] for a detailed analysis). Our treatment below takes a slightly different angle and provides some geometric insight by pointing as a key reason that the maximal number of vectors at an obtuse angle from one another can exceed the dimension of the ambient space by at most one (Corollary 6). We provide new proofs where we also utilize a dual formulation with an analogous decomposition of the inverse covariance.

4.1. A Geometric Insight

We begin with two basic lemmas for irreducible matrices in $M \in \mathbf{S}_{n,+}$. Recall that a matrix is reducible if by permutation of rows and columns it can be brought into a block-diagonal form; otherwise it is irreducible.

Lemma 4—Let $M > 0$ and be irreducible. Then

$$M \preceq_e 0 \Rightarrow M^{-1} \succ_e 0. \quad (4.4)$$

Lemma 5—Let $M \succeq 0$ and be irreducible. Then

$$M \preceq_e 0 \Rightarrow \text{nullity}(M) \leq 1. \quad (4.5)$$

Proof: It is easy to verify that for matrices of size 2×2 , (4.4) holds true. Assume that the statement also holds true for matrices of size up to $k \times k$ for a certain value of k , and consider a matrix M of size $(k + 1) \times (k + 1)$ with $M > 0$ and $M \preceq_e 0$. Partition

$$M = \begin{bmatrix} A & b \\ b' & c \end{bmatrix}$$

so that c is a scalar, and hence A is of size $k \times k$. Partitioning conformably,

$$M^{-1} = \begin{bmatrix} F & g \\ g' & h \end{bmatrix},$$

where

$$F = (A - bc^{-1}b')^{-1}, \quad g = -A^{-1}bh, \quad \text{and } h = (c - b'A^{-1}b)^{-1} > 0.$$

For the case where A is irreducible, because A has size $k \times k$ and $A \preceq_e 0$, invoking our hypothesis we conclude that $A^{-1} \succ_e 0$. Now, since b has only nonpositive entries and $b' \preceq_e 0$, $g = -A^{-1}bh$ has positive entries. Since $-bc^{-1}b' \preceq_e 0$ and $A \preceq_e 0$, then $A - bc^{-1}b' \preceq_e 0$ is also irreducible. Thus $F = (A - bc^{-1}b')^{-1}$ has positive entries by hypothesis.

For the case where A is reducible, permutation of columns and rows brings A into a block-diagonal form with irreducible blocks. Thus, A^{-1} is also a block-diagonal matrix with each block entrywise positive. Because M is irreducible, b must have at least one nonzero entry corresponding to the rows of each diagonal block of A . Then $A - bc^{-1}b'$ is irreducible and $\preceq_e 0$. Also $A^{-1}b$ has all of its entries negative. Therefore $F = (A - bc^{-1}b')^{-1}$ and $g = -A^{-1}bh$ have positive entries. Therefore $M^{-1} \succ_e 0$.

Proof: Rearrange rows and columns and partition

$$M = \begin{bmatrix} A & B \\ B' & C \end{bmatrix}$$

so that A is nonsingular and of maximal size, equal to the rank of M . Then

$$C = B'A^{-1}B. \quad (4.6)$$

We first show that $B'A^{-1}B \succeq_e 0$. Assume that A is irreducible. Then $A^{-1} \succ_e 0$. At the same time B has negative entries and not all zero (since M is irreducible). In this case, $B'A^{-1}B \succ_e 0$. If on the other hand A is reducible, Lemma 4 applied to the (irreducible) blocks of A implies that $A^{-1} \succeq_e 0$. Therefore, in this case, $B'A^{-1}B \succeq_e 0$.

Returning to (4.6), and in view of the fact that $C \preceq_e 0$ while $B'A^{-1}B \succeq_e 0$, we conclude that either C is a scalar (and hence there are no off-diagonal negative entries) or both C and B

$A^{-1}B$ are diagonal. The latter contradicts the assumption that M is irreducible. Hence, the nullity of M can be at most 1.

Lemma 5 provides the following geometric insight, stated as a corollary.

Corollary 6—*In any Euclidean space of dimension n , there can be at most $n+1$ vectors forming an obtuse angle with one another.*

Proof: The Grammian $M = [v'_k v_\ell]_{k,\ell=1}^{n+q}$ of a selection $\{v_k \mid k = 1, \dots, n+q\}$ of such vectors has off-diagonal entries which are negative. Hence, by Lemma 5, the nullity of M cannot exceed 1.

The necessity part of Theorem 3 is also a direct corollary of Lemma 5.

Corollary 7—*Let $\Sigma \in \mathbf{S}_{n,+}$ and be irreducible. Then*

$$\sum \leq_e 0 \Rightarrow \text{mr}(\Sigma) = n-1.$$

Proof: Let $\Sigma = \hat{\Sigma} + \tilde{\Sigma}$, with $\tilde{\Sigma}$ diagonal and $\hat{\Sigma} \geq 0$. $\tilde{\Sigma}$ is irreducible since Σ is irreducible. From Lemma 5, the nullity of $\tilde{\Sigma}$ is at most 1. Thus $\text{mr}(\Sigma) = n-1$.

4.2. A Dual Decomposition

The matrix inversion lemma provides a correspondence between an additive decomposition of a positive definite matrix and a decomposition of its inverse, albeit with a different sign in one of the summands. This is stated next.

Lemma 8—Let

$$\Sigma = D + FF' \quad (4.7)$$

with $\Sigma, D \in \mathbf{S}_{n,+}$, with $\Sigma, D > 0$ and $F \in \mathbb{R}^{n \times r}$. Then

$$S := \Sigma^{-1} = E - GG' \quad (4.8)$$

for $E = D^{-1}$ and $G = D^{-1}F(I + F'D^{-1}F)^{-1/2}$. Conversely, if (4.8) holds with $G \in \mathbb{R}^{n \times r}$, then so does (4.7) for $D = E^{-1}$ and $F = E^{-1}G(I - G'E^{-1}G)^{-1/2}$.

Proof: This follows from the identity $(I \pm MM')^{-1} = I \mp M(I \mp M'M)^{-1}M'$.

Application of the lemma suggests the following variation to Frisch's problem.

Problem 3 (the dual Frisch problem)—*Given a positive definite $n \times n$ symmetric matrix S , determine the dual minimum rank:*

$$\text{mr}_{\text{dual}}(S) := \min\{\text{rank}(\hat{S}) | S = E - \hat{S}, \hat{S}, E \geq 0, E \text{ is diagonal}\}.$$

Clearly, if $S = \Sigma^{-1} = E - GG'$ (as in (4.8)), then $E > 0$. Furthermore, a decomposition of S always gives rise to a decomposition $\Sigma = D + FF'$ (as in (4.7)) with the terms FF' and GG' having the same rank. Thus, it is clear that

$$\text{mr}_+(\Sigma) \leq \text{mr}_{\text{dual}}(\Sigma^{-1}), \quad (4.9)$$

and that the above holds with equality when an optimal choice of $D \equiv \tilde{\Sigma}$ in (4.1) is invertible. However, if D is allowed to be singular, the rank of the summands FF' and GG' may not agree. This can be seen using the following example. Take

$$\Sigma = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

It is clear that Σ admits a decomposition $\Sigma = \tilde{\Sigma} + \hat{\Sigma}$ in correspondence with (4.7), where $\tilde{\Sigma} = D = \text{diag}\{1, 1, 0\}$, while $\hat{\Sigma} = FF'$ as well as $F = [1, 1, 1]$ are of rank one. On the other hand,

$$S = \Sigma^{-1} = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ -1 & -1 & 3 \end{bmatrix}.$$

Taking $E = \text{diag}\{e_1, e_2, e_3\}$ in (4.8), it is evident that the rank of

$$GG' = E - S = \begin{bmatrix} e_1 - 1 & 0 & 1 \\ 0 & e_2 - 1 & 1 \\ 1 & 1 & e_3 - 3 \end{bmatrix}$$

cannot be less than 2 without violating the nonnegativity assumption for the summand GG' . The minimal rank for the factor G is 2 and is attained by taking $e_1 = e_2 = 2$ and $e_3 = 5$.

On the other hand, in general, if we perturb Σ to $\Sigma + \varepsilon I$ and, accordingly, D to $D + \varepsilon I$, then

$$\text{mr}_{\text{dual}}((\Sigma + \varepsilon I)^{-1}) \leq \text{mr}_+(\Sigma) \quad \forall \varepsilon > 0. \quad (4.10)$$

Equality in (4.10) holds for sufficiently small value of ε . Thus, mr_+ and mr_{dual} are closely related. However, it should be noted that $\text{mr}_{\text{dual}}(\cdot)$ fails to be lower semicontinuous since a small perturbation of the off-diagonal entries can reduce $\text{mr}_{\text{dual}}(\cdot)$. Yet, interestingly, an exact characterization of the $\text{mr}_{\text{dual}}(S) = n - 1$ can be obtained which is analogous to those

for mr_+ and mr being equal to $n - 1$; the condition for mr_{dual} will be used to prove the Reiersøl and Shapiro theorems.

Theorem 9—For $S \in \mathbf{S}_{n,+}$, with $S > 0$ and irreducible,

$$\text{mr}_{\text{dual}}(S) = n - 1 \iff S \succeq_e 0. \quad (4.11)$$

Proof: If $S \succeq_e 0$ and E is diagonal satisfying $E - S > 0$, then $E - S = GG^T \succeq_e 0$. By invoking Lemma 5 we deduce that if $E - S$ is singular, $\text{rank}(G) = n - 1$. Hence, $\text{mr}_{\text{dual}}(S) = n - 1$.

To establish that $\text{mr}_{\text{dual}}(S) = n - 1 \Rightarrow S \succeq_e 0$, we assume that the condition $S \succeq_e 0$ fails and show that $\text{mr}_{\text{dual}}(S) < n - 1$. We first argue the case for a 3×3 matrix $S = [s_{ij}]_{i,j=1}^3$. Provided $S \not\succeq_e 0$, we can assume that it has strictly negative off-diagonal entries (which can be done by reflecting the signs of rows and columns). We now let

$$e_i = s_{ii} - \frac{s_{ij}s_{ki}}{s_{jk}}$$

for $i \in \{1, 2, 3\}$, with (i, j, k) being permutations of $(1, 2, 3)$. These are all positive. Let $\tilde{S} = \text{diag}^*(e_1, e_2, e_3)$. It can be seen that $\tilde{S} - S \succeq 0$, while $\text{rank}(\tilde{S} - S) = 1$. To verify the latter observe that $\tilde{S} - S = \nu\nu'$ for

$$\nu' = [\sqrt{e_1 - s_{11}}, \sqrt{e_2 - s_{22}}, \sqrt{e_3 - s_{33}}].$$

This establishes the reverse implication for matrices of size 3×3 .

We now assume that the statement holds true for matrices of size up to $(n - 1) \times (n - 1)$ for some $n \geq 4$ and use induction. So let $S, \tilde{S}$ be of size $n \times n$ with $S \not\succeq_e 0$ and $\tilde{S}$ diagonal. We need to prove that $\text{mr}_{\text{dual}}(S) < n - 1$. We partition

$$S = \begin{bmatrix} A & b \\ b' & c \end{bmatrix}, \quad \tilde{S} = \begin{bmatrix} E & 0 \\ 0 & e \end{bmatrix},$$

with A, E being $(n - 1) \times (n - 1)$. For any $\tilde{S}$ such that $\tilde{S} - S \succeq 0$, e cannot be equal to c ; otherwise $b = 0$ and S is reducible. Further, $\tilde{S} - S \succeq 0$ if and only if $e > c$ and

$$M := E - (A + b(e - c)^{-1}b') \succeq 0.$$

The nullity of $S - S$ coincides with that of M . To prove our claim, it suffices to show that $A_e := A + b(e - c)^{-1}b' \not\geq_e 0$, or that A_e is reducible for some $e > c$ (since, in either case, by our hypothesis the nullity of M for a suitable E exceeds 1).

We now consider two possible cases where $S \succeq_e 0$ fails. First, we consider the case where already $A \not\geq_e 0$. Then obviously $A_e \not\geq_e 0$ for $e - c$ sufficiently large. The second possibility is $S \not\geq_e 0$, while $A \succeq_e 0$. But if A is (transformed into) elementwise nonnegative, then bb' must have at least one pair of negative off-diagonal entries. Then consider $A_e = A + \lambda bb'$ for $\lambda = (e - c)^{-1} \in (0, \infty)$. Evidently, for certain values of λ , entries of A_e change sign. If a whole row becomes zero for a particular value of λ , then A_e is reducible. In all other cases, there are values of λ for which $A_e \not\geq_e 0$. This completes the proof.

4.3. Proof of Reiersøl's Theorem (Theorem 2)

We first show that $\Sigma^{-1} \succ_e 0$ implies $\text{mr}_+(\Sigma) = n - 1$. From the continuity of the inverse, $(\Sigma + \varepsilon I)^{-1} \succ_e 0$ for sufficiently small $\varepsilon > 0$. Applying Theorem 9, we conclude that

$$\text{mr}_{\text{dual}}((\Sigma + \varepsilon I)^{-1}) = n - 1.$$

Since $\text{mr}_+(\Sigma) = \text{mr}_{\text{dual}}((\Sigma + \varepsilon I)^{-1})$ as in (4.10), we conclude that $\text{mr}_+(\Sigma) = n - 1$.

To prove that $\text{mr}_+(\Sigma) = n - 1 \Rightarrow \Sigma^{-1} \succ_e 0$, we show that assuming $\Sigma^{-1} \not\succ_e 0$ and $\text{mr}_+(\Sigma) = n - 1$ together leads to a contradiction. From the continuity of the inverse and the lower semicontinuity of $\text{mr}_+(\cdot)$ (Proposition 1), there exists a symmetric matrix Δ and an $\varepsilon > 0$ such that

$$(\Sigma + \varepsilon \Delta)^{-1} \not\geq_e 0 \quad \text{and} \quad \text{mr}_+(\Sigma + \varepsilon \Delta) = n - 1.$$

Then, from Theorem 9, $\text{mr}_{\text{dual}}((\Sigma + \varepsilon \Delta)^{-1}) < n - 1$, while from (4.9)

$$\text{mr}_+(\Sigma + \varepsilon \Delta) \leq \text{mr}_{\text{dual}}((\Sigma + \varepsilon \Delta)^{-1}).$$

Thus, we have a contradiction and therefore $\Sigma^{-1} \succ_e 0$.

It is well known that zero entries in the inverse covariance matrix of a multivariate normal distribution imply the conditional independence of the corresponding variables given the remaining ones; see [26, 23]. Interestingly, a single zero element in the inverse covariance violates the condition in the Reiersøl's theorem and in this case $\text{mr}_+(\Sigma) = n - 2$.

4.4. Proof of Shapiro's Theorem (Theorem 3)

Given $\Sigma \succ 0$, consider $\lambda > 0$ such that $\lambda I - \Sigma \succ 0$, a diagonal D , and let $E := \lambda I - D$. Since $\Sigma - D = E - (\lambda I - \Sigma)$,

$$\text{mr}(\Sigma) = \text{mr}_{\text{dual}}(\lambda I - \Sigma). \quad (4.12)$$

If Σ is irreducible and $\Sigma \preceq_e 0$, then $\lambda I - \Sigma$ is irreducible and $\lambda I - \Sigma \succeq_e 0$. It follows (Theorem 9) that $\text{mr}_{\text{dual}}(\lambda I - \Sigma) = n - 1$, and therefore $\text{mr}(\Sigma) = n - 1$ as well.

For the reverse direction, if $\text{mr}(\Sigma) = n - 1$, then $\text{mr}_{\text{dual}}(\lambda I - \Sigma) = n - 1$, which implies that $\lambda I - \Sigma \succeq_e 0$ and therefore that $\Sigma \preceq_e 0$.

The original proof in [69] claims that for any $\Sigma \succeq_e 0$ of size $n \times n$ with $n > 3$ and $\Sigma \not\preceq_e 0$, there exists an $(n - 1) \times (n - 1)$ principle minor that is $\not\preceq_e 0$. This statement fails for the following sign pattern:

$$\begin{bmatrix} + & 0 & - & - \\ 0 & + & - & + \\ - & - & + & 0 \\ - & + & 0 & + \end{bmatrix}.$$

This matrix cannot be transformed to have all nonpositive off-diagonal entries, yet all its 3×3 principle minors $\preceq_e 0$.

4.5. Parametrization of Solutions under Reiersøl's and Shapiro's Conditions

For either the Frisch or the Shapiro problem, a solution is not unique in general. The parametrization of solutions to the Frisch problem when $\text{mr}_+(\Sigma) = n - 1$ has been known and is briefly explained below (without proof). Interestingly, an analogous parametrization is possible for Shapiro's problem, and this is given in Proposition 11 below, and both are presented here for completeness of the exposition.

For the Frisch problem, and assuming that Σ satisfies Reiersøl's condition, the set of nonnegative definite diagonal matrices D such that $\Sigma - D \succeq_e 0$ and singular is homeomorphic to the simplex of probability vectors. The precise result is stated next (without proof).

Proposition 10—Let $\Sigma \in \mathbf{S}_{n,+}$ with $\Sigma \succ 0$ and $\Sigma^{-1} \succ_e 0$. The following hold:

- i. For $D \succeq_e 0$ diagonal with $\Sigma - D \succeq_e 0$ and singular, there is a probability vector ρ (ρ has entries ≥ 0 that sum up to 1) such that $(\Sigma - D)\Sigma^{-1}\rho = 0$.
- ii. For any probability vector ρ ,

$$D = \text{diag}^* \left(\left[\frac{[\rho]_i}{[\Sigma^{-1}\rho]_i}, i=1, \dots, n \right] \right)$$

satisfies $\Sigma - D \succeq_e 0$ and $\Sigma - D$ is singular.

Proof: See [45, 52].

Thus, solutions of Frisch's problem under Reiersøl's conditions are in bijective correspondence with probability vectors. A very similar result holds true for Shapiro's problem and is stated next.

Proposition 11—Let $\Sigma \in \mathbf{S}_{n,+}$ be irreducible and have 0 off-diagonal entries. The following hold:

- i. For D diagonal with $\Sigma - D \succeq 0$ and singular, there is a strictly positive vector v such that $(\Sigma - D)v = 0$.
- ii. For any strictly positive vector $v \in \mathbb{R}^{n \times 1}$,

$$D = \text{diag}^* \left(\left[\frac{[\Sigma v]_i}{[v]_i}, i=1, \dots, n \right] \right) \quad (4.13)$$

satisfies that $\Sigma - D \succeq 0$ and $\Sigma - D$ is singular.

Proof: To prove (i), we note that if $(\Sigma - D)v = 0$, then v is elementwise positive. To see this consider $(\Sigma - D + \varepsilon I)^{-1}$ for $\varepsilon > 0$. From Lemma 4,

$$(\Sigma - D + \varepsilon I)^{-1} \succ_e 0,$$

and since v is an eigenvector corresponding to its largest eigenvalue, a power iteration argument concludes that v is elementwise positive.

To prove (ii), it is easy to verify that the diagonal matrix D in (4.13) for $v \succ_e 0$ satisfies $(\Sigma - D)v = 0$. We only need to prove that $\Sigma - D \succeq 0$. Without loss of generality we assume that all the entries of v are equal. (This can always be done by scaling the entries of v and scaling accordingly rows and columns of Σ .) Since v is a null vector of $\Sigma - D$ and since $M := \Sigma - D$ has 0 off-diagonal entries,

$$[M]_{ii} = \sum_{j \neq i} |[M]_{ij}|.$$

Geršgorin's circle theorem (e.g., see [76]) now states that every eigenvalue of M lies within at least one of the closed discs $\{\text{Disk}([M]_{ii}, \sum_{j \neq i} |[M]_{ij}), i = 1, \dots, n\}$. No disc intersects the negative real line. Therefore $\Sigma - D \succeq 0$.

4.6. Decomposition of Complex-Valued Matrices

Complex-valued covariance matrices are commonly used in radar and antenna arrays [75]. The rank of $\Sigma - D$, for noise covariance D as in the Frisch problem, is an indication of the number of (dominant) scatterers in the scattering field. If this is of the same order as the number of array elements (e.g., $n - 1$), any conclusion about their location may be suspect. Thus, it is natural to seek conditions for $\text{mr}_+(\Sigma) = n - 1$ analogous to those given by

Reiersøl, for the case of complex covariances, as a possible warning. This we do next. We note in passing that a version of the Frisch problem for complex-valued covariance matrices is also relevant in the context of multivariate time series [67].

Consider complex-valued observation vectors $x_t = y_t + iz_t$, $t = 1, \dots, T$, where $i = \sqrt{-1}$ and $y_t, z_t \in \mathbb{R}^{m \times 1}$, and set

$$X = [x_1, \dots, x_T] = Y + iZ,$$

with $Y = [y_1, \dots, y_T]$, $Z = [z_1, \dots, z_T]$. The (scaled) sample covariance is

$$\Sigma = XX^* = \Sigma_r + i\Sigma_i \in \mathbf{H}_{n,+},$$

where the real part $\Sigma_r := YY' + ZZ'$ is symmetric, the imaginary part $\Sigma_i := ZY' - YZ'$ is antisymmetric, and “*” denotes complex-conjugate transpose. As before, we consider a decomposition

$$\Sigma = \hat{\Sigma} + D,$$

with $\hat{\Sigma} \geq 0$ singular and $D \geq 0$ diagonal. We refer the reader to [3, 19] for the special case where $\text{mr}_+(\Sigma) = 1$. In this section we present a sufficient condition for a Reiersøl case where $\text{mr}_+(\Sigma) = n - 1$.

Before we proceed we note that recasting the problem in terms of the real-valued

$$R := \begin{bmatrix} \Sigma_r & \Sigma_i \\ \Sigma_i & \Sigma_r \end{bmatrix} \in \mathbf{S}_{2n,+}$$

does not allow taking advantage of earlier results. The structure of R with antisymmetric off-diagonal blocks implies that if $[a', b']'$ is a null vector, then so is $[-b', a']'$ (since, accordingly, $a + ib$ and $ia - b$ are both null vectors of Σ). Thus, in general, the nullity of R is not 1 and the theorem of Reiersøl is not applicable. Further, the corresponding noise covariance is diagonal with repeated blocks.

The following lemmas for the complex case echo Lemmas 4 and 5.

Lemma 12—*Let $M \in \mathbf{H}_{n,+}$ be irreducible. If the argument of each nonzero off-diagonal entry of $-M$ is in $(-\frac{\pi}{2n}, \frac{\pi}{2n})$, then each entry of M^{-1} has an argument in $(-\frac{\pi}{2} + \frac{\pi}{2n}, \frac{\pi}{2} - \frac{\pi}{2n})$.*

Proof: It is easy to verify the lemma for 2×2 matrices. Assume that the statement holds for sizes up to $n \times n$ and consider an $(n+1) \times (n+1)$ matrix M that satisfies the conditions of the lemma. Partition

$$M = \begin{bmatrix} A & b \\ b^* & c \end{bmatrix},$$

where A is of size $n \times n$, and conformably

$$M^{-1} = \begin{bmatrix} F & g \\ g^* & h \end{bmatrix}.$$

By assumption nonzero entries of $-A$ and $-b$ have their argument in $(-\frac{\pi}{2^{n+1}}, \frac{\pi}{2^{n+1}})$. Then, by bounding the possible contribution of the respective terms, it follows that the argument of each of the entries of $-A + bc^{-1}b^*$ is in $(-\frac{\pi}{2^n}, \frac{\pi}{2^n})$. Then the argument of each entry of $F = (A - bc^{-1}b^*)^{-1}$ is in $(-\frac{\pi}{2} + \frac{\pi}{2^n}, \frac{\pi}{2} - \frac{\pi}{2^n})$; this follows by assumption since F is $n \times n$. Clearly, $(-\frac{\pi}{2} + \frac{\pi}{2^n}, \frac{\pi}{2} - \frac{\pi}{2^n}) \subset (-\frac{\pi}{2} + \frac{\pi}{2^{n+1}}, \frac{\pi}{2} - \frac{\pi}{2^{n+1}})$. Regarding g , by bounding the possible contribution of respective terms, we similarly conclude that the argument of each of its nonzero entries is in $(-\frac{\pi}{2} + \frac{\pi}{2^{n+1}}, \frac{\pi}{2} - \frac{\pi}{2^{n+1}})$.

Lemma 13—Let $M \in \mathbf{H}_{n,+}$ be irreducible. If the argument of each nonzero off-diagonal entry of $-M$ is in $(-\frac{\pi}{2^n}, \frac{\pi}{2^n})$, then $\text{rank}(M) = n - 1$.

Proof: First rearrange rows and columns of M , and partition as

$$M = \begin{bmatrix} A & B \\ B^* & C \end{bmatrix}$$

so that A is nonsingular and of size equal to the rank of M , which we denote by r . Then

$$C = B^* A^{-1} B \quad (4.14)$$

and has size equal to the nullity of M . We now compare the argument of the off-diagonal entries of C and $B^* A^{-1} B$ and show they cannot be equal unless C is a scalar. Since the off-diagonal entries of $-A$ have their argument in $(-\frac{\pi}{2^n}, \frac{\pi}{2^n}) \subset (-\frac{\pi}{2^r}, \frac{\pi}{2^r})$, the off-diagonal entries of A^{-1} have their argument in $(-\frac{\pi}{2} + \frac{\pi}{2^r}, \frac{\pi}{2} - \frac{\pi}{2^r})$ from Lemma 12. Now, the (k, ℓ) entry of $B^* A^{-1} B$ is

$$[B^* A^{-1} B]_{k\ell} = \sum_{i,j} [B^*]_{ki} [A^{-1}]_{ij} [B]_{j\ell},$$

and the phase of each summand is

$$\arg([B^*]_{ki}[A^{-1}]_{ij}[B]_{j\ell}) \in \left(-\frac{\pi}{2} + \frac{\pi}{2r} - \frac{\pi}{2^{n-1}}, \frac{\pi}{2} - \frac{\pi}{2r} + \frac{\pi}{2^{n-1}}\right).$$

Thus, the nonzero off-diagonal entries of $B^*A^{-1}B$ have positive real part, while

$$\arg(-[C]_{k\ell}) \in \left(-\frac{\pi}{2^n}, \frac{\pi}{2^n}\right).$$

Hence, either the off-diagonal entries of $B^*A^{-1}B$ and C are zero, in which case these are diagonal matrices and M must be reducible, or $B^*A^{-1}B$ and C are both scalars. This concludes the proof.

Theorem 14—Let $\Sigma \in \mathbf{H}_{n,+}$ be irreducible. If the argument of each nonzero off-diagonal entry of $-\Sigma$ is in $(-\frac{\pi}{2^n}, \frac{\pi}{2^n})$, then $\text{mr}(\Sigma) = n - 1$.

Proof: The matrix $\Sigma - D$ is irreducible since D is diagonal. If $\Sigma - D = 0$ and singular, and since the argument of each nonzero off-diagonal entry of $-(\Sigma - D)$ is in $(-\frac{\pi}{2^n}, \frac{\pi}{2^n})$, Lemma 13 applies and gives that $\text{rank}(\Sigma - D) = n - 1$.

Clearly, since $\text{mr}_+(\Sigma) = \text{mr}(\Sigma)$, under the condition of Theorem 14, $\text{mr}_+(\Sigma) = n - 1$. It is also clear that for $S \in \mathbf{H}_{n,+}$ irreducible with all nonzero off-diagonal entries having argument in $(-\frac{\pi}{2^n}, \frac{\pi}{2^n})$, we also conclude that $\text{mr}_{\text{dual}}(S) = n - 1$.

5. Trace-Minimization Heuristics

The rank of a matrix is a nonconvex function of its elements, and the problem of finding the matrix of minimal rank within a given set is a difficult one, in general. Therefore, certain heuristics have been developed over the years to obtain approximate solutions. In particular, in the context of factor analysis, trace minimization has been pursued as a suitable heuristic [56, 68, 69], thereby relaxing the Frisch problem into

$$\min_D \{\text{trace}(\Sigma - D) \mid \Sigma \geq D \geq 0\}$$

for a diagonal matrix D , with a relaxation of $D = 0$ corresponding to Shapiro's problem. The theoretical basis for using the trace and, more generally, the nuclear norm for nonsymmetric matrices as a surrogate for the rank was provided by Fazel, Hindi, and Boyd [28], who proved that these constitute convex envelopes of the rank function on bounded sets of matrices.

The relation between minimum-trace factor analysis and minimum-rank factor analysis goes back to Ledermann in [56] (see [22] and [66]). Herein we only refer to two propositions which characterize minimizers for the two problems, Frisch's and Shapiro's, respectively.

Proposition 15 (see [22])—Let $\Sigma = \hat{\Sigma}_1 + D_1 > 0$ for a diagonal $D_1 \geq 0$. Then

$$(\hat{\Sigma}_1, D_1) = \arg \min \{ \text{trace}(\hat{\Sigma}) \mid \Sigma = \hat{\Sigma} + D > 0, \hat{\Sigma} \geq 0, \text{diagonal } D \geq 0 \} \\ \iff \exists \Lambda_1 \geq 0: \hat{\Sigma}_1 \Lambda_1 = 0 \text{ and } \begin{cases} [\Lambda_1]_{ii} = 1 & \text{if } [D_1]_{ii} > 0, \\ [\Lambda_1]_{ii} \geq 1 & \text{if } [D_1]_{ii} = 0. \end{cases} \quad (5.1a)$$

Proposition 16 (see [66])—Let $\Sigma = \hat{\Sigma}_2 + D_2 > 0$ for a diagonal $D_2 \geq 0$. Then,

$$(\hat{\Sigma}_2, D_2) = \arg \min \{ \text{trace}(\hat{\Sigma}) \mid \Sigma = \hat{\Sigma} + D > 0, \hat{\Sigma} \geq 0, \text{diagonal } D \geq 0 \} \\ \iff \exists \Lambda_2 \geq 0: \hat{\Sigma}_2 \Lambda_2 = 0 \text{ and } [\Lambda_2]_{ii} = 1 \forall i. \quad (5.1b)$$

Evidently, when the solutions to these two problems differ and $D_1 \neq D_2$, then there exists $k \in \{1, \dots, n\}$ such that

$$[D_2]_{kk} < 0 \text{ and } [D_1]_{kk} = 0.$$

Further, the essence of Proposition 16 is that a singular $\hat{\Sigma}$ originates from such a minimization problem if and only if there is a correlation matrix in its null space. The matrices Λ_1 and Λ_2 appear as Lagrange multipliers in the respective problems.

Factor analysis is closely related to *low-rank matrix completion* as well as to *sparse and low-rank decomposition* problems. Typically, low-rank matrix completion asks for a matrix X which satisfies a linear constraint $\mathcal{A}(X) = b$ and has low/minimal rank ($\mathcal{A}(\cdot)$ denotes a linear map $\mathcal{A}: \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^p$). Thus, factor analysis corresponds to the special case where $\mathcal{A}(\cdot)$ maps X onto its off-diagonal entries. In a recent work by Recht, Fazel, and Parrilo [63], the nuclear norm of X was considered as a convex relaxation of $\text{rank}(X)$ for such problems, and a sufficient condition for exact recovery was provided. However, this sufficient condition amounts to the requirement that the null space of $\mathcal{A}(\cdot)$ contain no matrix of low rank. Therefore, since in factor analysis diagonal matrices are in fact contained in the null space of $\mathcal{A}(\cdot)$ and include matrices of low rank, the condition in [63] does not apply directly. Other works on low-rank matrix completion (see, e.g., [15]) mainly focus on assessing the probability of exact recovery and on constructing efficient computational algorithms for *large-scale* low-rank completion problems [49, 50]. On the other hand, since diagonal matrices are sparse (most of their entries are zero), the work on matrix decomposition into sparse and low-rank components by Chandrasekaran et al. [16] is very pertinent. In this, the ℓ_1 and nuclear norms were used as surrogates for sparsity and rank, respectively, and a sufficient condition for exact recovery was provided which captures a certain “rank-sparsity

incoherence”; an analogous but stronger sufficient “incoherence” condition which applies to problem (5.1b) is given in [66].

5.1. Weighted Minimum-Trace Factor Analysis

Both $\text{mr}(\Sigma)$ and $\text{mr}_+(\Sigma)$ in (4.1) and (4.2), respectively, remain invariant under scaling of rows and the corresponding columns of Σ by the same coefficients. On the other hand, the minimizers in (5.1a) and (5.1b) and their respective ranks are not invariant under scaling. This fact motivates weighted-trace minimization,

$$\min \left\{ \text{trace}(W\hat{\Sigma}) \mid \Sigma = \hat{\Sigma} + D, \hat{\Sigma} \geq 0, \text{diagonal } D \geq 0 \right\}, \quad (5.2)$$

given $\Sigma > 0$ and a diagonal weight $W > 0$. As before the characterization of minimizers relates to a suitable condition for the corresponding Lagrange multipliers.

Proposition 17 (see [69])—Let $\Sigma = \Sigma_0 + D_0 > 0$ for a diagonal matrix $D_0 \geq 0$ and consider a diagonal $W > 0$. Then

$$\begin{aligned} (\hat{\Sigma}_0, D_0) &= \arg \min \{ \text{trace}(W\hat{\Sigma}) \mid \Sigma = \hat{\Sigma} + D > 0, \hat{\Sigma} \geq 0, \text{diagonal } D \geq 0 \} \\ \iff \exists \Lambda_0 \geq 0: \hat{\Sigma}_0 \Lambda_0 &= 0 \text{ and } \begin{cases} [\Lambda_0]_{ii} = [W]_{ii} & \text{if } [D_0]_{ii} > 0, \\ [\Lambda_0]_{ii} \geq [W]_{ii} & \text{if } [D_0]_{ii} = 0. \end{cases} \end{aligned} \quad (5.3)$$

A corresponding sufficient and necessary condition for $(\hat{\Sigma}, D)$ to be a minimizer in Shapiro’s problem is that there exists a Grammian in the null space of $\hat{\Sigma}$ whose diagonal entries are equal to the diagonal entries of W .

Minimum-rank solutions may be recovered as solutions to (5.3) using suitable choices of weight. However, these choices depend on Σ and are not known in advance; this motivates a selection of certain canonical Σ -dependent weight as well as iteratively improving the choice of weight. One should note that since D is diagonal, letting W be a not necessarily diagonal matrix does not change the problem—only the diagonal entries of W determine the minimizer.

We first consider taking $W = \Sigma^{-1}$. A rationale for this choice is that the minimal value in (5.2) bounds $\text{mr}_+(\Sigma)$ from below, since for any decomposition $\Sigma = \hat{\Sigma} + D$,

$$\begin{aligned} \text{rank}(\hat{\Sigma}) &= \text{trace}(\hat{\Sigma}^\dagger \hat{\Sigma}) \\ &\geq \text{trace}((\hat{\Sigma} + D)^{-1} \hat{\Sigma}) \\ &= \text{trace}(\Sigma^{-1} \hat{\Sigma}), \end{aligned} \quad (5.4)$$

where † denotes the Moore–Penrose pseudoinverse. Continuing with this line of analysis,

$$\begin{aligned} \text{rank}(\hat{\Sigma}) &= \text{trace}(\hat{\Sigma}^\dagger \hat{\Sigma}) \\ &\geq \text{trace}((\hat{\Sigma} + \varepsilon I)^{-1} \hat{\Sigma}) \end{aligned} \quad (5.5)$$

for any $\varepsilon > 0$ suggests the iterative reweighting process

$$D_{(k+1)} := \arg \min_D \text{trace}((\hat{\Sigma} - D_{(k)} + \varepsilon I)^{-1} (\hat{\Sigma} - D)) \quad (5.6)$$

for $k = 1, 2, \dots$ and $D_{(0)} := 0$. In fact, as pointed out in [29], (5.6) corresponds to minimizing $\log \det(\hat{\Sigma} - D + \varepsilon I)$ by local linearization.

Next we provide a sufficient condition for $\hat{\Sigma}$ to be such a stationary point (5.6), i.e., for $\hat{\Sigma}$ to satisfy

$$\arg \min_D \text{trace}((\hat{\Sigma} + \varepsilon I)^{-1} (\hat{\Sigma} - D)) = 0. \quad (5.7)$$

The notation $\circ$ used below denotes the elementwise product between vectors or matrices which is also known as the *Schur product* [41] and, likewise, for vectors $a, b \in \mathbb{R}^{n \times 1}$, $a \circ b \in \mathbb{R}^{n \times 1}$ with $[a \circ b]_i = [a]_i [b]_i$.

Proposition 18—Let $\hat{\Sigma} \in \mathbf{S}_{n,+}$ and let the columns of U form a basis of $\mathcal{R}(\hat{\Sigma})$. If

$$\mathcal{R}(U \circ U) \subset \mathcal{R}(\Pi_{\mathcal{N}(\hat{\Sigma})} \circ \Pi_{\mathcal{N}(\hat{\Sigma})}), \quad (5.8)$$

then $\hat{\Sigma}$ satisfies (5.7) for all $\varepsilon \in (0, \varepsilon_1)$ and some $\varepsilon_1 > 0$.

We first need the following result which generalizes [70, Theorem 3.1].

Lemma 19—For $A \in \mathbb{R}^{n \times p}$ and $B \in \mathbb{R}^{n \times q}$ having columns $a_1, \dots, a_p$ and $b_1, \dots, b_q$ respectively, we let

$$\begin{aligned} C &= [a_1 \circ b_1, a_1 \circ b_2, \dots, a_2 \circ b_1, \dots, a_p \circ b_q] \in \mathbb{R}^{n \times pq}, \\ \phi: \mathbb{R}^n &\rightarrow \mathbb{R}^n \quad d \mapsto \text{diag}(AA' \text{diag}^*(d)BB'), \text{ and} \\ \psi: \mathbb{R}^{p \times q} &\rightarrow \mathbb{R}^n \quad \Delta \mapsto \text{diag}(A\Delta B'). \end{aligned}$$

Then $\mathcal{R}(\phi) = \mathcal{R}(\psi) = \mathcal{R}((AA') \circ (BB')) = \mathcal{R}(C)$.

Proof: Since $\text{diag}(AA' \text{diag}^*(d)BB') = ((AA') \circ (BB'))d$, it follows that

$$\mathcal{R}(\phi) = \mathcal{R}((AA') \circ (BB')).$$

Moreover, $\text{diag}(A\Delta B') = \sum_{i=1}^p \sum_{j=1}^q a_i \circ b_j [\Delta]_{ij}$, and then $\mathcal{R}(\psi) = \mathcal{R}(C)$. We only need to show that $\mathcal{R}(C) = \mathcal{R}((AA') \circ (BB'))$. This follows from

$$\begin{aligned} (AA') \circ (BB') &= \sum_{i=1}^p \sum_{j=1}^q (a_i a_i') \circ (b_j b_j') \\ &= \sum_{i=1}^p \sum_{j=1}^q (a_i \circ b_j)(a_i \circ b_j)' = CC'. \end{aligned}$$

Thus $\mathcal{R}(C) = \mathcal{R}((AA') \circ (BB'))$.

Proof of Proposition 18—Assume that $\hat{\Sigma}$ satisfies (5.7). If $\text{rank}(\hat{\Sigma}) = r$, let $\hat{\Sigma} = USU'$ be the eigendecomposition of $\hat{\Sigma}$ with $S = \text{diag}^*(s)$ with $s \in \mathbb{R}^r$. Let the columns of V be an orthogonal basis of the null space of $\hat{\Sigma}$, i.e., $\Pi_{\mathcal{N}(\hat{\Sigma})} = VV'$. Then

$$(\hat{\Sigma} + \varepsilon I)^{-1} = (\hat{\Sigma} + \varepsilon \Pi_{\mathcal{R}(\hat{\Sigma})} + \varepsilon \Pi_{\mathcal{N}(\hat{\Sigma})})^{-1} = (\hat{\Sigma} + \varepsilon \Pi_{\mathcal{R}(\hat{\Sigma})})^{\dagger} + \frac{1}{\varepsilon} \Pi_{\mathcal{N}(\hat{\Sigma})}$$

and

$$\arg \min_{D: \hat{\Sigma} \geq D} \text{trace} \left((\hat{\Sigma} + \varepsilon I)^{-1} (\hat{\Sigma} - D) \right) = \arg \min_{D: \hat{\Sigma} \geq D} \text{trace} \left(\left(\varepsilon (\hat{\Sigma} + \varepsilon \Pi_{\mathcal{R}(\hat{\Sigma})})^{\dagger} + \Pi_{\mathcal{N}(\hat{\Sigma})} \right) (\hat{\Sigma} - D) \right).$$

From Proposition 17, (5.7) holds if there is $M \in \mathbf{S}_{n-r,+}$ such that

$$\text{diag}(VMV') = \text{diag} \left(\varepsilon (\hat{\Sigma} + \varepsilon \Pi_{\mathcal{R}(\hat{\Sigma})})^{\dagger} + \Pi_{\mathcal{N}(\hat{\Sigma})} \right). \quad (5.9)$$

Obviously, if $\varepsilon = 0$, $M = I$ satisfies the above equation. We consider the matrix M of the form $M = I + \Delta$. When (5.9) holds, we need $\text{diag}((\hat{\Sigma} + \varepsilon \Pi_{\mathcal{R}})^{\dagger})$ to be in the range of ψ for

$$\psi: \mathbf{S}_n \rightarrow \mathbb{R}^n \quad \Delta \mapsto \text{diag}(V\Delta V').$$

From Lemma 19, $\mathcal{R}(\psi) = \mathcal{R}(\Pi_{\mathcal{N}(\hat{\Sigma})} \circ \Pi_{\mathcal{N}(\hat{\Sigma})})$. On the other hand, since

$$\varepsilon (\hat{\Sigma} + \varepsilon \Pi_{\mathcal{R}(\hat{\Sigma})})^{\dagger} = U \text{diag} \left(\left[\frac{\varepsilon}{[s]_1 + \varepsilon}, \dots, \frac{\varepsilon}{[s]_r + \varepsilon} \right] \right) U',$$

then $\text{diag}(\varepsilon (\hat{\Sigma} + \varepsilon \Pi_{\mathcal{R}(\hat{\Sigma})})^{\dagger}) \in \mathcal{R}(U \circ U)$. So if (5.8) holds, there is always a Δ such that $M = I + \Delta$ satisfies (5.9). Moreover, it is also required that $I + \Delta \geq 0$. Since the map from ε to Δ is

continuous, for small enough ε , i.e., in an interval $(0, \varepsilon_1)$, the condition $I+$ can always be satisfied.

We note that (5.8) is a sufficient condition for $\hat{\Sigma}$ to be a stationary point of (5.7) in both Frisch's and Shapiro's settings.

6. Certificates of Minimum Rank

It is of interest to know whether the attained rank of $\hat{\Sigma}$ which in turn is obtained by numerical optimization (e.g., using (5.6)) is minimal. Thus, in this section we are interested in determining bounds on $\text{mr}_+(\Sigma)$. These bounds, when met, can serve as "certificates" to ascertain minimality of rank.

The following two bounds were proposed in [78] and follow from Theorem 2. However, both of these bounds require exhaustive search, which may be prohibitively expensive when n is large.

Corollary 20—*Let $\Sigma \in \mathbf{S}_{n,+}$ and $\Sigma > 0$. If there is an $s_1 \times s_1$ principle minor of Σ whose inverse is positive, then*

$$\text{mr}_+(\Sigma) \geq s_1 - 1. \quad (6.1a)$$

If there is an $s_2 \times s_2$ principle minor of Σ^{-1} which is elementwise positive, then

$$\text{mr}_+(\Sigma) \geq s_2 - 1. \quad (6.1b)$$

Other criteria that also require exhaustive search are given in [6] for the special case when $\text{mr}_+(\Sigma) \leq \frac{n+1}{2}$. Next we discuss three bounds that are computationally tractable; the first two were proposed by Guttman [37]. Guttman's bounds are based on a conservative assessment for the admissible range of each of the diagonal entries of $D = \Sigma - \hat{\Sigma}$.

Proposition 21—*Let $\Sigma \in \mathbf{S}_{n,+}$ and let*

$$\begin{aligned} D_1 &:= \text{diag}^*(\text{diag}(\Sigma)), \\ D_2 &:= (\text{diag}^*(\text{diag}(\Sigma^{-1})))^{-1}. \end{aligned}$$

Then the following hold:

$$\text{mr}_+(\Sigma) \geq n_+(\Sigma - D_1), \quad (6.1c)$$

$$\text{mr}_+(\Sigma) \geq n_+(\Sigma - D_2). \quad (6.1d)$$

Further, $n_+(\Sigma - D_1) \leq n_+(\Sigma - D_2)$.

Proof: The proof follows from the fact that $\Sigma \succeq D$ implies $D \succeq D_2 \succeq D_1$. See [37] for details.

It is also easy to see that $\text{mr}(\Sigma) \geq n_+(\Sigma - D_1)$, which provides a lower bound for the minimum rank in Shapiro's problem. Next we return to a bound noted earlier in (5.4).

Proposition 22—Let $\Sigma \in \mathbf{S}_{n,+}$. Then the following holds:

$$\text{mr}_+(\Sigma) \geq \min_{\Sigma \succeq D \succeq 0} \text{trace}(\Sigma^{-1}(\Sigma - D)). \quad (6.1e)$$

Proof: The statement follows readily from (5.4).

Evidently an analogous statement holds for $\text{mr}(\Sigma)$. We note that (6.1c) and (6.1d) remain invariant under scaling of rows and corresponding columns, whereas (6.1e) does not; hence these two cannot be compared directly.

7. Correspondence between Decompositions

We now return to the decomposition of the data matrix $X = \hat{X} + \tilde{X}$ and its relation to the corresponding sample covariances. The decomposition of X into “noise-free” and “noisy” components implies a corresponding decomposition for the sample covariance, but in the converse direction, a decomposition $\Sigma = \hat{\Sigma} + \tilde{\Sigma}$ leads to a family of compatible decompositions for X , which corresponds to the boundary of a matrix ball. This is discussed next.

Proposition 23—Let $X \in \mathbb{R}^{n \times T}$ and $\Sigma := XX'$. If

$$\Sigma = \hat{\Sigma} + \tilde{\Sigma}, \quad (7.1)$$

with $\hat{\Sigma}, \tilde{\Sigma}$ symmetric and nonnegative definite, there exists a decomposition

$$X = \hat{X} + \tilde{X} \quad (7.2a)$$

for which

$$\hat{X}\tilde{X}' = 0, \quad (7.2b)$$

$$\hat{\Sigma} = \hat{X}\hat{X}' \quad (7.2c)$$

$$\tilde{\Sigma} = \tilde{X}\tilde{X}'. \quad (7.2d)$$

Further, all pairs $(\hat{X}, \tilde{X})$ that satisfy (7.2a)–(7.2d) are of the form

$$\hat{X} = \hat{\Sigma} \hat{\Sigma}^{-1} X + R^{1/2} V, \quad \tilde{X} = \tilde{\Sigma} \tilde{\Sigma}^{-1} X - R^{1/2} V, \quad (7.3)$$

with

$$R := \hat{\Sigma} - \hat{\Sigma} \hat{\Sigma}^{-1} \tilde{\Sigma} \quad (7.4a)$$

$$\begin{aligned} &= \tilde{\Sigma} - \tilde{\Sigma} \tilde{\Sigma}^{-1} \hat{\Sigma} \\ &= \tilde{\Sigma} \tilde{\Sigma}^{-1} \hat{\Sigma} \\ &= \tilde{\Sigma} \tilde{\Sigma}^{-1} \hat{\Sigma}, \end{aligned} \quad (7.4b)$$

and $V \in \mathbb{R}^{n \times T}$ such that $VV' = I$, $XV = 0$.

Proof: The proof relies on a standard lemma [24, Theorem 2] which states that if $A \in \mathbb{R}^{n \times T}$, $B \in \mathbb{R}^{n \times m}$ with $m \leq T$ such that $AA' = BB'$, then $A = BU$ for some $U \in \mathbb{R}^{m \times T}$ with $UU' = I$. Thus, we let $A := X$,

$$S := \begin{bmatrix} \hat{\Sigma} & 0 \\ 0 & \tilde{\Sigma} \end{bmatrix},$$

and $B := \begin{bmatrix} I & I \end{bmatrix} S^{1/2}$, where $S^{1/2}$ is the matrix-square root of S . It follows that there exists a matrix U as above for which $A = BU$, and therefore we can take

$$\begin{bmatrix} \hat{X} \\ \tilde{X} \end{bmatrix} := S^{1/2} U.$$

This establishes the existence of the decomposition (7.2a).

In order to parameterize all such pairs $(\hat{X}, \tilde{X})$, let U_o be an orthogonal (square) matrix such that

$$XU_o = \begin{bmatrix} \Sigma^{1/2} & 0 \end{bmatrix}.$$

Then $X\hat{U}_o$ and $X\tilde{U}_o$ must be of the form

$$\hat{X}U_o = \begin{bmatrix} \hat{X}_1 & \Delta \end{bmatrix}, \quad \tilde{X}U_o = \begin{bmatrix} \tilde{X}_1 & -\Delta \end{bmatrix}, \quad (7.5)$$

with $\hat{X}_1, \tilde{X}_1$ square matrices. Since

$$\begin{bmatrix} \hat{X} \\ \tilde{X} \end{bmatrix} \begin{bmatrix} \hat{X}' & \tilde{X}' \end{bmatrix} = \begin{bmatrix} \hat{\Sigma} & 0 \\ 0 & \tilde{\Sigma} \end{bmatrix},$$

then

$$\hat{X}_1 \hat{X}_1' + \Delta \Delta' = \hat{\Sigma}, \quad (7.6a)$$

$$\hat{X}_1 \tilde{X}_1' - \Delta \Delta' = 0, \quad (7.6b)$$

$$\tilde{X}_1 \tilde{X}_1' + \Delta \Delta' = \tilde{\Sigma}. \quad (7.6c)$$

Substituting $\hat{X}_1 \tilde{X}_1'$ for $\Delta \Delta'$ into (7.6a) and using the fact that $\tilde{X}_1 = X_1 - \hat{X}_1$ with $X_1 = \Sigma^{1/2}$, we obtain that

$$\hat{X}_1 = \hat{\Sigma} \Sigma^{-1/2}.$$

Similarly, using (7.6c) instead, we obtain that

$$\tilde{X}_1 = \tilde{\Sigma} \Sigma^{-1/2}.$$

Substituting into (7.6b), (7.6a), and (7.6c), we obtain the following three relations:

$$\begin{aligned} \Delta \Delta' &= \hat{\Sigma} \Sigma^{-1} \tilde{\Sigma} \\ &= \hat{\Sigma} - \hat{\Sigma} \Sigma^{-1} \hat{\Sigma} \\ &= \tilde{\Sigma} - \tilde{\Sigma} \Sigma^{-1} \tilde{\Sigma}. \end{aligned}$$

Since $\Delta \Delta'$ and the Σ 's are all symmetric,

$$\Delta \Delta' = \tilde{\Sigma} \Sigma^{-1} \hat{\Sigma}$$

as well. Thus, $\Delta = R^{1/2} V_1$ with $V_1 V_1' = I$. The proof is completed by substituting the expressions for $\hat{X}_1$ and $\tilde{X}_1$ into (7.5).

Interestingly,

$$\text{rank}(R) + \text{rank}(\hat{\Sigma}) = \text{rank} \left(\begin{bmatrix} \hat{\Sigma} & \hat{\Sigma} \\ \hat{\Sigma} & \Sigma \end{bmatrix} \right) = \text{rank} \left(\begin{bmatrix} \hat{\Sigma} & 0 \\ 0 & \tilde{\Sigma} \end{bmatrix} \right) = \text{rank}(\hat{\Sigma}) + \text{rank}(\tilde{\Sigma}),$$

and hence the rank of the “uncertainty radius” R of the corresponding $\hat{X}$ and $\tilde{X}$ matrix spheres is

$$\text{rank}(R) = \text{rank}(\hat{\Sigma}) + \text{rank}(\tilde{\Sigma}) - \text{rank}(\Sigma).$$

In cases where one identifies $\hat{X}$ from the data matrix X , different criteria may be used to quantify uncertainty. One such case is the rank of R , while another is its trace, which is the variance of estimation error in determining X . This topic is considered next, and its relation to the Frisch decomposition is highlighted.

8. Uncertainty and Worst-Case Estimation

The basic premise of the decomposition (7.1) is that, in principle, no probabilistic description of the data is needed. Thus, under the assumptions of Proposition 23, R represents a deterministic radius of uncertainty when interpreting the data. On the other hand, if “signal” and “noise” are assumed probabilistic in nature and jointly Gaussian, assuming that $[\hat{\mathbf{x}}, \mathbf{x}]'$ has zero mean and covariance

$$\begin{bmatrix} \hat{\Sigma} & \hat{\Sigma} \\ \hat{\Sigma} & \Sigma \end{bmatrix},$$

the conditional expectation of the noise-free component $\hat{\mathbf{x}}$ given observations $\mathbf{x}$ is $E\{\hat{\mathbf{x}}|\mathbf{x}\} = \hat{\Sigma}\hat{\Sigma}^{-1}\mathbf{x}$ and the variance of the estimation error is

$$\mathcal{E}\{(\hat{\mathbf{x}} - \hat{\Sigma}\hat{\Sigma}^{-1}\mathbf{x})(\hat{\mathbf{x}} - \hat{\Sigma}\hat{\Sigma}^{-1}\mathbf{x})'\} = \hat{\Sigma} - \hat{\Sigma}\hat{\Sigma}^{-1}\hat{\Sigma} = R,$$

which is precisely the radius of the deterministic uncertainty set; cf. [26, p. 116]. Either way, it is of interest to assess how this radius depends on the decomposition of Σ , and we discuss this next.

8.1. Uniformly Optimal Decomposition

The uncertainty radius R relates to the minimum mean-squared error (MMSE) estimation. Thus, suppose we seek to estimate $\hat{\mathbf{x}}$ using a linear estimator $K\mathbf{x}$ for $K \in \mathbb{R}^{n \times n}$. The mean-squared error (MSE) corresponds to the expression given by

$$\mathcal{E}\{\|\hat{\mathbf{x}} - K\mathbf{x}\|^2\} = \text{trace} \left(\hat{\Sigma} - K\hat{\Sigma} - \hat{\Sigma}K' + K\hat{\Sigma}K' \right).$$

If Σ , $\hat{\Sigma}$ are known and Σ is positive definite, then the optimal estimator corresponds to $K = \hat{\Sigma}\hat{\Sigma}^{-1}$ and, as noted earlier, the covariance matrix of the estimation error is precisely R .

Considering that $\hat{\Sigma}$ is unknown as in the Frisch problem, the MSE also depends on the decomposition of Σ . Thus, it is natural to consider the MSE loss function

$$L(K, \hat{\Sigma}, \tilde{\Sigma}) := \text{trace} \left(\hat{\Sigma} - K\hat{\Sigma} - \hat{\Sigma}K' + K(\hat{\Sigma} + \tilde{\Sigma})K' \right) \quad (8.1)$$

and seek an estimator that is uniformly optimal for all admissible pairs $(\hat{\Sigma}, \tilde{\Sigma})$ in the uncertain set

$$\mathcal{S}(\Sigma) := \{(\hat{\Sigma}, \tilde{\Sigma}) : \Sigma = \hat{\Sigma} + \tilde{\Sigma}, \hat{\Sigma}, \tilde{\Sigma} \geq 0 \text{ and } \tilde{\Sigma} \text{ is diagonal}\}.$$

This leads to the following min-max problem:

$$\min_K \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} L(K, \hat{\Sigma}, \tilde{\Sigma}). \quad (8.2)$$

Analogous min-max formulations for a variety of assumptions on the uncertainty sets have been widely used as a robust technique in signal processing [48, 77, 47, 51, 27]. Below, we focus on the relation between the optimal estimator and the corresponding decomposition of Σ as in the Frisch problem.

By switching the order of min and max in (8.2), we obtain

$$\min_K \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} L(K, \hat{\Sigma}, \tilde{\Sigma}) \geq \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \min_K L(K, \hat{\Sigma}, \tilde{\Sigma}) \quad (8.3a)$$

$$= \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \text{trace} \left(\hat{\Sigma} - \hat{\Sigma}\hat{\Sigma}^{-1}\hat{\Sigma} \right) \quad (8.3b)$$

$$= \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \text{trace} \left(\tilde{\Sigma} - \tilde{\Sigma}\hat{\Sigma}^{-1}\tilde{\Sigma} \right). \quad (8.3c)$$

Here, in (8.3b), we have used that the optimal K for a pair $(\hat{\Sigma}, \tilde{\Sigma})$ is $K = \hat{\Sigma}\hat{\Sigma}^{-1} = I - \hat{\Sigma}^{-1}\tilde{\Sigma}$. The functions to maximize in (8.3b) and (8.3c) are both strictly concave in $\hat{\Sigma}$ and $\tilde{\Sigma}$. Therefore the maximizer is unique. Thus, we denote

$$(K_{\text{opt}}, \hat{\Sigma}_{\text{opt}}, \tilde{\Sigma}_{\text{opt}}) := \arg \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \min_K L(K, \hat{\Sigma}, \tilde{\Sigma}), \quad (8.4)$$

where $K_{\text{opt}} = \hat{\Sigma}_{\text{opt}} \Sigma^{-1}$, and $K_{\text{opt}} \mathbf{x}$ is viewed as an estimate of $\hat{\mathbf{x}}$.

In general, the decomposition suggested by the uniformly optimal estimation problem does not lead to a singular signal covariance $\hat{\Sigma}$. The condition for when that happens is given next. Interestingly, the condition is expressed in terms of the candidate noise covariance in one of the Guttman bounds (Proposition 21).

Proposition 24—Let $\Sigma > 0$, and let

$$D_0 := \frac{1}{2} \text{diag}^*(\text{diag}(\Sigma^{-1}))^{-1} \quad (8.5)$$

(which is equal to $\frac{1}{2} D_2$ defined in Proposition 21). If $\Sigma - D_0 \geq 0$, then

$$\tilde{\Sigma}_{\text{opt}} = D_0 \text{ and } \hat{\Sigma}_{\text{opt}} = \Sigma - D_0. \quad (8.6a)$$

Otherwise,

$$\tilde{\Sigma}_{\text{opt}} \leq D_0 \text{ and } \hat{\Sigma}_{\text{opt}} \text{ is singular}. \quad (8.6b)$$

Proof: From (8.3c),

$$\begin{aligned} L(K_{\text{opt}}, \hat{\Sigma}_{\text{opt}}, \tilde{\Sigma}_{\text{opt}}) &= \max \left\{ \tilde{\Sigma} - \tilde{\Sigma} \Sigma^{-1} \tilde{\Sigma} \mid \Sigma \geq \tilde{\Sigma} \geq 0, \tilde{\Sigma} \text{ is diagonal} \right\} \\ &\leq \max \left\{ \tilde{\Sigma} - \tilde{\Sigma} \Sigma^{-1} \tilde{\Sigma} \mid \tilde{\Sigma} \text{ is diagonal} \right\} \\ &= \frac{1}{2} \text{trace}(D_0), \end{aligned} \quad (8.7)$$

with the maximum attained for $\tilde{\Sigma} = D_0$. Then (8.6a) follows. In order to prove (8.6b), consider the Lagrangian corresponding to (8.3c),

$$\mathcal{L}(\tilde{\Sigma}, \Lambda_0, \Lambda_1) = \text{trace}(\tilde{\Sigma} - \tilde{\Sigma} \Sigma^{-1} \tilde{\Sigma} + \Lambda_0(\Sigma - \tilde{\Sigma}) + \Lambda_1 \tilde{\Sigma}),$$

where Λ_0, Λ_1 are Lagrange multipliers. The optimal values satisfy

$$[I - 2 \Sigma^{-1} \tilde{\Sigma}_{\text{opt}} - \Lambda_0 + \Lambda_1]_{kk} = 0 \quad \forall k=1, \dots, n, \quad (8.8a)$$

$$\Lambda_0 \hat{\Sigma}_{\text{opt}} = 0, \Lambda_0 \geq 0, \quad (8.8b)$$

$$\Lambda_1 \tilde{\Sigma}_{\text{opt}} = 0, \Lambda_1 \geq 0 \text{ and is diagonal.} \quad (8.8c)$$

If $\Sigma - D_0 \not\geq 0$, we show that $\hat{\Sigma}_{\text{opt}}$ is singular. Assume the contrary, i.e., that $\hat{\Sigma}_{\text{opt}} > 0$. From (8.8b), we see that $\Lambda_0 = 0$, while from (8.8a), $[I - 2\Sigma^{-1}\tilde{\Sigma}_{\text{opt}}]_{kk} = 0$. This gives that

$$[\tilde{\Sigma}_{\text{opt}}]_{kk} \geq \frac{1}{2[\Sigma^{-1}]_{kk}} = [D_0]_{kk}$$

for all $k = 1, \dots, n$, which contradicts the fact that $\Sigma - D_0 \not\geq 0$. Therefore $\hat{\Sigma}_{\text{opt}}$ is singular. We now assume that $\Sigma \not\geq D_0$. Then there exists k such that $[\Sigma_{\text{opt}}]_{kk} > [D_0]_{kk}$. From (8.8c) and (8.8a), we have that

$$[\Lambda_1]_{kk} = 0 \text{ and } [I - 2\Sigma^{-1}\tilde{\Sigma}_{\text{opt}}]_{kk} \geq 0,$$

which contradicts the assumption that $[\Sigma_{\text{opt}}]_{kk} > [D_0]_{kk}$. Therefore $\hat{\Sigma}_{\text{opt}} = D_0$ and (8.6b) has been established.

We remark that while the matrix-valued error variance

$$\mathcal{E} \{(\hat{\mathbf{x}} - K\mathbf{x})(\hat{\mathbf{x}} - K\mathbf{x})'\} = \hat{\Sigma} - K\hat{\Sigma} - \hat{\Sigma}K' + K\hat{\Sigma}K'$$

is matrix-convex in K and has a unique minimum for $K = \hat{\Sigma}\hat{\Sigma}^{-1}$ for which the optimal value equals $\hat{\Sigma} - \hat{\Sigma}\hat{\Sigma}^{-1}\hat{\Sigma}$, the optimal error covariance $\hat{\Sigma} - \hat{\Sigma}\hat{\Sigma}^{-1}\hat{\Sigma}$ as a function of $\hat{\Sigma}$ may not have

a unique maximum in the positive semidefinite sense. To see this, consider $\hat{\Sigma} = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$

In this case $D_0 = \frac{3}{4}I$, $\hat{\Sigma}_{\text{opt}} = \begin{bmatrix} 5/4 & 1 \\ 1 & 5/4 \end{bmatrix}$, and

$$\hat{\Sigma}_{\text{opt}} - \hat{\Sigma}_{\text{opt}}\hat{\Sigma}^{-1}\hat{\Sigma}_{\text{opt}} = \begin{bmatrix} 3/8 & 3/16 \\ 3/16 & 3/8 \end{bmatrix}. \quad (8.9)$$

On the other hand, for $\hat{\Sigma} = \begin{bmatrix} 3/2 & 1 \\ 1 & 3/2 \end{bmatrix}$, then

$$\hat{\Sigma} - \hat{\Sigma} \hat{\Sigma}^{-1} \hat{\Sigma} = \begin{bmatrix} 1/3 & 1/12 \\ 1/12 & 1/3 \end{bmatrix},$$

which is neither larger nor smaller than (8.9) in the sense of semidefiniteness. This is a key reason for considering scalar loss functions of the error covariance as in (8.1).

Next we note that there is no gap between the min-max and max-min values in the two sides of (8.3a).

Proposition 25—For $\Sigma \in \mathbf{S}_{n,+}$, then

$$\min_K \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} L(K, \hat{\Sigma}, \tilde{\Sigma}) = \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \min_K L(K, \hat{\Sigma}, \tilde{\Sigma}). \quad (8.10)$$

Proof: We observe that for a fixed K , the function $L(K, \hat{\Sigma}, \tilde{\Sigma})$ is a linear function of $(\hat{\Sigma}, \tilde{\Sigma})$. For fixed $(\hat{\Sigma}, \tilde{\Sigma})$, the function is a convex function of K . Under these conditions it is standard that (8.10) holds; see, e.g., [13, p. 281].

It is interesting to note that when $D_0 = \frac{1}{2} \text{diag}^*(\text{diag}(\Sigma^{-1}))^{-1}$ is admissible as noise covariance, i.e., $\Sigma - D_0 \succeq 0$, the optimal signal covariance is $\hat{\Sigma}_{\text{opt}} = \Sigma - D_0$, and the gain matrix $K_{\text{opt}} = \hat{\Sigma}_{\text{opt}} \Sigma^{-1} = I - D_0 \Sigma^{-1}$ has all diagonal entries equal to $\frac{1}{2}$. Thus, with K_{opt} in (8.1) the MSE loss is independent of $\hat{\Sigma}$ and equal to $\text{trace}(2D_0 - \Sigma + K_{\text{opt}} \Sigma K_{\text{opt}}')$ for any admissible decomposition of Σ . We also point out that the key condition (Proposition 24)

$$\begin{aligned} \Sigma &\geq \frac{1}{2} \text{diag}^*(\text{diag}(\Sigma^{-1}))^{-1} \\ \iff 2 \text{diag}^*(\text{diag}(\Sigma^{-1})) &\geq \Sigma^{-1} \end{aligned}$$

can be equivalently written as $\Sigma^{-1} \circ (2I - \mathbf{1}\mathbf{1}') \succeq 0$ and, interestingly, amounts to the positive semidefiniteness of a matrix formed by changing the signs of all off-diagonal entries of Σ^{-1} . The set of all such matrices, $\{S \mid S \succeq 0, S \circ (2I - \mathbf{1}\mathbf{1}') \succeq 0\}$, is convex, invariant under scaling rows and corresponding columns, and contains the set of diagonally dominant matrices $\{S \mid S \succeq 0, [S]_{ii} \geq \sum_j [S]_{ij} \text{ for all } i\}$.

We conclude this section by showing that $\text{trace}(R_{\text{opt}})$, with

$$R_{\text{opt}} := \hat{\Sigma}_{\text{opt}} - \hat{\Sigma}_{\text{opt}} \hat{\Sigma}^{-1} \hat{\Sigma}_{\text{opt}},$$

represents the matrix radius confining all admissible matrices $\hat{\Sigma}$ about $\hat{\Sigma}_{\text{opt}}$. Thus, R quantifies the distance between admissible decompositions of Σ as stated next.

Proposition 26—For $\Sigma > 0$ and any pair $(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)$,

$$\text{trace} \left((\hat{\Sigma} - \hat{\Sigma}_{\text{opt}}) \Sigma^{-1} (\hat{\Sigma} - \hat{\Sigma}_{\text{opt}})' \right) \leq \text{trace}(R_{\text{opt}}).$$

Proof: Clearly $\text{trace}(\Sigma \hat{\Sigma}^{-1} \Sigma) = \text{trace}(\hat{\Sigma})$, while from Proposition 25,

$$L(K_{\text{opt}}, \hat{\Sigma}, \tilde{\Sigma}) = \text{trace}(\hat{\Sigma} - 2\hat{\Sigma}_{\text{opt}} \Sigma^{-1} \hat{\Sigma} + \hat{\Sigma}_{\text{opt}} \Sigma^{-1} \hat{\Sigma}_{\text{opt}}') \leq \text{trace}(R_{\text{opt}}). \quad (8.11)$$

$$\text{Thus, } \text{trace}(\hat{\Sigma} \Sigma^{-1} \hat{\Sigma} - 2\hat{\Sigma}_{\text{opt}} \Sigma^{-1} \hat{\Sigma} + \hat{\Sigma}_{\text{opt}} \Sigma^{-1} \hat{\Sigma}_{\text{opt}}') \leq \text{trace}(R_{\text{opt}}).$$

8.2. Uniformly Optimal Estimation and Trace Regularization

The decomposition of Σ in accordance with the min-max estimation problem of the previous section often produces an invertible signal covariance $\hat{\Sigma}$. On the other hand, as argued earlier, it is often desirable to obtain a decomposition where $\hat{\Sigma}$ is singular and of relatively low rank. Thus, following [28, 14, 59, 65, 16, 1], we may combine the MSE loss function with a regularization term that promotes low rank for the signal covariance $\hat{\Sigma}$. A corresponding problem is of the form

$$J := \min_K \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \left(L(K, \hat{\Sigma}, \tilde{\Sigma}) - \lambda \cdot \text{trace}(\hat{\Sigma}) \right), \quad (8.12)$$

with $\lambda \geq 0$.

As noted in Proposition 25 (see [13, p. 281]), here too there is no gap between the min-max and the max-min, and therefore

$$\begin{aligned} \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \min_K \{L(K, \hat{\Sigma}, \tilde{\Sigma}) - \lambda \text{trace}(\hat{\Sigma})\} &= \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \min_K \text{trace} \left((1-\lambda)\hat{\Sigma} - K\hat{\Sigma} - \hat{\Sigma}K' + K(\hat{\Sigma} + \tilde{\Sigma})K' \right) \\ &= \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \text{trace} \left((1-\lambda)\hat{\Sigma} - \hat{\Sigma}(\hat{\Sigma} + \tilde{\Sigma})^{-1} \hat{\Sigma} \right) \end{aligned} \quad (8.13a)$$

$$= \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{S}(\Sigma)} \text{trace} \left(-\lambda \Sigma + (1+\lambda) \hat{\Sigma} - \tilde{\Sigma}(\hat{\Sigma} + \tilde{\Sigma})^{-1} \tilde{\Sigma} \right). \quad (8.13b)$$

Note that (8.13a) is obtained by substituting $\Sigma(\hat{\Sigma} + \tilde{\Sigma})^{-1}$ for K . Since (8.13a) and (8.13b) are strictly concave functions of $\hat{\Sigma}$ and $\tilde{\Sigma}$, respectively, there is a unique set of optimal values $(K_{\lambda, \text{opt}}, \hat{\Sigma}_{\lambda, \text{opt}}, \tilde{\Sigma}_{\lambda, \text{opt}})$. For sufficiently large λ the optimal value for $\hat{\Sigma}$ is singular. The following provides one such condition.

Proposition 27—For $\Sigma > 0$, let

$$D_0 := \frac{1}{2}(\text{diag}^* \text{diag}(\Sigma^{-1}))^{-1},$$

let $\lambda_{\min}$ be the smallest eigenvalue of $D_0^{-\frac{1}{2}} \Sigma D_0^{-\frac{1}{2}}$, and let $(K_{\lambda, \text{opt}}, \hat{\Sigma}_{\lambda, \text{opt}}, \tilde{\Sigma}_{\lambda, \text{opt}})$ as defined above. If $\lambda = \lambda_{\min} - 1$, then $\hat{\Sigma}_{\lambda, \text{opt}}$ is singular.

Proof: The trace of $(-\lambda \Sigma + (1+\lambda) \hat{\Sigma} - \tilde{\Sigma} \tilde{\Sigma}^{-1} \tilde{\Sigma})$ is maximal for the diagonal choice $\tilde{\Sigma} = (1+\lambda) D_0$. For any $\lambda = \lambda_{\min} - 1$, $\tilde{\Sigma} - (1+\lambda) D_0$ fails to be positive semidefinite. Thus, the constraint $\tilde{\Sigma} - \Sigma \succeq 0$ in (8.13b) is active and $\hat{\Sigma}_{\lambda, \text{opt}}$ is singular.

Note that $\Sigma - 2D_0 \not\geq 0$ (unless Σ is diagonal), and therefore $\lambda_{\min} < 2$. Hence, for $\lambda = 1$, $\hat{\Sigma}_{\lambda, \text{opt}}$ is singular. When $\lambda \rightarrow 0$ we recover the solution in (8.4), whereas for $\lambda \rightarrow \infty$ we recover the solution in Proposition 15.

9. Accounting for Statistical Errors

The impetus for some of the early work on the Frisch problem was provided by the need for reliable models in econometrics—a subject which was characterized by Ragnar Frisch as the application of statistical and mathematical methods in economics [73, p. 266]. A history of the subject and the pivotal contributions of Frisch, Koopmans, Haavelmo, Anderson, and others, as well as a brief reference to the modern critique by R. E. Kalman, is recounted by Madala [73, pp. 265–283]. Perhaps the key stumbling block, going back to the work of Tintner [74] and an issue often raised by R. E. Kalman, has been the sensitivity of models to statistical uncertainty and the well-posedness of the Frisch problem (which was discussed earlier). In the present section we discuss how we may account for uncertainty in the data in the context of the Frisch paradigm. The section concludes with an example that highlights the difficulties of conclusively ascertaining the number of linear relations when statistical uncertainty is present. The examples help highlight tradeoffs and the relevance of the theory of the Frisch viewpoint in practice.

From an applications standpoint, Σ represents an empirical covariance which is obtained on the basis of the finite observation record in the data matrix X . Hence, the assumptions (3.3a) and (3.3b) as well as the decomposition $\Sigma = \hat{\Sigma} + \tilde{\Sigma}$ can only be expected to hold approximately; the sample covariance Σ may not have a low-rank decomposition even if the underlying true covariance matrix does. Statistical errors can then be taken into account in the min-max formulation (8.2) and (8.12) by replacing the uncertainty set $\mathcal{S}(\Sigma)$ with

$$\mathcal{S}_\varepsilon(\Sigma) := \{(\hat{\Sigma}, \tilde{\Sigma}) \mid d(\hat{\Sigma} + \tilde{\Sigma}, \Sigma) \leq \varepsilon, \text{ with } \hat{\Sigma}, \tilde{\Sigma} \geq 0, \text{ and } \tilde{\Sigma} \text{ diagonal}\}. \quad (9.1)$$

Here $\varepsilon \geq 0$ and $d(\cdot, \cdot)$ represents a distance between $\hat{\Sigma} + \tilde{\Sigma}$ and Σ . One should note that the distance between (sets of) coefficients of linear relations corresponding to different elements in $\mathcal{S}_\varepsilon(\Sigma)$ is on the order of ε [67, section 3.2].

Several candidate distance measures $d(\cdot, \cdot)$ have been used to compare covariance matrices [17, 11, 33, 62, 7, 57, 2]. Examples include the Frobenius norm [57], the affine invariant metric [11, 2, 7], Jensen–Bregman divergences [17], the Wasserstein distance between corresponding distributions [61], and others. While we do not make any claim as to the appropriateness of one distance versus the other, for specificity, we make use of the 2-Wasserstein distance and seek a Gaussian distribution closest to one which agrees with the sample statistics. Below we provide a brief account of the Wasserstein distance between Gaussian distributions. We express this distance as the solution to an optimization problem expressed in terms of the corresponding covariances [61].

Consider two zero-mean Gaussian random vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ having covariance matrices Σ_x and Σ_y , respectively. The squared 2-Wasserstein distance between the respective density functions is defined as

$$\min_{p_{xy}} \mathcal{E}_{p_{xy}}(\|\mathbf{x} - \mathbf{y}\|^2), \quad (9.2)$$

where $\mathcal{E}_{p_{xy}}$ denotes expectation with respect to a joint distribution p_{xy} ; i.e., the minimization is over Gaussian densities p_{xy} that agree with the given marginals specified by their respective covariances Σ_x and Σ_y . If we let $\Sigma_{xy} := \mathcal{E}_{p_{xy}}(\mathbf{x}\mathbf{y}')$ represent the unknown correlation between the two vectors,

$$\mathcal{E}_{p_{xy}}(\|\mathbf{x} - \mathbf{y}\|^2) = \text{trace}(\Sigma_x + \Sigma_y - \Sigma_{xy} - \Sigma_{xy}'),$$

while Σ_{xy} is only constrained to satisfy

$$\begin{bmatrix} \Sigma_x & \Sigma_{xy} \\ \Sigma_{xy}' & \Sigma_y \end{bmatrix} \geq 0.$$

It follows that the minimization in (9.2) can be expressed directly in terms of the corresponding covariance matrices [61]. The optimal value in (9.2) now also represents the square of a metric between the respective covariance matrices Σ_x and Σ_y . We use this distance measure to quantify dissimilarity between $\hat{\Sigma} + \tilde{\Sigma}$ and Σ , defining

$$d_{w_2}(\hat{\Sigma} + \tilde{\Sigma}, \Sigma) = \min_C \left(\text{trace}(\Sigma + \hat{\Sigma} + \tilde{\Sigma} - C - C') \mid \begin{bmatrix} \hat{\Sigma} + \tilde{\Sigma} & C \\ C' & \Sigma \end{bmatrix} \geq 0 \right), \quad (9.3)$$

with the optimization in the form of a semidefinite program.

In order to account for statistical errors, we may now modify the optimization problem (8.12) to include the uncertainty set $S_\epsilon(\Sigma)$ in the place of $S(\Sigma)$; note that in (9.1) we use d_{w_2} as the corresponding distance. Thus, we consider

$$J_\varepsilon := \min_K \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{J}_\varepsilon(\Sigma)} \left(L(K, \hat{\Sigma}, \tilde{\Sigma}) - \lambda \text{trace}(\hat{\Sigma}) \right),$$

which incorporates tradeoffs between the dimension of the signal subspace, MSE loss, and statistical errors. For reasons similar to those in (8.13), the order of max-min and min-max can be reversed. This gives that the optimal value is

$$J_\varepsilon := \max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{J}_\varepsilon(\Sigma)} \text{trace} \left((1-\lambda) \hat{\Sigma} - \hat{\Sigma} (\hat{\Sigma} + \tilde{\Sigma})^{-1} \hat{\Sigma} \right).$$

We reformulate the last optimization problem as follows:

$$\max_{(\hat{\Sigma}, \tilde{\Sigma}) \in \mathcal{J}_\varepsilon(\Sigma), Q} \left\{ \text{trace} \left((1-\lambda) \hat{\Sigma} - Q \right) \mid \begin{bmatrix} Q & \hat{\Sigma} \\ \hat{\Sigma} & \hat{\Sigma} + \tilde{\Sigma} \end{bmatrix} \geq 0 \right\}. \quad (9.4)$$

Putting together (9.3) and (9.4), we obtain that

$$J_\varepsilon = \max_{\hat{\Sigma}, \tilde{\Sigma}, Q, C} \text{trace} \left((1-\lambda) \hat{\Sigma} - Q \right) \quad (9.5a)$$

subject to

$$\begin{bmatrix} Q & \hat{\Sigma} \\ \hat{\Sigma} & \hat{\Sigma} + \tilde{\Sigma} \end{bmatrix} \geq 0, \quad \hat{\Sigma} \geq 0, \quad \tilde{\Sigma} \geq 0, \quad \text{and } \tilde{\Sigma} \text{ diagonal}, \quad (9.5b)$$

$$\text{trace}(\hat{\Sigma} + \tilde{\Sigma} + \hat{\Sigma} - C - C') \leq \varepsilon, \quad \begin{bmatrix} \hat{\Sigma} + \tilde{\Sigma} & C \\ C' & \hat{\Sigma} \end{bmatrix} \geq 0. \quad (9.5c)$$

Note that (9.5c) reflects the constraint $d_{W_2}(\hat{\Sigma} + \tilde{\Sigma}, \Sigma) \leq \varepsilon$.

An alternative approach to obtaining suitable decompositions is to seek pairs $(\hat{\Sigma}, \tilde{\Sigma})$ such that $d(\hat{\Sigma} + \tilde{\Sigma}, \Sigma)$ is minimal while $\tilde{\Sigma}$ has a prespecified rank; likelihood-based methods are based on such an approach [6, 43]. In this case, statistical tests can be used to identify a suitable value for the rank [6]; for high-dimensional observations and under mild assumptions the rank may be estimated using principle component analysis [30, 32].

We conclude with an example to highlight the potential and limitations of the framework. We generate data X in the form

$$X = FV + \tilde{X},$$

where $F \in \mathbb{R}^{n \times r}$, $V \in \mathbb{R}^{r \times T}$, and $X \in \mathbb{R}^{n \times T}$ with $n = 50$, $r = 10$, $T = 100$. The elements of F and V are generated from normal distributions with mean zero and unit covariance. The columns of X are generated from a normal distribution with mean zero and diagonal covariance, itself having (diagonal) entries which are uniformly drawn from interval $[1, 10]$. The matrix $\Sigma = XX'$ is subsequently scaled so that $\text{trace}(\Sigma) = 1$. We determine the optimal value of (9.5a) and tabulate below a typical set of values for the rank of the minimizer $\hat{\Sigma}$ (Table 9.1) as a function of λ and ε . We observe a “plateau” where the rank stabilizes at 10 over a small range of values for ε and λ . Naturally, such a plateau may be taken as an indication of a suitable range of parameters. Although the current setting where a small perturbation in the empirical covariance Σ is allowed, the bounds for the rank in (6.1d) and (6.1e) are still pertinent. In fact, for this example, in 7 out of 10 instances where the rank($\hat{\Sigma}$) = 10, the bound in (6.1d) (computed based on the perturbed covariance $\hat{\Sigma} + \Sigma$) has been tight and is thus a valid certificate. For the same range of parameters, the bound in (6.1e) has been lower than the actual rank of $\hat{\Sigma}$. In general, the bounds in (6.1d) and (6.1e) are not comparable as either one may be tighter than the other.

10. Conclusions

Our aim has been to provide an overview of the theory and techniques that are relevant for the problem of identifying linear relations among variables based on noisy measurements—a classical problem of major importance in the current era of “big data.” Novel numerical tools and increasingly powerful computers have made it possible to successfully treat a number of key issues in this topic in a unified manner.

Thus, the thrust of the paper has been to present and develop in a unified manner key ideas of the theory of noise-in-variables linear modeling including an account of key foundational contributions. In particular, we considered two complementing viewpoints for linear modeling under the assumption of independent noise. Starting from covariance data that are exact and known, we first developed classical results on the Frisch problem, which asks for the maximum number of simultaneous linear relations that is consistent with the data. Our analysis provides a geometric insight into a fundamental theorem of Reiersøl, results on the structure of solutions and the well-posedness of the Frisch and Shapiro problems, bounds on the corresponding minimum rank, and a generalization to complex-valued matrices. We then provide an account of trace minimization heuristics and of iterative reweighting trace minimization as a technique for obtaining solutions of low rank and of computationally tractable lower bounds which can serve as certificates to guarantee that the minimum rank has been achieved. Addressing statistical uncertainty, we discuss min-max estimation problems and how to quantify uncertainty while seeking minimum rank models in a convex optimization setting. In this we integrate various objectives (low-rank, minimal worst-case estimation error) and explain their effectiveness and caveats in a numerical example.

In recent years, techniques such as the ones presented in this work are becoming increasingly important in subjects where one has very large noisy data sets. Typical examples include medical imaging, genomics/proteomics, and finance. It is our hope that the material we presented in this paper will be found useful in this context. It must be noted that throughout the present work we emphasized independence of noise in individual variables.

Evidently, more general and versatile structures for the noise statistics can be treated in a similar manner, and these may become important when dealing with large data sets.

An important topic for future research is that of dealing with errors in estimating empirical statistics. It is common to quantify distances using standard matrix norms— as is done in the present paper. Alternative distance measures such as the Wasserstein distance mentioned in section 9 and others [2, 7, 33, 17, 60]) may become increasingly important. Finally, we raise the question of the asymptotic performance of certificates such as those presented in section 6. It is important to know how the tightness of the certificate to the minimal rank of linear models relates to the size of the problem.

References

1. Agarwal A, Negahban S, Wainwright M. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *Ann Statist.* 2012; 40:1171–1197.
2. Amari, S-I. *Lecture Notes in Statist.* Vol. 28. Springer-Verlag; Berlin: 1985. *Differential-Geometrical Methods in Statistics.*
3. Anderson BDO, Deistler M. Dynamic errors-in-variables systems with three variables. *Automatica J IFAC.* 1987; 23:611–616.
4. Anderson, BDO.; Deistler, M. Generalized linear dynamic factor models: A structure theory. *Proceedings of the 47th IEEE Conference on Decision and Control*; 2008. p. 1980-1985.
5. Anderson TW. Estimating linear statistical relationships. *Ann Statist.* 1984; 12:1–45.
6. Anderson TW, Rubin H. Statistical inference in factor analysis. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability.* 1956; 5:111–150.
7. Barbaresco, F. *Matrix Information Geometry.* Springer; Heidelberg: 2013. Information geometry of covariance matrix: Cartan-Siegel homogeneous bounded domains, Mostow/Berger fibration and Frechet median; p. 199-255.
8. Bartholomew D. The foundations of factor analysis. *Biometrika.* 1984; 71:221–232.
9. Bartholomew, D.; Knott, M.; Moustaki, I. *Wiley Ser Probab Stat.* Vol. 899. John Wiley & Sons; Chichester, UK: 2011. *Latent Variable Models and Factor Analysis: A Unified Approach.*
10. Bekker P, ten Berge J. Generic global identification in factor analysis. *Linear Algebra Appl.* 1997; 264:255–263.
11. Bhatia, R. *Positive Definite Matrices.* Princeton University Press; Princeton, NJ: 2009.
12. Bottegal, G.; Picci, G. A note on generalized factor analysis models. *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*; 2011. p. 1485-1490.
13. Boyd, S.; Vandenberghe, L. *Convex Optimization.* Cambridge University Press; Cambridge, UK: 2004.
14. Candès E, Plan Y. Matrix completion with noise. *Proc IEEE.* 2010; 98:925–936.
15. Candès E, Recht B. Exact matrix completion via convex optimization. *Found Comput Math.* 2009; 9:717–772.
16. Chandrasekaran V, Sanghavi S, Parrilo PA, Willsky AS. Rank-sparsity incoherence for matrix decomposition. *SIAM J Optim.* 2011; 21:572–596.
17. Cherian, A.; Sra, S.; Banerjee, A.; Papanikolopoulos, N. Efficient similarity search for covariance matrices via the Jensen-Bregman LogDet divergence. *IEEE International Conference on Computer Vision*; Washington, DC: IEEE; 2011. p. 2399-2406.
18. Choi J, Zou H, Oehlert G. A penalized maximum likelihood approach to sparse factor analysis. *Stat Interface.* 2011; 3:429–436.
19. Deistler M, Anderson BDO. Linear dynamic errors-in-variables models: Some structure theory. *J Econometrics.* 1989; 41:39–63.
20. Deistler M, Anderson BDO, Filler A, Zinner C, Chen W. Generalized linear dynamic factor models: An approach via singular autoregressions. *Eur J Control.* 2010; 16:211–224.

21. Deistler, M.; Scherrer, W. *New Directions in Time Series Analysis*. Springer; New York: 1993. Identification of linear systems from noisy data; p. 21–42.
22. Della Riccia G, Shapiro A. Minimum rank and minimum trace of covariance matrices. *Psychometrika*. 1982; 47:443–448.
23. Dempster A. Covariance selection. *Biometrics*. 1972; 28:157–175.
24. Douglas R. On majorization, factorization, and range inclusion of operators on Hilbert space. *Proc Amer Math Soc*. 1966; 17:413–415.
25. Durbin J. Errors in variables. *Rev Inst Internat Statist*. 1954; 22:23–32.
26. Eaton, M. *Multivariate Statistics: A Vector Space Approach*. Wiley; New York: 1983.
27. Eldar Y, Merhav N. A competitive minimax approach to robust estimation of random parameters. *IEEE Trans Signal Process*. 2004; 52:1931–1946.
28. Fazel M, Hindi H, Boyd S. A rank minimization heuristic with application to minimum order system approximation. *Proceedings of the 2001 American Control Conference*. 2001; 6:4734–4739.
29. Fazel M, Hindi H, Boyd S. Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices. *Proceedings of the 2003 American Control Conference*. 2003; 3:2156–2162.
30. Forni M, Hallin M, Lippi M, Reichlin L. The generalized dynamic-factor model: Identification and estimation. *Rev Econom Statist*. 2000; 82:540–554.
31. Forni M, Hallin M, Lippi M, Zaffaroni P. Dynamic factor models with infinite-dimensional factor spaces: One-sided representations. *J Econometrics*. 2015; 185:359–371.
32. Forni M, Lippi M. The general dynamic factor model: One-sided representation results. *J Econometrics*. 2011; 163:23–28.
33. Förstner, W.; Moonen, B. *Geodesy: The Challenge of the 3rd Millennium*. Springer; New York: 2003. A metric for covariance matrices; p. 299–309.
34. Friedland S. Matrices with prescribed off-diagonal elements. *Israel J Math*. 1972; 11:184–189.
35. Frisch, R. *Statistical Confluence Analysis by Means of Complete Regression Systems*. Vol. 5. Universitets Økonomiske Institut; 1934.
36. Guidorzi R. Identification of the maximal number of linear relations from noisy data. *Systems Control Lett*. 1995; 24:159–165.
37. Guttman L. Some necessary conditions for common-factor analysis. *Psychometrika*. 1954; 19:149–161.
38. Harman, H. *Modern Factor Analysis*. University of Chicago Press; Chicago, IL: 1976.
39. Harman H, Jones W. Factor analysis by minimizing residuals (minres). *Psychometrika*. 1966; 31:351–368. [PubMed: 5221131]
40. Heij C, Scherrer W, Deistler M. System identification by dynamic factor models. *SIAM J Control Optim*. 1997; 35:1924–1951.
41. Horn, R.; Johnson, C. *Matrix Analysis*. Cambridge University Press; Cambridge, UK: 1990.
42. Jolliffe, I. *Principal Component Analysis*. Wiley Online Library; 2005.
43. Jöreskog K. Some contributions to maximum likelihood factor analysis. *Psychometrika*. 1967; 32:443–482.
44. Jöreskog K. A general approach to confirmatory maximum likelihood factor analysis. *Psychometrika*. 1969; 34:183–202.
45. Kalman, R. System identification from noisy data. In: Bednarek, A.; Cesari, L., editors. *Dynamical Systems II*. Academic Press; New York: 1982. p. 135–164.
46. Kalman R. Identification of noisy systems. *Russian Math Surveys*. 1985; 40:25–42.
47. Kassam S, Poor H. Robust signal processing for communication systems. *IEEE Commun Mag*. 1983; 21:20–28.
48. Kassam S, Poor H. Robust techniques for signal processing: A survey. *Proc IEEE*. 1985; 73:433–481.
49. Keshavan R, Montanari A, Oh S. Matrix completion from a few entries. *IEEE Trans Inform Theory*. 2010; 56:2980–2998.

50. Keshavan R, Montanari A, Oh S. Matrix completion from noisy entries. *J Mach Learn Res.* 2010; 99:2057–2078.
51. Kim SJ, Boyd S. A minimax theorem with applications to machine learning, signal processing, and finance. *SIAM J Optim.* 2008; 19:1344–1367.
52. Klepper S, Leamer E. Consistent sets of estimates for regressions with errors in all variables. *Econometrica.* 1984; 52:163–183.
53. Koopmans, T. *Linear Regression Analysis of Economic Time Series.* Harrlem-de Erwen, F.; Bohn, NV., editors. Netherlands Economic Institute; 1937.
54. Lam C, Yao Q. Factor modeling for high-dimensional time series: Inference for the number of factors. *Ann Statist.* 2012; 40:694–726.
55. Ledermann W. On the rank of the reduced correlational matrix in multiple-factor analysis. *Psychometrika.* 1937; 2:85–93.
56. Ledermann W. On a problem concerning matrices with variable diagonal elements. *Proc Roy Soc Edinburgh.* 1940; 60:1–17.
57. Lin F, Jovanovi M. Least-squares approximation of structured covariances. *IEEE Trans Automat Control.* 2009; 54:1643–1648.
58. Los C. Identification of a linear system from inexact data: A three-variable example. *Comput Math Appl.* 1989; 17:1285–1304.
59. Negahban S, Wainwright M. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Ann Statist.* 2011; 39:1069–1097.
60. Ning, L.; Jiang, X.; Georgiou, T. Geometric methods for estimation of structured covariances. *Proceedings of the American Control Conference; Montreal, Canada.* 2012.
61. Olkin I, Pukelsheim F. The distance between two random vectors with given dispersion matrices. *Linear Algebra Appl.* 1982; 48:257–263.
62. Pennec X, Fillard P, Ayache N. A Riemannian framework for tensor computing. *Internat J Comput Vision.* 2006; 66:41–66.
63. Recht B, Fazel M, Parrilo PA. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.* 2010; 52:471–501.
64. Reiersøl O. Confluence analysis by means of lag moments and other methods of confluence analysis. *Econometrica.* 1941; 9:1–24.
65. Rohde A, Tsybakov A. Estimation of high-dimensional low-rank matrices. *Ann Statist.* 2011; 39:887–930.
66. Saunderson J, Chandrasekaran V, Parrilo PA, Willsky AS. Diagonal and low-rank matrix decompositions, correlation matrices, and ellipsoid fitting. *SIAM J Matrix Anal Appl.* 2012; 33:1395–1416.
67. Scherrer W, Deistler M. A structure theory for linear dynamic errors-in-variables models. *SIAM J Control Optim.* 1998; 36:2148–2175.
68. Shapiro A. Rank-reducibility of a symmetric matrix and sampling theory of minimum trace factor analysis. *Psychometrika.* 1982; 47:187–199.
69. Shapiro A. Weighted minimum trace factor analysis. *Psychometrika.* 1982; 47:243–264.
70. Shapiro A. Identifiability of factor analysis: Some results and open problems. *Linear Algebra Appl.* 1985; 70:1–7.
71. Söderström T. Errors-in-variables methods in system identification. *Automatica.* 2007; 43:939–958.
72. Spearman C. General intelligence, objectively determined and measured. *Amer J Psychol.* 1904; 15:201–292.
73. Székely, G. *Statistics for the 21st Century: Methodologies for Applications of the Future.* CRC Press; 2000.
74. Tintner G. Static econometric models and their empirical verification, illustrated by a study of the American meat market. *Metroeconomica.* 1950; 2:172–181.
75. Van Trees, H. *Optimum Array Processing.* Wiley-Interscience; New York: 2002.
76. Varga, R. *Geršgorin and His Circles.* Springer-Verlag; Berlin: 2004.

77. Verdú S, Poor H. On minimax robustness: A general approach and applications. *IEEE Trans Inform Theory*. 1984; 30:328–340.
78. Woodgate K. An upper bound on the number of linear relations identified from noisy data by the Frisch scheme. *Systems Control Lett*. 1995; 24:153–158.
79. Woodgate, K. On Computing the Maximum Corank in the Frisch Scheme. Citeseer: 2007. preprintavailable online from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.55.2691&rep=rep1&type=pdf>

Table 9.1

$\widehat{\text{rank}}(\Sigma)$ as a function of λ and ε .

ε	0				0.08	0.10	0.12	0.14	0.16
	λ								
1	46				26	24	23	22	22
5	46				17	14	10	10	9
10	45				16	12	10	10	8
20	45				15	12	10	10	8
50	45				15	12	10	10	8
100	45				15	11	10	10	8