

**A framework for spatiotemporal prediction
with small and heterogeneous data
—and an application to consumer price indexes—**

Dissertation zur Erlangung des Grades eines
Doktors der Wirtschaftswissenschaften (Dr. rer. pol.)
der Fakultät für Wirtschaftswissenschaften
der Universität Bielefeld

vorgelegt von Christian Heinze

1. Gutachter: Prof. Dr. Harry Haupt

2. Gutachter: Prof. Dr. Dietmar Bauer

Datum der mündlichen Prüfung: 06.10.2016

Gedruckt auf alterungsbeständigem Papier gemäß ISO 9706

Contents

Preface	ii
1. Price level prediction	1
1.1. Price indexes	1
1.2. Price index prediction	7
Comments and references	13
2. Euclidean space basics	15
2.1. Fundamentals	15
2.2. Unitary maps	20
2.3. Projectors	24
2.4. Gramians	26
2.5. Singular values	31
Comments and references	35
<i>Appendix</i>	38
3. Regularized least-squares estimation	42
3.1. Basic convex analysis	42
3.2. Regularized least-squares	49
3.3. A gradient descent algorithm	54
3.4. A poor man's factor model	56
3.5. Transition matrix estimation	61
Comments and references	73
<i>Appendix</i>	78
4. Prediction techniques	89
4.1. Oblique approximation	89
4.2. Subordinate directions	94
4.3. A prediction framework	99
4.4. Prediction algorithms	105
Comments and references	119
<i>Appendix</i>	121
Global bibliography	125
Index	128

Preface

Contribution

This text considers the prediction of consumer price indexes which allow to compare the consumer price level across time and space. For specificity, the discussion is in terms of German counties and the years 1993–2014, but all methods apply more generally.

Section 1 explains that prediction of these spatiotemporal price indexes—after a suitable reformulation—amounts to the prediction of a “long” vector x based on a “short” vector of observations y . These observations equal weighted sums of the entries of x and may only be observed with error. More specifically, the setup takes the form

$$\begin{array}{c} \text{“short” vector} \\ \text{of observations} \end{array} \left\{ \begin{array}{c} y \\ \end{array} \right\} = \underbrace{\begin{array}{c} \boxed{A} \\ \text{given} \\ \text{nonrandom matrix} \end{array}}_{\text{given nonrandom matrix}} \left\{ \begin{array}{c} x \\ \end{array} \right\} + \left\{ \begin{array}{c} e \\ \end{array} \right\} \begin{array}{c} \text{“long” target vector} \\ \text{observation errors} \end{array} .$$

Therein, A symbolizes a given nonrandom matrix, and e denotes a vector of observation errors which may be (partially) equal to zero. Section 3 and 4 develop some statistical methodology to tackle this prediction task. Hence, this text naturally splits into two parts: the first part (section 1) outlines an application, and the second part (sections 2, 3, and section 4) presents some corresponding theory. The remainder of this section briefly summarizes each of the four main sections and highlights the individual contributions. In general, most sections may be read independently after acquainting oneself with the basic notation presented in section 2.1 as well as the first parts of section 2.3 and 2.4.1.

Section 1 presents the available—from official statistics—price indexes for spatial and temporal price level comparisons between German counties and states. The second part of this section develops a formal framework which integrates the different index types and defines an additional price index which allows spatiotemporal price comparison. Finally, the section sketches a data-based procedure for the prediction of the latter index.

Section 2 gathers the prerequisites for the subsequent investigation of the suggested prediction procedure. The presentation has its peculiarities, but the material is standard.

Section 3 initially focuses on an abstract regularized least-squares criterion defined on the space of symmetric matrices. The restriction to symmetric matrices requires an adapted duality analysis alongside an associated investigation of the set of minimizers. The presented (proximal) gradient algorithm is standard except for the stopping rule, which derives from the previous study of duality. The subsequent discussion introduces a stripped-down factor model and uses the minimizer of a special case of the abstract least-squares criterion for estimation. The latter allows the integration of spatial information. Herein, the individual building blocks stem from literature; their integration and application does not. This section finishes with a novel investigation of statistical properties of the proposed estimator. The specific form of the results in this final part, in

particular, those in section 3.5.2 and the final part of section 3.5.3, should be considered preliminary. A fully satisfactory analysis remains a topic for future work.

The final section 4 develops the prediction framework alluded to in section 1 in-depth. The investigation contains a population-level quality assurance and closes with a discussion of suitable computational techniques. Once more, the integration of the individual components into a coherent whole is novel; the ingredients are not.

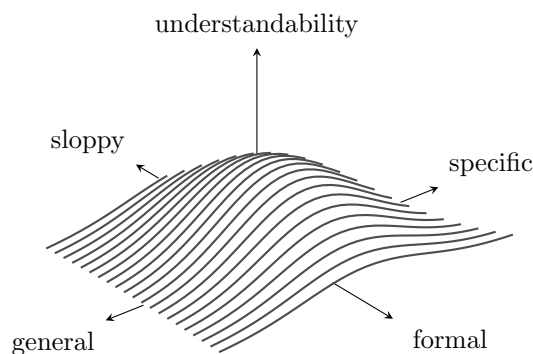
Organization

Several major and minor sections structure the content of this text. Within major sections, propositions, lemmas, and corollaries share their counter. The same applies to figures and tables. Frequently, definitions are recalled upon use to aid skimming the text. Moreover, an index at the end of the document lists most keywords. Markers in the right margin point to the location of the indexed definitions.

References gather in a literature discussion at the close of the corresponding major section. The listed sources often supply a more detailed presentation including additional definitions and results omitted from this text. However, the bibliography by no means constitutes a comprehensive list nor a compilation of the original sources.

This text adopts a tutorial style. In particular, it includes many of the required definitions and proves most of its assertions—the sole notable exception being the Hanson-Wright inequality in appendix 3.a. However, this text is far from being a self contained introduction to the individual topics, which are only developed to the extend needed in the subsequent derivations. Many important results are accompanied by a discussion indicating how and why the result works instead of a short formal proof. Mere technicalities come with a proof in traditional form (with the above mention exception). Included proofs are relegated to an appendix of the corresponding major section.

In general, the presentation is meant to be *sloppy enough* such that the key ideas are not hidden by formalities, *formal enough* to make sense, *general enough* to reveal what really matters, and *specific enough* to not feel like abstract nonsense. The following sketch visualizes this trade-off perceived by the author during the creation of this text.



As an example, intuitive notions are left undefined, which is indicated by enclosing quotation marks “...”, whenever an exact meaning is not needed.

Most notation is standard or otherwise explained. A notable exception are sets: these are symbolized by labels surrounded by braces—as commonplace in probability theory—if the meaning is uncontroversial. For example, $\{\|\bullet\| = 1\}$ contains all elements of length one; $\{f = a\}$ gathers all preimages of a under the function f ; and so forth. In addition, the symbol \mathbb{N} denotes the set of positive integers $\{1, 2, 3, \dots\}$. Likewise integer-valued indexes start at one unless indicated otherwise. Accordingly, the summation shorthand $\sum_{t \leq k}$ abbreviates the more verbose $\sum_{t=1}^k$. Sometimes integer-valued quantities are not explicitly labeled as such; then their nature follows from the context.

Acknowledgments

On the professional side, I'd like to thank my two advisers Prof. Dr. Harry Haupt and Prof. Dr. Dietmar Bauer for their patient support, for generously giving advice, and creating a pleasant working environment. Apl. Prof. Dr. Peter Wolf chaired my defense and I'm grateful he did. During my time at Bielefeld University, I have been fortunate to enjoy the company of many wonderful colleagues. In particular, I want to thank my former office mate Oliver Jones for many benefiting and often entertaining discussions. Last but not least, I want to thank Helga Radtke for advice on and help with the many organizational aspects surrounding this project.

On the personal side, I am greatly in debt to my parents Eckhardt and Gitta and my girlfriend Teresa for their invaluable support and encouragement.

1. Price level prediction

1.1. Price indexes

1.1.1. Overview

German price statisticians use Laspeyeres-type price indexes to implement price comparison in time at state and national level as well as price comparison in space at the city level. Price index calculations are complex, but are roughly summarized by

$$\text{cpi}_{t/b_t, C'_k} = \sum_j \frac{\sum_{i \in C'_k} \widehat{\text{price}}_{t,i,j}}{\sum_{i \in C'_k} \widehat{\text{price}}_{b_t,i,j}} \widehat{\text{wgt}}_{b_t,j}, \quad \text{cpi}_{t/b_t, C'} = \sum_{k \in S} \widehat{\text{swgt}}_{b_t,k} \text{cpi}_{t/b_t, C'_k}, \quad \langle 1.1a \rangle$$

$$\text{and } \text{smpi}_{t,i/c} = 100 \times \sum_j \frac{\widehat{\text{price}}_{t,i,j}}{\widehat{\text{price}}_{t,c,j}} \widehat{\text{wgt}}_{t,j}. \quad \langle 1.1b \rangle$$

Indexes calculated according to $\langle 1.1a \rangle$ serve as inflation measures and are referred to as consumer price indexes (**cpi**). The third index type $\langle 1.1b \rangle$ implements price comparison in space and is called a spatial consumer price index (**smpi**).

The real number $\widehat{\text{price}}_{t,i,j}$ in $\langle 1.1a \rangle$ and $\langle 1.1b \rangle$ refers to the price observed in location i during year t of a good/service—a price representative—representing the j -th goods/service category in a basket of goods considered for index calculation. The basket of goods is meant to reflect average consumption choices. Here, rent is not part of this selection. Prior to index calculation, prices are adjusted to take care of temporary unavailability and special offers—these are features of the price representative and not the category it is meant to represent. In some cases, more than one price representative is observed, and then $\widehat{\text{price}}_{t,i,j}$ equals an average of the observed prices.

The 402 German counties subdividing the 16 German states—as of the December 31st, 2013—form the spatial entities of interest. Subsequently, these counties are represented by the elements of $C = \{1, \dots, 402\}$; these numbers also serve to index other objects associated with these counties. Similarly, states are represented by the labels in $S = \{\text{BB}, \text{BE}, \dots, \text{TH}\}$. The caption of table 1.1 lists all 16 state labels and the corresponding full names. The subsets $C_{\text{BB}}, \dots, C_{\text{TH}}$ partition C such that $i \in C_k, i \in C, k \in S$, indicates that county i lies in state k ; thus, $C_k \cap C_{k'} = \emptyset$ if $k \neq k'$ and $\cup_{k \in S} C_k = C$. The relevant time points are given by the years 1993–2014, which also serve as time indexes.

The lefthand side of $\langle 1.1a \rangle$ describes **cpi** calculation at the state level. Therein, the average price $\frac{1}{m'_k} \sum_{i \in C'_k} \widehat{\text{price}}_{t,i,j}$ across a subset $C'_k \subset C_k$ of counties of state $k \in S$ observed for good/service category j in year t is compared to the respective average price in a basis year b , wherein m'_k symbolizes the number of elements of C'_k . The **cpi** results as a weighted average of the category specific ratios across all categories in the basket of goods. The weight $\widehat{\text{wgt}}_{b_t,j} \in (0, 1)$ attached to the j -th category is an estimate of the average—at the national level—expenditure share of category j in the basis year b_t . Section 1.1.2 contains further details and also explains the association $t \mapsto b_t$.

A weighted average of the state **cpi**s shown on the righthand side of $\langle 1.1a \rangle$ addresses

price comparison in time at the national level. The share of the German population living in state k during the basis year b_t serves as the weight $\widehat{\text{swgt}}_{b_t,k}$ for the state index $\text{cpi}_{t/b_t,C'_k}$ in the national consumer price indexes $\text{cpi}_{t/b_t,C'}$ with $C' = \cup_{k \in S} C'_k$ calculated up to (but not including) the year 2000. Later cpis use the respective state's share of the German private consumption expenditures in the respective basis year b_t .

Cross-sectional price comparison implemented via the scpi formula compares the price of the j -th representative to the respective price observed in a reference location c . Estimated average expenditure shares relate to the same year as the price ratios. The basket of goods is identical to the one used for cpi calculations in the respective year; in particular, rent is excluded here. Section 1.1.3 provides further information on the scpis .

1.1.2. Consumer price indexes

Surveys for cpi calculation are implemented on a monthly basis to inform about consumer price inflation rates at the state and the national level. The corresponding index calculation at state level proceeds as shown on the lefthand side of <1.1a> but with t and b_t replaced by the respective month and a fixed month of the basis year b_t . The subsequent discussion considers averages of these monthly indexes across the respective year t . Hence, the numerator of the lefthand side of <1.1a> amounts to an unweighted average of the prices observed for the j -th good/service category during the twelve months of the year t ; the denominator refers to the given basis month.

The state cpi compares prices collected in one state to prices collected in the same state during a basis period. Therefore, the cpi becomes meaningful if the price representatives are fixed in the basis year and remain unaltered subsequently. In practice, the representatives are chosen by local staff and remain unaltered up to the next basis year. This decentralized selection procedure ensures that price representatives reflect local taste. Local price collection is organized by each state's own statistical bureau and is limited to a selection of counties $C'_k \subset C_k$, $k \in S$. The price collection within $i \in C'_k$ takes place in one or more of its cities. Subsequently, the inflation measured in the selected cities is taken as representative for the respective county. Goods and services having a common nation-wide price are observed by the national statistical bureau. The latter supplies these data to the state bureaus and obtains the state indexes in return.

The data collection and processing methodology, in particular, the weighting scheme and the selection of price representatives, is revised roughly every five years—about three year after the new basis year—alongside a basis year change. The revision process may involve changes in the selection of cities used for price observation, that is, the sets C'_k , $k \in S$, may change with t , but this possibility is ignored here due to data availability. Upon revision, the whole series of published indexes is adjusted to hide the resulting structural break. The cpis for all years following and including the new basis year are recalculated using the new methodology and weights. Indexes published prior to the new basis year are adjusted by a heuristic manipulation. After the revision only the adjusted indexes are available. Thus, cpis for year t are calculated using the methodology, price representatives, and weights corresponding to the basis year $b_t = \max\{l \leq t \mid l \text{ is a basis year}\}$, but may later be adjusted to a different basis year. Here, all cpis refer to the basis




Figure 1.2

The figure shows the observed growth rates of the `cpi`s and `cpiR`s at state and national level, wherein `cpiR` represents indexes of the type `<1.1a>` but based on a basket of goods including rent. The `cpi` inflation rates result as $\text{cpi}_{t/b,C'_k} / \text{cpi}_{t-1/b,C'_k} - 1$ at state level and $\text{cpi}_{t/b,C'} / \text{cpi}_{t-1/b,C'} - 1$, $b = 2010$, at national level; the calculation for `cpiR` is in analogy. Solid squares \blacksquare symbolize `cpi` growth rates at state level; growth rates of `cpiR`s at state level correspond to \square . If either of the index values needed for calculation is missing, then the respective symbol is absent. Inflation rates at the national level are visualized as a solid (`cpi`) and dashed (`cpiR`) line, respectively. Vertical solid lines illustrate the spread of the growth rates at state level in the respective year.

year 2010. In addition, the present analysis acts—for most purposes—as if $b_t = 2010$ for all $t \leq 2014$, that is, the adjusted indexes receive no special treatment.

Table 1.1 summarizes the availability of `cpi`s. Indexes based on a basket of goods without rent are not available for the years 1993 and 1994. For those years, the available indexes based on a basket of goods including rent—denoted by `cpiR`—may be used as a surrogate. Data availability improves in 1995 with full availability for years follow-

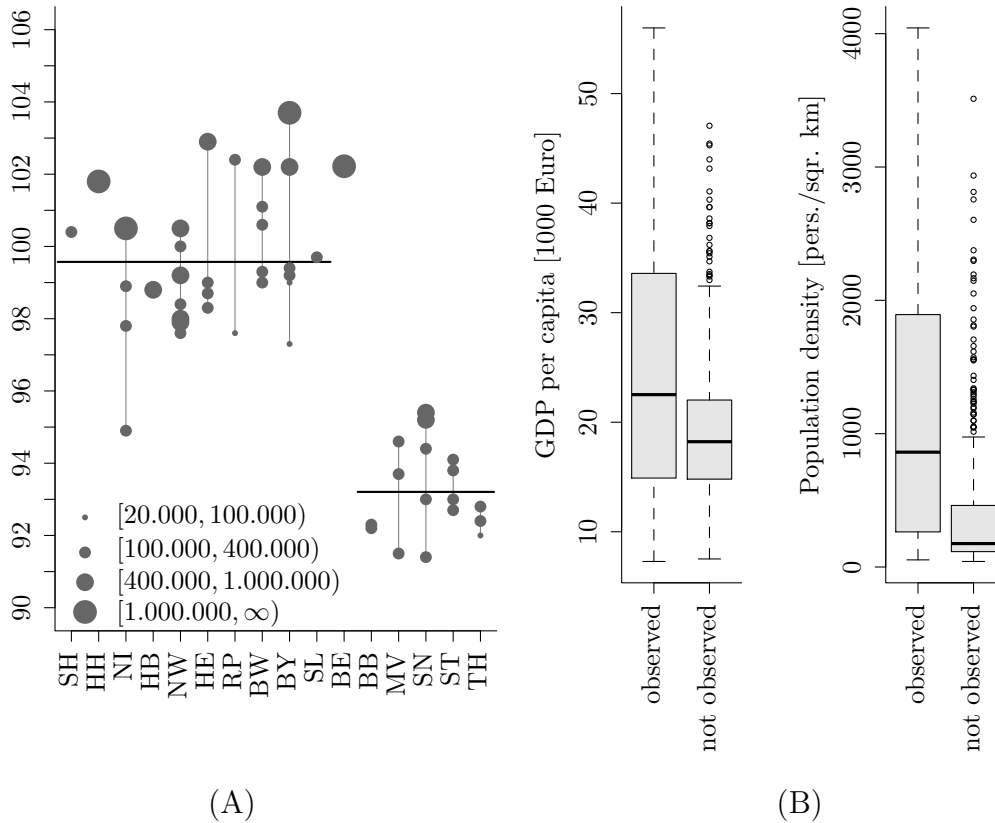


Figure 1.3

The figure shows the *scpis* corresponding to 50 counties and $t = 1993$ (Panel (A)) as well as the spatial variation of the gross domestic product (GDP) per capita and the population density separately for sampling and non-sampling locations (Panel (B)). The *scpis* in panel (A) are visualized by \bullet ; the size of these points reflects the respective of the four size categories. Horizontal solid lines indicate the average *scpi* across the surveyed Western German counties and the surveyed East German counties—both excluding Berlin. Panel (B) shows boxplots.

ing 2000 except for Schleswig-Holstein, Hamburg, and Bremen. Similarly, national *cpis* including rent are available for all years and all years except 1993 and 1994 if rent is excluded. These indexes amount to a weighted sum of the state *cpis* corresponding to the same year but convey additional information as the state *cpis* are never fully available.

Figure 1.2 shows the growth rates of *cpis* and *cpis*R at state and national level. It shows a considerable spatial variation of the growth rates in the individual years—indicated by vertical solid lines—of one to two percentage points. However, the figure also creates the impression of a joint movement in time of the state level growth rates.

1.1.3. Spatial consumer price indexes

Products of the same category observed for *cpis* calculation usually differ considerably across sampling locations in product quality, package size, and so forth due to the




Figure 1.4

The figure shows the `scpis` calculated from the surveys in 1993 and 2007. Equally sized \bullet indicate index values for locations included in the former survey. Vertical solid lines highlight the spread of the indexes calculated from the latter survey across the respective state. Horizontal solid lines visualize the average index value separately for East German states and West German states—both excluding Berlin—as well as the two surveys.

decentralized choice of the price representatives. As mentioned in section 1.1.2, this is not only acceptable but also desirable for `cpi` calculation. However, `scpi` calculation based on `<1.1b>` and such data potentially leads to a comparison—across space—of prices of rather different products. Meaningful `scpi` calculation therefore requires a survey of its own, which ensures comparable choices of price representatives across sampling locations.

Such surveys resembles those implemented for `cpi` calculation to the extend that local price collectors visit suppliers and observe prices of selected representatives. However, they do so equipped with a more detailed descriptions of providers and goods/services used for price representation. Such surveys are implemented infrequently and only for a small number of locations—mostly for cost reasons. More specifically, the present investigation uses only the latest data set of this type. This data set relates to $t = 1993$ and 50 Germany cities, which are subsequently identified with their surrounding counties.

The respective indexes are based on a subset of the basket of goods and slightly adjusted good/service category weights $\widehat{\text{wgt}}_{t,j}$ of the `cpi`s in 1993. The former German capital Bonn—a county of its own—serves as the reference location c . Surveyed cities are

split into four groups based on their population size: 20 000 to 100 000, 100 000 to 400 000, 400 000 to 1 000 000, and above one million. If the price for a good/service category cannot be observed in a county—due to the lack of shopping facilities, then the average price for that good/service category across locations of the same size group fills this gap.

Panel (A) of figure 1.3 shows the corresponding `scpis`. The indexes exhibit a notable east/west divide, which comes as no real surprise as the two parts re-unified only a few years earlier. Panel (B) of that figure presents the spatial variation of the gross domestic product per capita and the population density separately for surveyed and non-surveyed counties. The shown boxplots indicate that the survey locations are not representative for all counties with respect to economic potency and settlement structure.

Finally, spatial price indexes from a non-official survey are available in addition to the `scpis` provided by official statistics. The index calculation is according to <1.1b> and based on a subset of the `cpi` basket of goods in 2007 as well as the `cpi` category weights. Bonn serves as the reference location c . However, the survey was implemented by a single person and during the years 2005–2009 with most of the data being from 2006–2008. Herein, these `scpis` are taken to refer to $t = 2007$.

Figure 1.4 compares the `scpis` obtained from the two surveys. Hence, the focus is on the locations included in the survey of 1993, whose indexes are symbolized by \bullet . The presentation does not allow to distinguish between locations, however, shows a larger spread of the index values for $t = 2007$ —even within the limited set of counties surveyed in 1993. Vertical solid lines represent the overall spread of these indexes. Herein, the substantially larger spread for $t = 2007$ reflects the inclusion of rural areas not surveyed in 1993. Index averages are included in the form of horizontal solid lines. These lines show that Bonn’s `scpi` enjoyed an increase relative to the other West German as well as East German counties. Moreover, these averages indicate a moderate price level convergence of the East German states to the West German states.

1.2. Price index prediction

1.2.1. A framework for price index prediction

This section develops a framework which integrates price indexes of the form <1.1a> alongside indexes of the form <1.1b>. In section 1.2.2, this setup guides the development of a prediction strategy for price indexes implementing spatiotemporal price level comparison. The framework involves several numerical characteristics of the $m = 402$ German counties mentioned in section 1.1.1 during the years $t \in \{1993, \dots, 2014\}$. Random variables—defined on a common measurable space (Ω, \mathcal{F}) —represent the values of these characteristics for all counties $i \leq 402$ and years t . The elements ω of Ω embody a priori imaginable “states of the world”. The numbers shown in section 1.1 correspond to one of these “states”, that is, equal the images of an element $\omega \in \Omega$ under the corresponding random variables. Herein, characteristics expressing “spatial properties”, the membership in C' , as well as the state weights $\widehat{\text{swgt}}_{b_t, k}$ in <1.1a> and the data availability are taken to be constant across $\omega \in \Omega$. Economic preknowledge is represented by a set \mathcal{P} of probability measures on (Ω, \mathcal{F}) —a statistical model, which contains the probability

measure \mathbb{P} on (Ω, \mathcal{F}) describing the “data generation”, that is, the choice of $\omega \in \Omega$ underlying the observations. Elements of \mathcal{P} satisfy the requirements outlined below and, in addition, ensure square-integrability of all mentioned random variables.

The price level of county i during year t provides the corner stone of the subsequent developments. This price level—a positive quantity—incorporates the vague idea of the “cost of living” in county i during year t . The associated random variables form the bridge that connects the various types of price indexes and also provide the ingredients for a spatiotemporal price index. This purely formal role requires no further conceptual considerations—such as a discussion of the meaning of “cost of living”.

The rationale behind the spatial consumer price index in <1.1b> is to compare the price level—in form of the price of a basket of goods—in county i during year t with that of a reference location c during the same year. The two equivalent equalities

$$\frac{\text{scpi}_{t,i/c}}{100} = \frac{\text{price level at } (t, i)}{\text{price level at } (t, c)} \quad \text{and} \quad \log \left[\frac{\text{scpi}_{t,i/c}}{100} \right] = x_{t,i} - x_{t,c} \quad <1.2>$$

formalize this thinking, wherein $x_{t,i}$ symbolizes the (natural) logarithm of the price level in county i during year t . The variables $x_{t,i}$, $(t, i) \in \{1993, \dots, 2014\} \times \{1, \dots, 402\} = I_x$, also give rise to the spatiotemporal index $\text{scpi}_{(t,i)/(d,c)}^\circ = 100 \times e^{x_{t,i} - x_{d,c}}$. Indexes of this type compare (the price level of) county i during year t with a fixed reference time/space point (d, c) . Therein, the symbol e represents Euler’s number. Choosing c equal to Bonn and $d = 1993$ allows a direct comparison with the available **scpis** for that year.

The lefthand side of <1.2> expresses a transformation of the **scpis** available in $t = 1993$ and $t = 2007$ as linear combinations of $x_{t,i}$, $(t, i) \in I_x$. A similar representation as in <1.2> of the growth rates of the state **cpis** in <1.1a> is possible if these growth rates are taken as measures of the average growth rate of price levels $e^{x_{t,i}}$ across the respective state. These quantities usually differ from the growth rates of the average price level $\sum_{i \in C'_k} e^{x_{t,i}} / m'_k$. The interpretation as an average growth rate implies

$$\left[\frac{\text{cpi}_{t/b, C'_k}}{\text{cpi}_{t-1/b, C'_k}} - 1 \right] = \frac{1}{m'_k} \sum_{i \in C'_k} \underbrace{\left[\frac{\text{price level at } (t, i)}{\text{price level at } (t-1, i)} - 1 \right]}_{\approx x_{t,i} - x_{t-1,i}} \quad <1.3a>$$

$$\approx \frac{1}{m'_k} \sum_{i \in C'_k} x_{t,i} - \frac{1}{m'_k} \sum_{i \in C'_k} x_{t-1,i} \quad <1.3b>$$

for states $k \in S = \{\text{BB}, \dots, \text{TH}\}$ and with C'_k , m'_k , and b being the set of surveyed counties, the number of its elements, and the common base year $b = 2010$, respectively. Herein, $k! = 1 \times 2 \times \dots \times k$ and the second step relies on $a/b - 1 = e^{\log(a/b)} - 1 = \sum_{k \geq 1} (\log^k(a/b)/k!) \approx \log a - \log b$ with $a, b > 0$ such that $|a/b - 1|$ is “small”. The previous display relates to the underlying random variables. Figure 1.2 shows that the available inflation rates—at state level as on the lefthand side of <1.3a>—are sufficiently small. However, the possibility of cancellation when summing positive and negative inflation rates for counties contained in the same state prevents a direct conclusion for the county level inflation rates. Moreover, these numbers correspond to a single




Figure 1.5

The figure shows the absolute values of the changes of state weights $|\widehat{\text{swgt}}_{b_t,k} - \widehat{\text{swgt}}_{b_{t-1},k}|$ (Panel (A)) and the deviation of ratios $\text{cpi}_{t/b,C'_k} / \text{cpi}_{t/b,C'}$ of the state cpi s and the national cpi from one (Panel (B)). Both quantities are represented by “•”. If one of the ingredients needed for the calculation is not available, then the respective symbol is absent.

image under the random variables. Thus, the available evidence supports the claim of approximate equality only to the extent that it does not contradict that equality.

Finally, the growth rates of the cpi at the national level approximately equal

$$\frac{\text{cpi}_{t/b,C'}}{\text{cpi}_{t-1/b,C'}} - 1 = \frac{\sum_{k \in S} \widehat{\text{swgt}}_{b_t,k} \text{cpi}_{t/b,C'_k}}{\text{cpi}_{t-1/b,C'}} - 1 \quad \langle 1.4a \rangle$$

$$\approx \sum_{k \in S} \widehat{\text{swgt}}_{b_{t-1},k} \frac{\text{cpi}_{t-1/b,C'_k}}{\text{cpi}_{t-1/b,C'}} \left[\frac{\text{cpi}_{t/b,C'_k}}{\text{cpi}_{t-1/b,C'_k}} - 1 \right] \quad \langle 1.4b \rangle$$

$$\approx \sum_{k \in S} \widehat{\text{swgt}}_{b_t,k} \left[\frac{\text{cpi}_{t/b,C'_k}}{\text{cpi}_{t-1/b,C'_k}} - 1 \right] \approx \sum_{k \in S} \sum_{i \in C'_k} \frac{\widehat{\text{swgt}}_{b_t,k}}{m'_k} [x_{t,i} - x_{t-1,i}], \quad \langle 1.4c \rangle$$

wherein $b_t = \max\{l \leq t \mid l \text{ is a basis year}\}$ provides the original basis year of the national indexes $\text{cpi}_{t/b, C'}$ and determines the state weights $\widehat{\text{swgt}}_{b_t, k}$, $k \in S$, used on the righthand side of <1.1a>. The first approximative equality—indicated by \approx —in <1.4b> relies on $\widehat{\text{swgt}}_{b_t, k} \approx \widehat{\text{swgt}}_{b_{t-1}, k}$, which holds with equality unless t amounts to one of the original basis years 1995, 2000, 2005, and 2010. Panel (A) of figure 1.5 shows the absolute values of the differences of these two weights for all states $k \in S$ and years $t \geq 1994$. The shown numbers justify the latter approximative equality as these quantities are taken as constant across $\omega \in \Omega$. The conclusion in <1.4c> reuses the latter and $\text{cpi}_{t/b, C'_k} \approx \text{cpi}_{t/b, C'}$, $k \in S$. Therein, equality holds—by definition—if t equals the common basis year 2010. Panel (B) of figure 1.5 shows the translated ratios $\text{cpi}_{t/b, C'_k} / \text{cpi}_{t/b, C'} - 1$ for all $k \in S$ —subject to the data availability shown in table 1.1—and years $t \geq 1995$. These quantities increase in absolute value at $|t - b|$ increases, but are still small even for the case $t = 1995$; thus, these observations do not contradict <1.4c>.

Subsequently, the available scpis —in the form shown on the righthand side of <1.2>—and cpi inflation rates—at state level as well as national level and calculated as shown on the lefthand side of <1.3a> and <1.4a>, respectively—are denoted by $y_{t,i}$, $i \leq k_t$. Therein, $y_{t,1}, \dots, y_{t, k_t''}$, $k_t'' \in \mathbb{N} \cup \{0\}$, represent the observed inflation rates; $y_{t, k_t''+1}, \dots, y_{t, k_t''+k_t'}$, $k_t' \in \mathbb{N} \cup \{0\}$, symbolize the observed and transformed scpis . Thus, $k_t = k_t' + k_t''$ equals the number of observations in t . The cases $k_t'' = 0$ and $k_t' = 0$ are possible and indicate the absence of observed inflation rates—as in $t = 1993$ —and the absence of observed scpis —as in $t = 1994$, respectively. However, the case $k_t = 0$ does not occur if cpiR inflation rates replace the non-available cpi inflation rates for $t \leq 1994$.

In total, the observed scpis and cpi inflation rates exhibit the representation

$$y_t = \begin{pmatrix} y_{t,1} \\ \vdots \\ y_{t, k_t''} \\ y_{t, k_t''+1} \\ \vdots \\ y_{t, k_t''+k_t'} \end{pmatrix} = \underbrace{\begin{bmatrix} -I & B_t \\ & J_t \end{bmatrix}}_{A_t} \underbrace{\begin{pmatrix} B_t x_{t-1} \\ x_t \end{pmatrix}}_{z_t} + S_t \begin{pmatrix} \bar{v}_{t,1} \\ \vdots \\ \bar{v}_{t, k_t''} \\ \bar{v}_{t, k_t''+1} \\ \vdots \\ \bar{v}_{t, k_t''+k_t'} \end{pmatrix} \quad \text{with } x_t = \begin{pmatrix} x_{t,1} \\ \vdots \\ x_{t,m} \end{pmatrix}. \quad <1.5>$$

Herein, I denotes the $k_t'' \times k_t''$ identity matrix, and the j -th row $b_j^{(t)} = (b_{j,1}^{(t)}, \dots, b_{j,m}^{(t)})$ of $B_t \in \mathbb{R}^{k_t'' \times m}$ corresponds either to the observed inflation rate of a state $k \in S$ or the observed national inflation rate. In the former case, its i -th entry $b_{j,i}^{(t)}$ equals $1/m'_k$ if $i \in C'_k$ and zero otherwise. In the latter case, the entries are in accordance with <1.4c>. In particular, $C'_k \cap C'_{k'} = \emptyset$ whenever $k \neq k'$ implies that the rows of B_t corresponding to state inflation rates are pairwise orthogonal. Table 1.1 shows that the complete set of state inflation rates is never observed. The latter together with $\widehat{\text{swgt}}_{b_t, k} > 0$ ensures that the rows of B_t are linearly independent for all t . Furthermore, the matrix $J_t \in \mathbb{R}^{k_t' \times m}$ is of the form $J_t^{(1)} - J_t^{(2)}$, wherein the rows of $J_t^{(1)}$ amount to distinct standard basis elements e_i of \mathbb{R}^m —defined in example (a) in section 2.1.1—with $i \neq c$ —the reference location (Bonn), and all entries of $J_t^{(2)}$ are zero except for those in the c -th column which equal -1 . In particular, the rows of J_t are linearly independent. In case $k_t' = 0$,

the second block row of the aggregation matrix A_t disappears; if $k_t'' = 0$, then $A_t = J_t$. In particular, the quantities $x_{1992,j}$ do not occur in <1.5>. Finally, the second summand $S_t \bar{v}_t$, $\bar{v}_t = (\bar{v}_{t,1}, \dots, \bar{v}_{t,k_t})$, on the righthand side of <1.5> embodies errors associated with the observation process. Therein, the random variables $\bar{v}_{t,i}$, $(t, i) \in I_{\text{obs}} = \bigcup_{t=1993}^{2014} (\{t\} \times \{1, \dots, k_t\})$, exhibit zero mean, unit variance, (pairwise) zero covariances, and zero covariances with $x_{t,j}$, $(t, j) \in I_x$. The matrices $S_t \in \mathbb{R}^{k_t \times k_t}$, $1993 \leq t \leq 2014$, determine the contemporaneous (co)variances of the observation errors in year t given by the entries of $S_t \bar{v}_t$. Hence, nonzero rows of S_t should be present for $t \in \{1994, 1995\}$ to capture the deviations between the unobserved `cpis` and their surrogates `cpIR` as well as in $t = 2007$ to represent doubts concerning the survey implementation.

In summary, the equation <1.5> expresses transformations of the two types of price indexes <1.1a> and <1.1b> in form of linear combinations of the underlying quantities $x_{t,j}$, $(t, j) \in I_x$. These quantities also provide the basic building blocks of the spatiotemporal price indexes $\widehat{\text{scpi}}_{(t,i)/(d,c)}^\circ$, $(t, i) \in I_x$, with reference time/space point $(d, c) \in I_x$. Moreover, the equation <1.5> parallels the specification <4.21b>; example (d) in section 2.1.1 and example (e) in section 2.4.1 bridge the differences in the notation.

1.2.2. A strategy for price index prediction

This section sketches a strategy for the prediction of the underlying quantities $x_{t,j}$, $(t, j) \in I_x$, and thus of the spatiotemporal indexes $\widehat{\text{scpi}}_{(t,i)/(d,c)}^\circ$, $(t, i) \in I_x$, as defined in section 1.2.1 for a given reference time/space point $(d, c) \in I_x$. To this end, the present section implicitly develops a representation of $x_{t,j}$, $(t, j) \in I_x$, of the form <4.21a>. This representation can be transformed into a corresponding representation of the entries $z_{t,j}$ of z_t —defined in <1.5>—as shown in <4.22>. Then, computations of the form <4.24>, <4.26>, and <4.27> lead to predictions $\hat{z}_{t,j}$ of $z_{t,j}$, that is, best guesses for the unobserved $z_{t,j}(\omega)$. Herein, $\omega \in \Omega$ denotes the argument corresponding to the available observations $y_{t,i}(\omega)$. Finally, predictions of $x_{t,j}$ follow from $x_{t,j} = z_{t,k_t''+j}$, $j \leq m$.

Section 4.3.1 explains that the equality $y_{t,i}(\omega) = \bar{a}_{t,i}^\top \hat{z}_t$ holds whenever $\bar{a}_{t,i} \in \mathbb{R}^{k_t''+m}$ denotes a row of A_t such that the corresponding row of S_t equals the zero vector. Therein, the form of $\hat{z}_t \in \mathbb{R}^{k_t''+m}$ parallels that of z_t —as defined in <1.5>—but with the random variables $z_{t,j}$, $j \leq k_t'' + m$, replaced by the predictions $\hat{z}_{t,j} \in \mathbb{R}$. In particular, the observed `scpis` satisfy $\log(\widehat{\text{scpi}}_{t,i/c}/100) = \hat{x}_{t,i} - \hat{x}_{t,c}$. Consequently, the plug-in predictions $\widehat{\text{scpi}}_{(t,i)/(d,c)}^\circ = 100 \times e^{\hat{x}_{t,i} - \hat{x}_{d,c}}$ interpolate the `scpis` observed in $t = 1993$ if c equals Bonn, $d = 1993$, and these `scpis` are (treated as) observed without error. An analogous results applies to `cpi` inflation rates, but only if state and national inflation rates implied by the predictions $\widehat{\text{scpi}}_{(t,i)/(d,c)}^\circ$, $(t, i) \in I_x$, of the spatiotemporal indexes are calculated according to the respective of the approximate equalities <1.3b> and <1.4c>.

Additional numerical characteristics $u_{i,t,j}$, $i \leq s$, of county i in year t enter the construction of the representation of $x_{t,j}$, $(t, j) \in I_x$, of the form <4.21a>. Examples of such variables included the gross domestic product per capita and the population density mentioned in panel (B) of figure 1.3. Section 4.3.1 contains a comprehensive account of the construction process. Therein, the characteristics are denoted by $z_{i,t,j}$ instead of $u_{i,t,j}$.

Here, the different symbol is justified to prevent confusion with $z_{t,j}$. The remainder of this section indicates how section 4.3.1 applies to the present prediction task and also comments on the estimation steps needed to obtain a complete prediction algorithm.

The additional variables are split into two groups: variables with first index $i < s'$ and variables with first index $i \geq s'$, wherein $1 \leq s' \leq s$, and the case $s' = 1$ indicates the absence of the first group. These variables are taken to satisfy the linear independence requirements specified in section 4.3.1. If $s = 2$ with $u_{1,t,j}$ and $u_{2,t,j}$ denoting the population density and gross domestic product per capita, then economic intuition suggests that the linear independence requirement—no linear combination of $\alpha_{t,j} + u_{2,t,j}\beta_2$, $(t, j) \in I_x$, is constant across $\omega \in \Omega$ —is satisfied whenever $\beta_2 \neq 0$. However, the presence of temporally or spatially lagged characteristics requires additional attention. This is further explained below <1.6>. If such lags are included, then lemma 4.4 may prove useful.

The price level variables $x_{t,j}$ and the additional characteristics $u_{i,t,j}$ are connected via

$$x_{t,j} = \alpha_{t,j} + \sum_{i \leq s} u_{i,t,j} \beta_i + v_{t,j}, \quad (t, j) \in I_x, \quad <1.6>$$

wherein each of the final summands $v_{t,j}$, $(t, j) \in I_x$, exhibits zero mean and zero covariance with every $u_{i,t,j}$, $i \leq s$, $(t, j) \in I_x$. Hence, this representation implies that the inequality $\mathbb{E}(x_{t,j} - \alpha_{t,j} - \sum_{i \leq s} u_{i,t,j} \beta_i)^2 \leq \mathbb{E}(x_{t,j} - a - \sum_{i,t,j} u_{i,t,j} b_{i,t,j})^2$ holds for all real numbers a , $b_{i,t,j}$, $i \leq s$, $(t, j) \in I_x$, and with \mathbb{E} denoting the \mathbb{P} -expectation. Sections 2.1.3, 2.2, and 2.3 in connection with example (e) of section 2.1 justify these inequalities and reveal that the restriction hiding in <1.6> amounts to the invariance of the projection coefficients β_1, \dots, β_s across $(t, j) \in I_x$. In fact, characteristics of other counties $j' \neq j$ and years $t' \neq t$ —such as the above mentioned spatial or temporal lags, can be added in the form $u_{i,t,j} = u_{i',t',j'}$ and $u_{i,t,j} = u_{i',t',j}$, respectively, with $i' \neq i$. In this case, the mapping $(i, t, j) \mapsto u_{i,t,j}$ is not injective, and restrictions on $\beta_b = (\beta_{s'}, \dots, \beta_s)$ beyond $\beta_b \neq 0$ may be needed to ensure the required linear independence. The invariance of β_1, \dots, β_s across time and space allows for the estimation of these coefficients, that is, the partial identification of the “data generating” $\mathbb{P} \in \mathcal{P}$. More specifically, combining the specification <1.6> with the observation equation <1.5> yields a system of equations which may be used for (generalized) least-squares estimation of β_1, \dots, β_s . If the elements of \mathcal{P} grant sufficient stochastic homogeneity and/or independence to the random vectors $(x_{t,j}, u_{1,t,j}, \dots, u_{s,t,j})$ —such that some law of large numbers becomes relevant, then least-squares estimation identifies the projection coefficients β_1, \dots, β_s . That is, these coefficients—elements of \mathbb{R} —are the limits in probability—or equivalently in distribution—of their estimators in a suitable asymptotic setting. The same applies to the (estimation of the) coefficients $\alpha_{t,j}$ if some specific structure of the correspondence $(t, j) \mapsto \alpha_{t,j}$ with “sufficiently low complexity” is added to the specification <1.6>.

If estimation uncertainty, that is, the use of estimates $\hat{\alpha}_{t,j}$ and $\hat{\beta}_1, \dots, \hat{\beta}_s$ in place of the unknown coefficients $\alpha_{t,j}$, $(t, j) \in I_x$, and β_1, \dots, β_s is ignored—at least as far as algorithmic design is concerned, then the machinery of section 4.3.1 becomes applicable. This approach requires the estimation of the expectations of the auxiliary quantities $\hat{x}_{t,j} = \alpha_{t,j} + \sum_{i=s'}^s u_{i,t,j} \beta_i$, $(t, j) \in I_x$, as well as the estimation of their covariance

matrix. If the deviations of $\hat{x}_{t,j}$, $(t, j) \in I_x$, from their expectations fit the framework in section 3.4, then the covariance estimation may proceed via the approach of section 3.5.1 with an implementation as in <3.7>. Section 3.5.3—in particular, proposition 3.13—reveals that the latter estimation strategy is in principle robust against departures from the autoregressive dynamics. The estimate $\hat{\Theta}$ and an estimate $\hat{\rho}$ of the variance parameter ρ lead to—as explained in section 4.3.1 in connection with <3.12>—an approximate representation of the form <4.21a> of the deviations $\bar{x}_{t,j} - \mathbb{E}\hat{x}_{t,j}$, $(t, j) \in I_x$, wherein $\bar{x}_{t,j} = x_{t,j} - \sum_{i < s'} u_{i,t,j} \beta_i$. The final prediction of $x_{t,j}$ is then obtained by adding the estimate of the expectation term $\mathbb{E}\hat{x}_{t,j}$ as well as $\sum_{i < s'} u_{i,t,j}(\omega) \hat{\beta}_i$ to the predictions of the deviations $\bar{x}_{t,j} - \mathbb{E}\hat{x}_{t,j}$. Section 4.3.2 assesses the accuracy of this approach. However, departures from the autoregressive dynamics, the—possibly inappropriate—use of a simple innovation covariance matrix $\rho^2 I$ with I symbolizing the $m \times m$ identity matrix, and the overall sampling uncertainty are not reflected by its key inequality <4.14>.

Comments and references

Section 1.1 The *Statistisches Bundesamt* (national statistical bureau) provides the national **cpi** figures for all years as well as state **cpi**s starting with 1995 on their homepage alongside a product description ([Statistisches Bundesamt, 2012](#)). State **cpi**s for the years 1993 and 1994 are available in the statistical yearbooks of the *Statistische Landesämter* (statistical bureaus of the states) but with respect to the basis year 1991. The available indexes amount to the quantities in <1.1a> multiplied by 100, but this scaling is irrelevant here. [Elbel \(1995\)](#), [Elbel \(1999\)](#), [Egner \(2003\)](#), [Elbel and Egner \(2008\)](#), and [Egner \(2013\)](#) discuss the changes in methodology and weights of the index revisions corresponding to relevant basis years 1991, 1995, 2000, 2005, and 2010, respectively.

[Ströhl \(1994\)](#) describes the price survey and the associated **scpi** calculation for $t = 1993$. This studies provides separate **scpi**s for East Berlin and West Berlin. Herein, a weighted average of both indexes provides Berlin’s **scpi**. The weights $w_{\text{West Berlin}}$ and $w_{\text{East Berlin}}$ derive from solving an overdetermined system of linear equations

$$\text{cpi}_{\text{West Berlin},t} w_{\text{West Berlin}} + \text{cpi}_{\text{East Berlin},t} w_{\text{East Berlin}} = \text{cpi}_{\text{Berlin},t}, \quad t \in \{1991, \dots, 1997\},$$

by least-squares. Therein, $\text{cpi}_{\text{West Berlin},t}$, $\text{cpi}_{\text{East Berlin},t}$, and $\text{cpi}_{\text{Berlin},t}$ denote the **cpi** for West Berlin, East Berlin, and the entire city for year t and basis year 1991. These numbers are provided in [Statistisches Landesamt Berlin \(1997, sec. XVIII, p. 524–526\)](#) and [Statistisches Landesamt Berlin \(1998, sec. XVIII, p. 520–522\)](#); the resulting residuals are below the accuracy of the published **cpi**s. [Kawka \(2010\)](#) documents the non-official price survey during 2005–2009. This study concerns different spatial entities (*Kreisregionen*). These entities either coincide with a county or equal the merger of two counties; in the latter case, the published **scpi** is used for both counties.

Section 1.2 [Roos \(2006\)](#) considers the prediction of price indexes in the German setting and based on a subset of the data presented in section 1.1. However, his predictions mainly derive from a linear model (as in <1.6>) for the **scpi**s available in $t = 1993$

and with least-squares estimates replacing the (unknown) coefficients. Kosfeld et al. (2008) and Blien et al. (2009) extend and refine this approach. Hill (2016, sec. 2) models available price indexes via an underlying quantity in a similar way as in <1.2>. Rao et al. (2010) present a comparable—to the approach outlined in section 1.2—prediction strategy but based on more “parametric assumptions” than are need in section 1.2.

- Blien, U., H. Gartner, H. Stüber, and K. Wolf (2009). Regional price levels and the agglomeration wage differential in western Germany. *The Annals of Regional Science* 43(1), 71–88.
- Egner, U. (2003). Umstellung des Verbraucherpreisindex auf Basis 2000. *Wirtschaft und Statistik* 5(2003), 423–432.
- Egner, U. (2013). Verbraucherpreisstatistik auf neuer Basis 2010. *Wirtschaft und Statistik* 5, 329–344.
- Elbel, G. (1995). Zur Neuberechnung des Preisindex für die Lebenshaltung auf Basis 1991. *Wirtschaft und Statistik* 11(1995), 801–809.
- Elbel, G. (1999). Die Berechnung der Wägungsschemata für die Preisindizes für die Lebenshaltung. *Wirtschaft und Statistik* 3, 171–178.
- Elbel, G. and U. Egner (2008). Verbraucherpreisstatistik auf neuer Basis 2005. *Wirtschaft und Statistik* 4(2008), 339–50.
- Hill, R. J. (2016). A least squares approach to imposing within-region fixity in the international comparisons program. *Journal of Econometrics* 191, 407–413.
- Kawka, R. (2010). Regionale Preisunterschiede in den alten und neuen Ländern. *ifo Dresden berichtet* 17(02), 5–16.
- Kosfeld, R., H.-F. Eckey, and J. Lauridsen (2008). Disparities in prices and income across German NUTS 3 regions. *Applied Economics Quarterly* 54(2), 123–141.
- Rao, P. D. S., A. Rambaldi, and H. Doran (2010). Extrapolation of purchasing power parities using multiple benchmarks and auxiliary information: A new approach. *Review of Income and Wealth* 56, 59–98.
- Roos, M. (2006). Regional price levels in Germany. *Applied Economics* 38(13), 1553–1566.
- Statistisches Bundesamt (2012, July). *Consumer Price Index: Quality Report*. Wiesbaden: Statistisches Bundesamt.
- Statistisches Landesamt Berlin (1997). *Statistisches Jahrbuch 1997*. Statistisches Landesamt Berlin.
- Statistisches Landesamt Berlin (1998). *Statistisches Jahrbuch 1998*. Statistisches Landesamt Berlin.
- Ströhl, G. (1994). Zwischenörtlicher Vergleich des Verbraucherpreisniveaus in 50 Städten. *Wirtschaft und Statistik* 6(1994), 415–434.

2. Euclidean space basics

2.1. Fundamentals

2.1.1. Finite dimension and linearity

Let y_1, \dots, y_k be k real-valued functions defined on a common set Ω . Further functions arise by pointwise addition and multiplication with real numbers as in $(\sum_{i \leq k} c_i y_i)(\omega) = \sum_{i \leq k} c_i y_i(\omega)$ for $\omega \in \Omega$. Such a weighted sum is called a *linear combination* of y_1, \dots, y_k . The set of all linear combinations of y_1, \dots, y_k forms a real linear/vector space V , which is referred to as the *span* $\text{span}\{y_1, \dots, y_k\}$ of the functions y_1, \dots, y_k .

linear
combination
span

The sequence y_1, \dots, y_k is called a *spanning sequence* of V . It provides a *basis* of this space if its elements y_1, \dots, y_k are *linearly independent*, that is, $\sum_{i \leq k} c_i y_i = 0$ implies $c_1 = c_2 = \dots = c_k = 0$. The coefficients c_1, \dots, c_k of a linear combination $\sum_{i \leq k} c_i y_i$ are called its *coordinates* with respect to y_1, \dots, y_k . Elements of V exhibit multiple such coordinate sequences unless the sequence y_1, \dots, y_k forms a basis of V .

basis
spanning
sequence
linearly
independent
coordinates

If V is *nontrivial*, that is, $V \neq \{0\}$, then all spanning sequences contain a basis as a subsequence. Usually several choices lead to a basis, but all of these subsequences share the number of their elements, which is called the *dimension* $\dim V (\geq 1)$ of V . If $V = \{0\}$, then one strategically sets $\dim\{0\} = 0$. In either case, the dimension $\dim V$ of V does not exceed the integer k , and V is therefore termed *finite dimensional*. Conversely, a sequence of linearly independent elements x_1, \dots, x_q of some finite dimensional linear space W may be extended to a basis of W , which then implies the inequality $\dim W \geq q$.

nontrivial
dimension
finite
dimensional

Below the concept of a finite dimensional linear space W is met in the guise of

- (a) elements of $c \in \mathbb{R}^k$, which are real-valued functions on $\{1, \dots, k\}$. The set \mathbb{R}^k equals $\text{span}\{e_1, \dots, e_k\}$, wherein $e_i(j) = 0$ unless $i = j$ and then $e_i(i) = 1$. The *standard basis* e_1, \dots, e_k of \mathbb{R}^k is a basis of this space.
- (b) Elements of the set of real $m \times k$ matrices $\mathbb{R}^{m \times k}$ are real-valued functions on $\Omega = \{1, \dots, m\} \times \{1, \dots, k\}$. The matrices $B_{i,j} \in \mathbb{R}^{m \times k}$, $(i, j) \in \Omega$, $B_{i,j}(p, q) = 0$ if $(p, q) \neq (i, j)$, $B_{i,j}(i, j) = 1$, form the *standard basis* of this mk dimensional space.

standard basis

standard basis

The term *vector* is reserved for $c \in \mathbb{R}^k$; singular vectors—see section 2.5—provide an exception to this rule. The image $c(i)$ of i under c is usually denoted by c_i and called the i -th entry of c . The latter suggests suggest writing $c = (c_1, \dots, c_k)$. The entry c_i also equals the i -th coordinate of c with respect to the standard basis defined in (a).

vector

Likewise if $A \in \mathbb{R}^{m \times k}$, then its i, j -th entry is given by $A(i, j) = a_{i,j}$, that is, the coordinate of A with respect to the i, j -th standard basis element $B_{i,j}$ shown in (b). The entries $a_{i,j}$ are displayed in the common array/block form when considered jointly. In that case, zero entries—if identified as such by the context—are replaced by white space.

If needed, in particular for the purpose of matrix-vector multiplication, $c \in \mathbb{R}^k$ is identified with its corresponding element in $\mathbb{R}^{k \times 1}$. Accordingly, the matrix representation of a linear map $f : \mathbb{R}^k \rightarrow \mathbb{R}^m$ is with respect to the standard bases of \mathbb{R}^k and \mathbb{R}^m . That is, the entries $a_{1,j}, \dots, a_{m,j}$ of the matrix representation $A \in \mathbb{R}^{m \times k}$ of f are given by the coordinates of the image $f(e_j)$ of $e_j \in \mathbb{R}^k$ with respect to the standard basis of \mathbb{R}^m .

A nonempty subset U of W is termed a *subspace* (of W) if it includes all linear combinations of its elements. Then U —considered in isolation—forms a linear space and the above terminology applies in analogy. In addition, W is referred to as a *superspace* of U . An important subspace of the *quadratic* matrices $\mathbb{R}^{m \times m}$ consists of

- (c) the *symmetric* matrices \mathbb{S}^m , that is, matrices $A \in \mathbb{R}^{m \times m}$ which satisfy $a_{i,j} = a_{j,i}$. The matrices $\bar{B}_{i,i} = B_{i,i}$, $i \leq m$, together with the matrices $\bar{B}_{i,j} = (B_{i,j} + B_{j,i})/\sqrt{2}$, $i < j \leq m$, provide the *standard basis* of this $m(m+1)/2$ dimensional space.

Every element $A \in \mathbb{R}^{m \times k}$ induces a linear map $\mathbb{R}^k \ni c \mapsto Ac = \sum_{i \leq k} c_i a_i \in \mathbb{R}^m$, which—at least with respect to the complexity of the notation—facilitates the study of the linear relations between its k *columns* $a_i = Ae_i$, $i \leq k$. Replacing these functions $a_i : \{1, \dots, m\} \rightarrow \mathbb{R}$ (see (a)) by real-valued functions y_1, \dots, y_k on a set Ω lifts this amenity to a higher level of generality. This construction generates a linear map Y from \mathbb{R}^k to a superspace W of span $\{y_1, \dots, y_k\}$. In fact, all linear maps $X : \mathbb{R}^k \rightarrow W$ are of the form $c \mapsto \sum_{i \leq k} c_i x_i$ with *columns* $x_i = Xe_i \in W$ and admit the structure of a linear space:

- (d) the real vector space $W^{\times k}$ with pointwise defined linear operations, that is, $(aY + bY')c = a(Yc) + b(Y'c)$ for $a, b \in \mathbb{R}$, $c \in \mathbb{R}^k$, and $Y, Y' \in W^{\times k}$.

This linear space conforms to the above general framework if $Y \in W^{\times k}$ is identified with $(\omega, i) \mapsto (Ye_i)(\omega)$ as in (b). The second perspective rightly suggests that $W_{i,j}$, wherein $W_{i,j}e_k$ equals the i -th element of a basis of the m' dimensional space W for $j = k$ and zero otherwise, form a basis of this $m'k$ dimensional space.

In general, elements of $W^{\times k}$ are symbolized by uppercase letters; corresponding lowercase letters represent the corresponding columns; and the block notation $Y = [y_1 \ \cdots \ y_k]$ honors the equality $\mathbb{R}^{m \times k} = (\mathbb{R}^m)^{\times k}$. Moreover, writing $Y = [Y_1 \ Y_2]$ with $Y_1 = [y_1 \ \cdots \ y_j]$, $Y_2 = [y_{j+1} \ \cdots \ y_k]$, and $j < k$ mimics the usual notation for partitioned matrices. Finally, the identification $W^{\times 1} = W$ resembles the above mentioned case $\mathbb{R}^{k \times 1} = \mathbb{R}^k$.

Two subspaces directly derive from $Y \in W^{\times k}$, namely, its *image* $\text{img } Y \subset W$, which is also referred to as its *column space*, and its *kernel* $\ker Y \subset \mathbb{R}^k$. The kernel or *null space* gathers all $c \in \mathbb{R}^k$ with $Yc = 0$. The column space/image consists of all linear combinations of its columns y_1, \dots, y_k . The dimension of its image is known as the *rank* $\text{rk } Y$ of Y . The latter satisfies $\text{rk } Y + \dim \ker Y = k$. If $\ker Y \neq \{0\}$, then this equality results from the ability to extend a basis of $\ker Y$ to a basis of \mathbb{R}^k . Otherwise, it holds by definition as $\ker Y = \{0\}$ is tantamount to linear independence of y_1, \dots, y_k .

Another relevant example of a finite dimensional linear space comes in the form of

- (e) the span of a finite sequence of \mathbb{P} -square integrable *random variables* on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, that is, real-valued and $\mathcal{F}/\mathcal{R}^1$ -measurable functions y on Ω with $\int y^2(\omega) \mathbb{P}(d\omega) < \infty$. Herein, \mathcal{R}^1 denotes the Borel σ -field corresponding to the $|\cdot|$ -topology—the smallest σ -field containing all open intervals.

The dimension of this space is bounded from above by the sequence length. Moreover, all elements of this linear space, that is, all linear combinations of the spanning sequence elements, are \mathbb{P} -square integrable.

The examples (b)–(e) as well as all other finite dimensional real linear spaces V mimic example (a) to the extent that elements of V share their linear relations with their coordinate vectors with respect to a basis. More formally, a basis y_1, \dots, y_k leads to a bijective linear map $Y = [y_1 \ \cdots \ y_k]$ from \mathbb{R}^k to V whose inverse map Y^{-1} is linear.

2.1.2. Norm topology

A norm $\|\cdot\|$ endows every element x of a finite dimensional linear space W with a length $\|x\|$. A pair consisting of a finite dimensional linear space and a norm forms a finite dimensional normed space. Relevant examples—numbered as in section 2.1.1—include

- (a) the real k -tuples \mathbb{R}^k together with the *Euclidean norm* $\|c\| = (\sum_{i \leq k} c_i^2)^{1/2}$, *Euclidean norm*
- (b) the real $m \times k$ -matrices with the *Frobenius norm* $\|A\| = (\sum_{i \leq m, j \leq k} a_{i,j}^2)^{1/2}$, and *Frobenius norm*
- (e) the span of a finite set of \mathbb{P} -square integrable random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ joined by the *L^2 -norm* $\|x\| = (\int x(\omega)^2 \mathbb{P}(d\omega))^{1/2}$. *L^2 -norm*

In this example, the existence of $F \in \mathcal{F}$ with $\mathbb{P}F = 0$ potentially reduces $\|\cdot\|$ to a seminorm. Appendix 2.a deals with this nuisance.

Further important instances of finite dimensional normed spaces are best discussed in connection with singular values; their treatment is deferred to section 2.5.2.

A norm induces a metric $d(x, y) = \|x - y\|$, which endows every finite dimensional normed space W with a metric space structure. The resulting metric space (W, d) exhibits a few notable features. More specifically, the finite dimension guarantees that linear maps are continuous and subspaces are closed. Furthermore, the Heine-Borel theorem asserts that closed and bounded subsets of such spaces are (sequentially) compact. The *unit sphere* $\{\|\cdot\| = 1\}$ provides the most relevant example. *unit sphere*

These properties are inherited from \mathbb{R}^k . In fact, if y_1, \dots, y_k are a basis of W , then $Y = [y_1 \ \cdots \ y_k]$ forms a bijective and continuous (linear) map with continuous (and linear) inverse Y^{-1} . The two maps Y and Y^{-1} transfer topological properties from W to the coordinates with respect to y_1, \dots, y_k in \mathbb{R}^k and vice versa.

Compactness of $\{\|\cdot\| = 1\}$ ensures that $\{\ell(y) \mid \|y\| = 1\}$ is (sequentially) compact whenever $\ell : \{\|\cdot\| = 1\} \rightarrow \mathbb{R}$ is continuous with respect to the metric d . As a consequence, $\{\ell(y) \mid \|y\| = 1\}$ contains its supremum and infimum. Any other norm $\|\cdot\|'$ on W provides an example of a continuous (with respect to d) function. Hence, there exists a *lower compatibility constant* $c = c(W, \|\cdot\|, \|\cdot\|') > 0$ and an *upper compatibility constant* $C = C(W, \|\cdot\|, \|\cdot\|') > 0$ such that $c\|\cdot\| \leq \|\cdot\|' \leq C\|\cdot\|$. These numbers usually depend on the dimension of the linear space W . Nonetheless, $\|\cdot\|$ and $\|\cdot\|'$ and thereby their induced metrics are (strongly) equivalent. Thus, any two norms on W induce the same topology, which is called the *norm topology*. Most notably, all metric notions such as convergence and continuity coincide across norms on W . *lower compatibility constant*
upper compatibility constant
norm topology

Compactness is key to many qualitative results; quantitative statements require a quantitative analog. If $\varepsilon > 0$, then an ε -*net* of a compact subset S of W is a finite subset $\{x_1, \dots, x_q\}$ of S such that $S \subset \cup_{i \leq q} \{\|x_i - \cdot\| \leq \varepsilon\}$. Compactness of S implies *ε -net*

total boundedness, which ensures the existence of ε -nets. The well-ordering principle guarantees the existence of a minimal number of elements of S needed to form an ε -net for a given $\varepsilon > 0$. This number is referred to as the ε -covering number $N = N(S, \|\cdot\|, \varepsilon)$ of S . In case of the unit sphere of \mathbb{R}^k , the translation invariance and scaling property of the Lebesgue measure lead to the upper bound on $N(\{\|\cdot\| = 1\}, \|\cdot\|, \varepsilon)$ in the following lemma, which is proved on page 39 of appendix 2.b.

ε -covering
number

Lemma 2.1. *For any $\varepsilon \in (0, 1)$, the ε -covering number $N(\{\|\cdot\| = 1\}, \|\cdot\|, \varepsilon)$ of the unit sphere $\{\|\cdot\| = 1\}$ of \mathbb{R}^k with $\|\cdot\|$ as in (a) satisfies*

$$N(\{\|\cdot\| = 1\}, \|\cdot\|, \varepsilon) \leq \left(1 + \frac{2}{\varepsilon}\right)^k .$$

Covering numbers and ε -nets serve to bound the supremum and infimum of an infinite set by the maximum and minimum over a finite ε -net. In particular, if $\ell : S \rightarrow \mathbb{R}$ exhibits a finite $(\|\cdot\|)$ -Lipschitz constant $L = \sup_{x \neq y} |\ell(x) - \ell(y)| / \|x - y\|$ and $\{x_1, \dots, x_N\}$ denotes a (\mathbb{C}) -minimal ε -net of S , then it follows that

Lipschitz
constant

$$\inf_{x \in S} \ell(x) \geq \min_{i \leq N} \ell(x_i) - L\varepsilon \quad \text{and} \quad \sup_{x \in S} \ell(x) \leq \max_{i \leq N} \ell(x_i) + L\varepsilon . \quad <2.1>$$

2.1.3. Geometry

The similarity in appearance of the Euclidean, the Frobenius, and the L^2 -norm is no coincidence. In fact, in all these cases the respective domain Ω of the functions exhibits a natural measure space interpretation with finite measure μ such that elements y_1, \dots, y_k of the resulting function space are μ -square-integrable:

- (a) the set $\{1, \dots, m\}$ coupled with its power set and the counting measure leads to the integral $\int x(i)\mu(di) = \sum_{i \leq m} x_i$ underlying the Euclidean norm.
- (b) The identical construction extended to fit $\mathbb{R}^{m \times k}$ produces $\int A(i, j)\mu(d(i, j)) = \sum_{i \leq m, j \leq k} a_{i, j}$ and thus the Frobenius norm. Such integration is also feasible if
- (c) the symmetric $m \times m$ matrices \mathbb{S}^m —a subset of $\mathbb{R}^{m \times m}$ —are considered in isolation.

The measure space structure moreover suggests an inner product $\langle \cdot, \cdot \rangle$ defined by $\langle x, y \rangle = \int x(\omega)y(\omega)\mu(d\omega)$ for any two elements x, y . This inner product induces the corresponding of the just mentioned norms via $\|x\| = \sqrt{\langle x, x \rangle}$. In this text, the symbol $\|\cdot\|$ is used exclusively for norms following this recipe. Herein, the finiteness of the measure μ is dispensable but convenient. Appendix 2.a states its relevant implications.

More generally, if W equals the span of a finite sequence of real-valued μ -square-integrable functions on a (finite) measure space $(\Omega, \mathcal{F}, \mu)$, then interpreting

- (d) elements Y of $W^{\times k}$ as functions $(\omega, i) \mapsto (Ye_i)(\omega) = y_i(\omega)$ allows their integration using the product ν of μ with the counting measure on $\{1, \dots, k\}$. By Fubini's theorem, the resulting integral is $\int y_i(\omega)\nu(d(\omega, i)) = \sum_{i \leq k} \int y_i(\omega)\mu(d\omega)$ and induces the inner product $\langle Y, Y' \rangle = \sum_{i \leq k} \int y_i(\omega)y'_i(\omega)\mu(d\omega) = \sum_{i \leq k} \langle y_i, y'_i \rangle$.




Figure 2.1

The figure visualizes the relation between the geometric notions of angles and length to inner products and norms. Panel (A) links the Euclidean norm $\|x\|$ to the length a of x . Panel (B) connects the cosine $\cos \theta$ of the (small) angle θ between x and y to their inner product.

If $W = \mathbb{R}^m$, then the inner product in (d) recovers the inner product underlying the Frobenius norm in (b). In particular, (d) points to the alternative expression $\langle A, B \rangle = \text{tr}(A^T B)$ for $A, B \in \mathbb{R}^{m \times k}$. Therein, a superscript T marks the *transpose* A^T of A —given by $A^T(i, j) = a_{j, i}$. In addition, tr denotes the *trace*—the sum of the *diagonal entries* $(A^T B)_{i, i}$, whose *cyclic property* $\text{tr}(A^T B) = \text{tr}(B A^T)$ is often used without further notice.

transpose
trace
diagonal entries
cyclic property
Euclidean space

In general, finite dimensional real linear spaces furnished with an inner product are called *Euclidean spaces*. Such spaces are finite dimensional normed spaces furnished with a notion of (small) angle between any two of their elements. The two concepts—length and (small) angle—give rise to a simple geometry, which parallels that of the plane \mathbb{R}^2 .

Figure 2.1 considers two nonzero elements x and y of the plane \mathbb{R}^2 with inner product as in (a). This space provides the archetypal Euclidean space, wherein length and angles occur in the usual sense. In Panel (A), Pythagoras’s theorem shows that $\|x\| = \sqrt{x_1^2 + x_2^2} = a$ coincides with the traditional understanding of length. Panel (B) illustrates the relations $\|y - x\|^2 = a^2 + c^2$, $\|y\|^2 = a^2 + b^2$, $\|x\| = b + c$, and thereby $b = \langle x, y \rangle / \|x\|$. Figure 2.2 ties b to the cosine of the (small) angle θ between x and y . Its panel (A) illustrates the cosine of the (small) angle θ' between a point z of length r and the first standard basis element e_1 . Therein, dropping a perpendicular from z to the first coordinate axis yields a right triangle such that the length of its leg adjacent to θ' equals $r \cos \theta'$. A comparison with panel (B) of figure 2.1 shows that the cosine of θ equals $\cos \theta = b / \sqrt{a^2 + b^2} = b / \|y\|$, and therefore one has $\cos \theta \|x\| \|y\| = \langle x, y \rangle$.

Figure 2.2 also contains a geometric characterization of the sine and tangent of θ' . In particular, the relation $\cos^2 \theta + \sin^2 \theta = a^2 / \|y\|^2 + b^2 / \|y\|^2 = 1$ is notable.

If x, y are nonzero, then the relation $\cos \theta \|x\| \|y\| = \langle x, y \rangle$ determines the value of the angle θ . Panel (A) of figure 2.2 shows that a (small) angle θ' is an element of $[0, \pi]$. As z moves along the upper half of the r -sphere in counterclockwise direction, $\cos \theta'$ decreases steadily from 1 to -1 . More specifically, \cos is continuous and monotone decreasing—thus bijective—on $[0, \pi]$. The neighboring panel (B) shows the relevant part of its graph. The




Figure 2.2

The figure shows the geometric significance of the functions \cos , \sin , \tan and their graphs on the intervals $[0, \pi]$, $[0, \pi]$, and $[0, \pi/2)$, respectively. Panel (A) expresses the cosine of the angle θ' between a point z of length r and the first standard basis element e_1 in terms of the first coordinate of z with respect to the standard basis. The sine $\sin \theta$ and tangent $\tan \theta$ exhibit similar representations. Panel (B) contains parts of the graphs of the resulting continuous functions \cos , \sin , and \tan on the above mentioned intervals.

latter indicates that the equalities $\langle x, y \rangle = \|x\| \|y\|$ and $\langle x, y \rangle = -\|x\| \|y\|$ occur if and only if x and y reside on the same line through the origin. Moreover, an angle of $\pi/2$ is tantamount to $\langle x, y \rangle = 0$. Then x and y are said to be *orthogonal*, and this term is also applied if $y = 0$. Thereby 0 becomes the sole element being orthogonal to all $x \in \mathbb{R}^2$.

orthogonal

The relation of $\langle \cdot, \cdot \rangle$ with the geometric concepts of length and angles in \mathbb{R}^2 allows to transport these notions to more general spaces W . Bijective linear maps Q with $\langle Qc, Qc' \rangle = \langle c, c' \rangle$ for any two points c, c' —so-called *unitary maps*—provide the bridge.

unitary map

Section 2.2.2 constructs a unitary map Q from \mathbb{R}^h to an abstract h dimensional Euclidean space V . Consequently, two elements x, y of an at least two dimensional Euclidean space W may be identified with their preimages under a unitary map Q from \mathbb{R}^2 to a two dimensional superspace V of their span $\text{span}\{x, y\}$. Then, the equality $\langle x, y \rangle = \langle Q^{-1}x, Q^{-1}y \rangle$ suggest thinking of $\|x\| = \sqrt{\langle x, x \rangle} = \|Q^{-1}x\|$ and the number θ , which is defined for nonzero elements x, y by $\cos \theta \|x\| \|y\| = \langle x, y \rangle$, as the length of x and the (small) angle between x and y , respectively. This identification forces two dimensional subspaces of W to resemble the geometry of \mathbb{R}^2 .

2.2. Unitary maps

2.2.1. Orthonormal bases

The interpretation of $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ in section 2.1 characterizes unitary maps as linear bijections preserving length and angles. The *polarization identity* $\langle x, y \rangle = (\|x + y\|^2 - \|x - y\|^2)/4$ shows that the preservation of length implies preservation of angles.

polarization identity

The sequence of preimages $Q^{-1}y_1, \dots, Q^{-1}y_k$ of Euclidean space elements y_1, \dots, y_k under a unitary map Q forms a *representation* of y_1, \dots, y_k . Economy of notation excuses using this term for the composition $Q^{-1}Y$ in connection with the linear map $Y = [y_1 \cdots y_k]$. The identity map on any Euclidean space is unitary; the same applies to the inverse of any unitary map and the composition of any two suitable unitary maps.

representation

Let V be a $h(\geq 1)$ dimensional Euclidean space and Q be unitary from \mathbb{R}^h to V . Then Q lies in $V^{\times h}$ and its columns—the images $q_i = Qe_i$ of the standard basis elements e_1, \dots, e_h of \mathbb{R}^h under Q —exhibit unit length and are pairwise orthogonal. Conversely, if q'_1, \dots, q'_h are pairwise orthogonal and of unit length, then $\langle \sum_{i \leq h} c_i q'_i, \sum_{i \leq h} c'_i q'_i \rangle = \langle c, c' \rangle$ ensures linear independence of q'_1, \dots, q'_h and identifies $Q' = [q'_1 \cdots q'_h]$ as a unitary map from \mathbb{R}^h to $V' = \text{span}\{q'_1, \dots, q'_h\}$. In particular, this sequence forms a basis of V' , and $q'_i = Q'e_i$ may be thought of as (perpendicular) coordinate axes in V' .

Such bases—consisting of unit length and pairwise orthogonal elements—are called *orthonormal basis*. The standard bases in (a), (b), and (c) are of this kind. Therein, the standard basis of $S^m \subset \mathbb{R}^{m \times m}$ may be grown into another orthonormal basis of $\mathbb{R}^{m \times m}$ by adding the $m(m-1)/2$ elements $B_{i,j} = (B_{i,j} - B_{j,i})/\sqrt{2}$, $j < i$. The latter span the space of *skew-symmetric* matrices, that is, quadratic matrices A with $A^T = -A$.

orthonormal basis

skew-symmetric

Section 2.2.2 presents a recipe—called *Gram-Schmidt orthogonalization*—for the extraction of an orthonormal basis of a h dimensional space from a spanning sequence y_1, \dots, y_k . Lemma 2.2 summarizes this finding. Its statement relies on the concept of a *row echelon matrix* $R \in \mathbb{R}^{h \times k}$. Such matrices satisfy $I_i = \{j \leq k \mid r_{i,j} \neq 0\} \neq \emptyset$, $i \leq h$, and $\inf I_i < \inf I_{i+1}$, $i < h$. Thus, $h \leq k$, and the rows of R —interpreted as elements of \mathbb{R}^k —are linearly independent. If y_1, \dots, y_k form a basis, that is, $h = k$, then R is an *upper triangular matrix*—its entries satisfy $r'_{i,j} = 0$ if $i > j$ —with nonnegative diagonal entries.

Gram-Schmidt orthogonalization

row echelon matrix

upper triangular matrix

Lemma 2.2. *Let y_1, \dots, y_k be a nontrivial sequence of Euclidean space elements, that is, with $y_j \neq 0$ for at least one $j \leq k$, spanning a $h(\geq 1)$ dimensional space V . There exists an orthonormal basis q_1, \dots, q_h of V and a row echelon matrix R such that*

nontrivial

$$\underbrace{[y_1 \cdots y_k]}_Y = \underbrace{[q_1 \cdots q_h]}_Q \underbrace{\begin{pmatrix} \overbrace{r_{1,1}}^{\text{coordinates of } y_1} & \overbrace{r_{1,2}}^{\text{coordinates of } y_2} & \cdots & \overbrace{r_{1,h}}^{\text{coordinates of } y_h} & \cdots & \overbrace{r_{1,k}}^{\text{coordinates of } y_k} \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \overbrace{r_{h,h}}^{\text{coordinates of } y_h} & \cdots & \overbrace{r_{h,k}}^{\text{coordinates of } y_k} \end{pmatrix}}_R \begin{matrix} \left. \begin{matrix} r_{1,1} \\ \vdots \\ r_{1,k} \end{matrix} \right\} \text{with resp.} \\ \left. \begin{matrix} r_{2,2} \\ \vdots \\ r_{2,k} \end{matrix} \right\} \text{to } q_1 \\ \left. \begin{matrix} \vdots \\ r_{h,h} \\ \vdots \\ r_{h,k} \end{matrix} \right\} \text{to } q_2 \\ \left. \begin{matrix} r_{h,h} \\ \vdots \\ r_{h,k} \end{matrix} \right\} \text{to } q_h \end{matrix} \quad . \quad \langle 2.2 \rangle$$

2.2.2. Gram-Schmidt orthogonalization

The discussion initially stresses geometric intuition and considers linearly independent $x, y \in \mathbb{R}^2$ and $x, y, z \in \mathbb{R}^3$ for clarity. The general case is tackled afterwards.

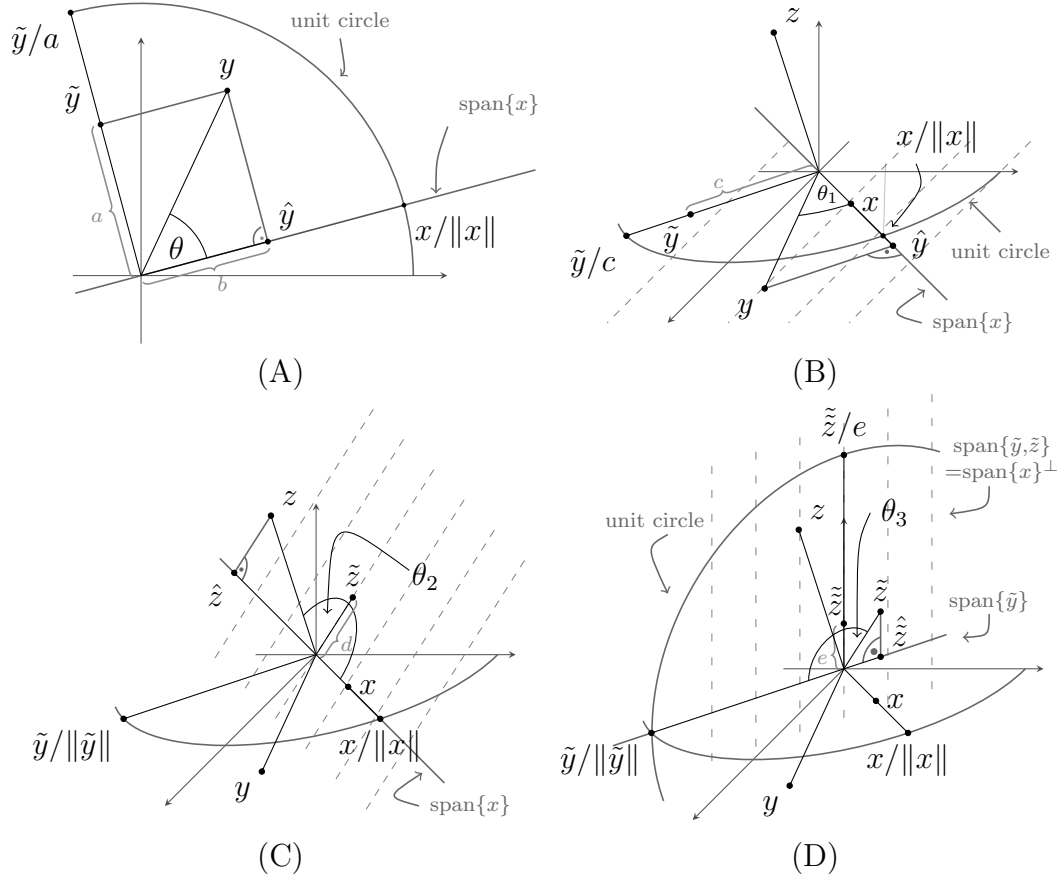


Figure 2.3

The figure illustrates the sequential construction of an orthonormal basis. Panel (A) considers a basis x and y of \mathbb{R}^2 . The transformation has two steps: first x is scaled to form an orthonormal basis of $\text{span}\{x\}$; then \tilde{y} is determined as the part of y orthogonal to $\text{span}\{x\}$; a final scaling to obtain \tilde{y}/a completes the orthogonalization. Panel (B) to (D) repeat this construction in a recursive manner to transform a basis x, y and z of \mathbb{R}^3 into an orthonormal basis $q_1 = x/\|x\|$, $q_2 = \tilde{y}/\|\tilde{y}\|$, and $q_3 = \tilde{z}/\|\tilde{z}\|$ of that space.

Panel (B) of figure 2.1 suggests how to construct q_1 and q_2 from an arbitrary basis x, y of \mathbb{R}^2 . In fact, the second basis element y equals the sum of a multiple of x with length $b > 0$ and an element \tilde{y} of length $a > 0$, which is orthogonal to x . Thus, x and \tilde{y} provide suitable elements q_1 and q_2 up to scaling. Panel (A) of figure 2.3 extends panel (B) in figure 2.1 to illustrate this approach. Therein, x is scaled to $q_1 = x/\|x\|$; the alternative choice $-x/\|x\|$ is not displayed. Subsequently, panel (A) decomposes y into two mutually orthogonal parts $\hat{y} \in \text{span}\{x\} = \text{span}\{q_1\}$ and an element \tilde{y} of $(\text{span}\{x\})^\perp$. The first part equals $r_{1,2}q_1$, wherein $r_{1,2} \in \mathbb{R}$ is given by $\|r_{1,2}q_1\| = r_{1,2} = b = \langle q_1, y \rangle$ as in section 2.1. Consequently, the second part results as $\tilde{y} = y - q_1 \langle q_1, y \rangle$ with length a , which reveals the two possible choices $q_2 = +\tilde{y}/a$ and $q_2 = -\tilde{y}/a$ for extending q_1 to an orthonormal basis of $\text{span}\{x, y\}$. Panel (A) presents the case $q_2 = \tilde{y}/a = \tilde{y}/\|\tilde{y}\|$.

Every element of \mathbb{R}^2 has a unique representation as the sum of an element of $\text{span}\{q_1\}$

and one from $\text{span}\{q_2\}$. In this sense, $\text{span}\{q_2\}$ *complements* $\text{span}\{q_1\} = \text{span}\{x\}$ in \mathbb{R}^2 . An analogous consideration identifies $\text{span}\{y\}$ as another complement of $\text{span}\{x\}$. However, each element of $\text{span}\{q_2\}$ is orthogonal to every element of $\text{span}\{q_1\}$, and it is the only complement with this property. Speaking of $\text{span}\{q_2\}$ as the *orthogonal complement* $(\text{span}\{x\})^\perp$ of $\text{span}\{x\}$ in \mathbb{R}^2 is therefore legitimate. This subspace relation exhibits symmetry to the extent that $((\text{span}\{x\})^\perp)^\perp = \text{span}\{x\}$.

complement

orthogonal complement

Pythagoras's theorem guarantees the inequality $\|y - \hat{y}\| \leq \|y - x\|$ in panel (B) of figure 2.1. The latter remains valid if x is replaced by any element of $\text{span}\{x\}$; it is strict unless one chooses \hat{y} as a replacement, which uniquely characterizes \hat{y} . In addition, the use of Pythagoras's theorem implies that $\hat{y} \in \text{span}\{x\}$ together with $\tilde{y} = y - \hat{y} \in (\text{span}\{x\})^\perp$ offers an equivalent characterization. The latter endows \hat{y} with its name *orthogonal projection* of y onto $\text{span}\{x\}$. The symmetry of the second characterization ensures that the *residual* \tilde{y} equals the orthogonal projection of y onto $(\text{span}\{x\})^\perp = \text{span}\{q_2\}$.

orthogonal projection

residual

Panel (B) to (D) of figure 2.3 extend the procedure of panel (A), including the concept of orthogonal complement and projection, to \mathbb{R}^3 . This extension comprises three major steps each turning a further element of a basis x, y, z into one of q_1, q_2, q_3 . The first two major steps—shown in panel (B)—replicate the \mathbb{R}^2 version of the construction process: the first major step provides q_1 by scaling x ; the second major step consists of an orthogonalization step to obtain $\tilde{y} = y - q_1\langle q_1, y \rangle$ and thereby $q_2 = \tilde{y}/\|\tilde{y}\|$.

Panel (C) brings in z only after an orthonormal basis q_1, q_2 of $\text{span}\{x, y\}$ has been found. Two orthogonalization steps follow. The first relieves z from its dependence with x yielding $\tilde{z} = z - q_1\langle q_1, z \rangle$. The two points \tilde{y} and \tilde{z} form a basis of $(\text{span}\{x\})^\perp = \{\langle x, \cdot \rangle = 0\}$. The intermediate residual \tilde{z} equals the sum of $\hat{z} = q_2\langle q_2, z \rangle \in (\text{span}\{x\})^\perp \cap \text{span}\{\tilde{y}\}$ and $\tilde{\tilde{z}} = \tilde{z} - \hat{z} \in (\text{span}\{x\})^\perp \cap (\text{span}\{\tilde{y}\})^\perp$. The second summand $\tilde{\tilde{z}} = \tilde{z} - q_2\langle q_2, \tilde{z} \rangle \neq 0$ spans $(\text{span}\{x, y\})^\perp$. Scaling the latter to unit length completes the construction of q_3 . The individual steps of this orthogonalization are confined to two dimensional subspaces. Figure 2.3 highlights these subspaces by dashed lines.

The generalization <2.3> to k linearly independent elements y_1, \dots, y_k of a Euclidean space W yields an orthonormal basis q_1, \dots, q_k of their span. In addition, one obtains for each $j \leq k$ a coordinate vector $r_j = (r_{1,j}, \dots, r_{j,j}, 0, \dots, 0)$ —shown in <2.2>—of y_j with respect to that basis. The construction requires k major steps, which are indexed by j and consist of $j - 1$ orthogonalization steps—indexed by i —alongside a final scaling step.

$$\begin{array}{l|l}
 1 & \tilde{y}_j^{(0)} = y_j, \quad j \leq k \\
 2 & \mathbf{for} \quad j = 1, \dots, k \\
 3 & \quad \mathbf{for} \quad i = 1, \dots, j - 1 \\
 4 & \quad \quad r_{i,j} = \langle q_i, \tilde{y}_j^{(i-1)} \rangle \\
 5 & \quad \quad \tilde{y}_j^{(i)} = \tilde{y}_j^{(i-1)} - r_{i,j}q_i \\
 6 & \quad r_{j,j} = \pm \|\tilde{y}_j^{(j-1)}\| \\
 7 & \quad q_j = \tilde{y}_j^{(j-1)} / r_{j,j}
 \end{array} \tag{2.3}$$

The sign in line 6 of <2.3> may be chosen at will. The choice affects the outcome as

$\tilde{y}_j^{(j-1)}$ equals a linear combination of y_1, \dots, y_j with at least one nonzero coordinate and is therefore nonzero. Consequently, $\|\tilde{y}_j^{(j-1)}\| \neq 0$ and the final division is uncontroversial.

The unit length elements q_1, \dots, q_k generated by line 7 are linear combinations of y_1, \dots, y_k . Furthermore, the construction ensures that for $1 \leq j < j' \leq k$ one has

$$\langle q_j, \tilde{y}_{j'}^{(j)} \rangle = \left\langle q_j, \tilde{y}_{j'}^{(j-1)} - q_j \langle q_j, \tilde{y}_{j'}^{(j-1)} \rangle \right\rangle = \langle q_j, \tilde{y}_{j'}^{(j-1)} \rangle - \langle q_j, \tilde{y}_{j'}^{(j-1)} \rangle = 0. \quad \langle 2.4 \rangle$$

In terms of the above vocabulary, the equalities $\langle 2.4 \rangle$ identify $\tilde{y}_{j+1}^{(j)}, \dots, \tilde{y}_k^{(j)}$, $j < k$, as elements of the orthogonal complement of $\text{span}\{q_1, \dots, q_j\} = \text{span}\{y_1, \dots, y_j\}$. In fact, these residuals equal the orthogonal projections of y_{j+1}, \dots, y_k onto $(\text{span}\{y_1, \dots, y_j\})^\perp$ and form a basis of this orthogonal complement in $\text{span}\{y_1, \dots, y_k\}$. Subsequently constructed basis elements q_{j+1}, \dots, q_k are linear combinations of $\tilde{y}_{j+1}^{(j)}, \dots, \tilde{y}_k^{(j)}$, thus, also lie in the orthogonal complement $(\text{span}\{q_1, \dots, q_j\})^\perp$ of $\text{span}\{q_1, \dots, q_j\}$.

Lemma 2.2 allows linear dependence between y_1, \dots, y_k . In that case, a zero residual $\tilde{y}_j^{(j-1)}$ occurs in at least one major step. The latter indicates that the j -th element y_j does not add to $\text{span}\{y_1, \dots, y_k\}$ in the given order and should be omitted in the basis construction. Setting $q_j = 0$ and thereby $r_{j,k} = 0$, $k > j$, effectively skips y_j and avoids the use of a further indexing variable. Replacing the seventh line of $\langle 2.3 \rangle$ by

$$\begin{array}{l|l} 7 & \text{if } r_{j,j} = 0 \\ 8 & \quad q_j = 0 \\ 9 & \text{else} \\ 10 & \quad q_j = \tilde{y}_j^{(j-1)} / r_{j,j} \end{array}$$

implements this approach. In this form, the procedure generates a preliminary sequence $\bar{q}_1, \dots, \bar{q}_k$ of mutually orthogonal elements of W and a corresponding $k \times k$ upper triangular matrix \bar{R} . The h nonzero \bar{q}_j are of unit length and form a basis of $\text{span}\{y_1, \dots, y_k\}$. The zero \bar{q}_j correspond to zero rows in the coordinate matrix \bar{R} . Dropping these zero elements \bar{q}_i and the corresponding rows in \bar{R} yields a representation as in $\langle 2.2 \rangle$.

2.3. Projectors

An orthonormal basis q_1, \dots, q_h of a h dimensional subspace V of a Euclidean space W yields the closed form expression $\sum_{i \leq h} q_i \langle q_i, x \rangle$ for the *orthogonal projection* \hat{x} of $x \in W$ onto V . As an example, the standard basis $\bar{B}_{i,j}$, $i \leq j$, of \mathbb{S}^m —see (c) of section 2.1.1—identifies the orthogonal projection \hat{A} of $A \in \mathbb{R}^{m \times m}$ onto \mathbb{S}^m as $\sum_{i \leq j \leq m} \bar{B}_{i,j} \langle \bar{B}_{i,j}, A \rangle = (A + A^\top)/2$. Alternatively, the orthogonal projection is uniquely defined as the best approximation of x in V in the sense that $\hat{x} \in V$ and $\|x - \hat{x}\| \leq \|x - x'\|$ for all $x' \in V$. In particular, the projection \hat{x} equals x if and only if x is itself an element of V .

The approximation error $\tilde{x} = x - \hat{x}$ is called the *residual* of that projection. It lies in the *orthogonal complement* $V^\perp = \bigcap_{y \in V} \{\langle y, \cdot \rangle = 0\}$ of V in W . The skew-symmetric residual $(A - A^\top)/2$ from orthogonally projecting $A \in \mathbb{R}^{m \times m}$ onto \mathbb{S}^m exemplifies this




Figure 2.4

The figure illustrates the projection onto a subspace V of \mathbb{R}^3 along its orthogonal complement V^\perp as well as another complement U of V in \mathbb{R}^3 . Panel (A) shows the two resulting projections \hat{x}_V and $\hat{x}_{V/U}$ for some $x \notin V$. Panel (B) introduces another element y such that $x + y \in V$. The latter is projected onto V along V^\perp as well as onto V along U .

inclusion. Symmetry and bilinearity of $\langle \cdot, \cdot \rangle$ ensure that V^\perp is a subspace with $V = (V^\perp)^\perp$. Moreover, the second equality in $\|x - x'\|^2 = \|\hat{x} + \tilde{x} - x'\|^2 = \|x - \tilde{x}\|^2 + \|\tilde{x} - x'\|^2$, which holds whenever $x' \in V^\perp$, shows that \tilde{x} provides the orthogonal projection of x onto V^\perp . The equality $\sum_{i>j} \bar{B}_{i,j} \langle \bar{B}_{i,j}, A \rangle = (A - A^\top)/2$ verifies this equivalence. The above definitions and observations generalize the geometric discussion of section 2.2.2.

The linear map $\hat{\cdot} = \sum_{i \leq h} q_i \langle q_i, \cdot \rangle$ is called the *orthogonal projector* onto V . The alternative notation $\hat{\cdot}_V$ highlights the target space V if its clarification is needed; the operator-type expression P_V provides yet another alternative. The meaning of $\tilde{\cdot}_V$ is in analogy, which generates the identity $\tilde{\cdot}_V = \hat{\cdot}_{V^\perp}$. The equality $\|x\|^2 = \|\hat{x}\|^2 + \|\tilde{x}\|^2$ proves that orthogonal projectors are *contractions*, that is, $\|\hat{x}_V\| \leq \|x\|$ for all $x \in W$. The *idempotence* of orthogonal projectors, that is, $P_V P_V = P_V$, follows from $\text{img } P_V = V$.

The subspace V^\perp is complementary to V in the sense that every $x \in W$ exhibits a unique decomposition $x = \hat{x}_V + \hat{x}_{V^\perp}$ with $\hat{x}_V \in V$ and $\hat{x}_{V^\perp} = \tilde{x}_V \in V^\perp$. In analogy, another subspace U of W is *complementary* to V (in W) if every $x \in W$ features a unique representation $x = \hat{x}_{V/U} + \hat{x}_{U/V}$ with $\hat{x}_{V/U} \in V$ and $\hat{x}_{U/V} \in U$. This relation between two subspaces holds if and only if their *sum* $U + V = \{u + v \mid u \in U, v \in V\}$ —a subspace of W —coincides with W and their intersection $U \cap V$ —another subspace of W —equals $\{0\}$. Then, the subspace U is called a *complement* of V (in W). Furthermore, this relation between V and U implies that the orthogonal complement U^\perp complements V^\perp , and that $\dim W = \dim V + \dim U$. In particular, all complements of V possess the same dimension, which is also referred to as the *codimension* $\text{codim } V$ of V (in W).

Every pair U, V of complementary subspaces comes with two associated projections: the *oblique projection* $\hat{x}_{V/U} \in V$ of $x \in W$ onto V along U and the oblique projection $\hat{x}_{U/V} \in U$ onto U along V . The latter are defined through the unique decomposition $x = \hat{x}_{V/U} + \hat{x}_{U/V}$ of x associated with the two complementary subspaces; panel (A) of figure 2.4 justifies the geometric naming. The associated residuals reflect the symmetry of this definition via $\tilde{x}_{V/U} = x - \hat{x}_{V/U} = \hat{x}_{U/V}$ and $\tilde{x}_{U/V} = x - \hat{x}_{U/V} = \hat{x}_{V/U}$, respectively.

oblique
projection

The corresponding *oblique projectors* $\hat{\cdot}_{V/U} = P_{V/U}$ and $\hat{\cdot}_{U/V} = P_{U/V}$ —either of which is denoted by $\hat{\cdot}_j$ if V and U as well as their role is clear from the context—are linear and idempotent. Moreover, one has the equalities $V = \text{img } P_{V/U} = \ker P_{U/V}$ and $U = \text{img } P_{U/V} = \ker P_{V/U}$, respectively, wherein the *image* $\text{img } F$ and *kernel* $\ker F$ of a linear map F on W are given by the subspaces $\{Fx | x \in W\}$ and $\{F = 0\}$, respectively. Finally, every linear and idempotent map $W \rightarrow W$ equals an oblique projector. In particular, an orthogonal projector P is an oblique projector with $\text{img } P = (\ker P)^\perp$.

oblique
projector

image
kernel

Figure 2.4 exemplifies the actions of the orthogonal projector $\hat{\cdot}_V$ onto a two dimensional subspace V of \mathbb{R}^3 as well as the projector $\hat{\cdot}_{V/U}$ onto the same space but along a complement $U \neq V^\perp$ in \mathbb{R}^3 . Its panel (A) shows the respective projections \hat{x}_V and $\hat{x}_{V/U}$ of some $x \notin V$. In particular, while \hat{x}_V and $\hat{x}_{V/U}$ are both elements of V , the orthogonal projection provides the closer substitute to x in terms of $\|\cdot\|$. This property rests on Pythagoras's theorem, that is, the relation $\|x\| = \sqrt{\langle x, x \rangle}$ is essential. Panel (B) concerns the linearity and (a consequence of) the idempotency of projectors. More specifically, two points x and y such that $x + y \in V$ are projected onto V along V^\perp as well as onto V along U . The sum of either pair of projections along the same complement equals the respective projection of the sum, which equals $x + y$ as $x + y \in V$.

For notational convenience, the composition $P_V X = [\hat{x}_1 \cdots \hat{x}_q]$ will subsequently be symbolized by $\hat{X}_V = \hat{X}$, wherein the latter is used only if the linear space V is clear from the context. Similarly, the two terms $\hat{X}_{V/U}$ and \hat{X}_j replace the composition $P_{V/U} X$.

2.4. Gramians

2.4.1. Inner product matrices

The $k \times q$ matrix $\langle Y, X \rangle$ with i, j -th entry $\langle y_i, x_j \rangle$ summarizes the geometric relations between two column sequences y_1, \dots, y_k and x_1, \dots, x_q in a Euclidean space W . If $k = q$, then the inner product $\langle \cdot, \cdot \rangle$ on $W^{\times k}$ equals the trace $\text{tr } \langle \cdot, \cdot \rangle$; in particular, $\langle \cdot, \cdot \rangle = \langle \cdot, \cdot \rangle$ for $k = q = 1$. If $W = \mathbb{R}^m$, then the equality $\langle A, \cdot \rangle = A^\top$ mimics $\langle a, \cdot \rangle = a^\top$. More generally, the similar appearance of $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle$ acknowledges the overlap of the respective feature sets of $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle$: firstly, every instance of $\langle \cdot, \cdot \rangle$ —defined on $W^{\times k} \times W^{\times q}$ —is bilinear; secondly, it exhibits symmetry to the extend that $\langle X, Y \rangle = \langle Y, X \rangle^\top$; thirdly, the *Gramian* $\langle Y, Y \rangle$ of $Y = [y_1 \cdots y_k]$ —or ‘of y_1, \dots, y_k ’—is *positive semidefinite*, that is $\langle a, \langle Y, Y \rangle a \rangle \geq 0$ for all $a \in \mathbb{R}^k$.

Gramian

positive
semidefinite

positive definite

The final property follows from $\langle a, \langle Y, Y \rangle a \rangle = \|Ya\|^2$, which also implies $\ker \langle Y, Y \rangle = \ker Y$. Thus, *positive definiteness* of $\langle Y, Y \rangle$ — $\langle a, \langle Y, Y \rangle a \rangle > 0$ for all $a \neq 0$ —is tantamount to $\ker Y = \{0\}$. Furthermore, $\ker \langle Y, \cdot \rangle = (\text{img } Y)^\perp$, and $\langle Y, Y \rangle \in \mathbb{S}^k$ guarantees $\langle \langle Y, Y \rangle a, b \rangle = \langle a, \langle Y, Y \rangle b \rangle$, $a, b \in \mathbb{R}^k$, which leads to $\text{img } \langle Y, \cdot \rangle = \text{img } \langle Y, Y \rangle = (\ker Y)^\perp$.

Example (e) in section 2.1 has its own customary nomenclature and notation regarding inner product matrices and more specifically Gramians. In this example,

- (e) the functions y_1, \dots, y_k and x_1, \dots, x_q are \mathbb{P} -square integrable random variables defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The *random vectors* $y = (y_1, \dots, y_k)$ and $x = (x_1, \dots, x_q)$, that is, $\mathcal{F}/\mathcal{R}^j$ -measurable functions $\Omega \ni \omega \mapsto (z_1(\omega), \dots, z_j(\omega)) \in \mathbb{R}^j$ with \mathcal{R}^j symbolizing the Borel σ -field of the norm topology on \mathbb{R}^j , allow the alternative representation $c \mapsto y^\top c$ and $c \mapsto x^\top c$ of the linear maps $Y = [y_1 \cdots y_k]$ and $X = [x_1 \cdots x_q]$, respectively. Moreover, the inner product $\langle \cdot, \cdot \rangle$ defined on $W = \text{img}[1 \ Y \ X]$ —after adjusting the representation as in appendix 2.a if needed—has the form of the \mathbb{P} -*expectation* $\langle x, y \rangle = \mathbb{E}xy = \int x(\omega)y(\omega)\mathbb{P}(\omega)$ of the pointwise product xy of $x, y \in W$. Herein, 1 denotes the constant function $\omega \mapsto 1$.

Consequently, the inner product matrix $\langle Y, X \rangle$ and the Gramian $\langle Y, Y \rangle$ equal $\mathbb{E}yx^\top$ and $\mathbb{E}yy^\top$, respectively. Therein, the expectations of the *random matrices* yx^\top and yy^\top , that is, $\mathcal{F}/\mathcal{R}^{d_1 \times d_2}$ -measurable maps $\Omega \rightarrow \mathbb{R}^{d_1 \times d_2}$ with $d_1, d_2 \in \mathbb{N}$ as well as $\mathcal{R}^{d_1 \times d_2}$ symbolizing the Borel σ -field of the respective norm topology, are defined entry-wise, that is, $\mathbb{E}y_i x_j = \langle y_i, x_j \rangle$ provides the i, j -th entry of $\mathbb{E}yx^\top$.

The projection of $z \in W$ onto the subspace $\text{span}\{1\}$ of W equals the expectation or *mean* $\langle 1, z \rangle 1 = \mathbb{E}z$ of z . The latter equality implicitly identifies a function of the type $\omega \mapsto c$ with $c \in \mathbb{R}$; this convention is applied throughout the text. The corresponding residual $z - \mathbb{E}z$ embodies the part of z that varies across different arguments ω ; its squared length $\mathbb{E}(z - \mathbb{E}z)^2 = \text{var}(z)$ is therefore called the *variance* of z . The Gramian of the residuals $y_1 - \mathbb{E}y_1, \dots, y_k - \mathbb{E}y_k$ provides the *variance matrix* $\text{var}(y)$ of the sequence y_1, \dots, y_k or (equivalently) the random vector y .

Gramians succinctly summarize the superiority of the composition \hat{X}_V of the linear map $X = [x_1 \cdots x_q]$ with the orthogonal projector onto a subspace V of W over the composition $\hat{X}_{V/U}$ of X with the oblique projector onto V along a complement $U \neq V^\perp$. More specifically, the residual maps $\tilde{X}_V = \tilde{X} = X - \hat{X}$ and $\tilde{X}_{V/U} = \tilde{X}_/ = X - \hat{X}_/$ satisfy

$$\langle a, \langle \tilde{X}_/, \tilde{X}_/ \rangle a \rangle = \|\tilde{X}a + \hat{X}a - \hat{X}_/a\|^2 = \langle a, \langle \tilde{X}, \tilde{X} \rangle a \rangle + \|(P_V - P_{V/U})Xa\|^2 \quad <2.5>$$

for all $a \in \mathbb{R}^q$. This equality directly follows from the linearity of projectors and the connection $\|x\| = \sqrt{\langle x, x \rangle}$. A general comment on the role of the latter is in order.

Complementary subspaces and oblique projectors are purely linear concepts in the sense that these notions are meaningful in the absence of a norm or an inner product. The inner product $\langle \cdot, \cdot \rangle$ or—by polarization—its induced norm determines the meaning of orthogonality. It thereby singles out a specific complement V^\perp of a subspace V of W and a single projector $P_V = P_{V/V^\perp}$ onto V as the orthogonal complement of V and orthogonal projector onto V , respectively. This projector enjoys the $\|\cdot\|$ -optimality in <2.5>. A different inner product $\langle \cdot, \cdot \rangle_*$ on W gives rise to another Euclidean space usually with a different complement $V^{\perp*}$ and projector $P_{V/V^{\perp*}}$ being the orthogonal ones. However, within the space $(W, \langle \cdot, \cdot \rangle_*)$ the projector $P_{V/V^{\perp*}}$ is (in general) merely oblique. This consideration of alternative inner products $\langle \cdot, \cdot \rangle_*$ provides an important source

of oblique projectors in this text and facilitates the associated computations as iterative schemes as in section 2.2.2 become applicable. Section 4.2.2 continues this line of argument. Sections 4.1.1 and 4.2.1 further investigate the final summand in <2.5>.

An inner product $\langle \cdot, \cdot \rangle$ on W endows every sequence y_1, \dots, y_k with a Gramian $\langle Y, Y \rangle$, that is, a positive semidefinite element of \mathbb{S}^k with kernel $\ker Y$. Lemma 2.3 provides a converse statement and an important source of further inner products $\langle \cdot, \cdot \rangle_*$ on $\text{span}\{y_1, \dots, y_k\}$. A proof of this assertion starts on page 40 in appendix 2.b.

Lemma 2.3. *If $G \in \mathbb{S}^k$ is positive semidefinite and $\ker G = \ker Y$, then there exists an inner product $\langle \cdot, \cdot \rangle_*$ on $\text{span}\{y_1, \dots, y_k\}$ with $\langle y_i, y_j \rangle_* = g_{i,j}$.*

The required kernel equality in lemma 2.3 is tantamount to $\text{img} \langle Y, \cdot \rangle = (\ker Y)^\perp = (\ker G)^\perp = \text{img} G$. In particular, every row $Y(\omega) = (y_1(\omega), \dots, y_k(\omega)) \in \mathbb{R}^k$ of Y exhibits a representation of the form $\langle Y, Q \rangle c$, wherein $Q \in \text{img} Y^{\times h}$ is $\langle \cdot, \cdot \rangle$ -unitary with columns q_1, \dots, q_h and $c = (q_1(\omega), \dots, q_h(\omega))$. Therefore, the kernel condition in lemma 2.3 requires the column space $\text{img} G$ of G to contain all rows of Y .

Lemma 2.2 yields a *Cholesky decomposition* of the Gramian $\langle Y, Y \rangle$, that is,

$$\underbrace{\begin{pmatrix} \langle y_1, y_1 \rangle & \cdots & \langle y_1, y_k \rangle \\ \vdots & \ddots & \vdots \\ \langle y_k, y_1 \rangle & \cdots & \langle y_k, y_k \rangle \end{pmatrix}}_{\text{Gramian } \langle Y, Y \rangle} = \begin{pmatrix} \boxed{r_{1,1}} \\ \vdots \quad \ddots \\ \boxed{r_{1,h} \cdots r_{h,h}} \\ \vdots \quad \ddots \quad \vdots \\ \boxed{r_{1,k} \cdots r_{h,k}} \end{pmatrix} \underbrace{\begin{pmatrix} \boxed{r_{1,1}} \cdots \boxed{r_{1,h}} \cdots \boxed{r_{1,k}} \\ \vdots \quad \ddots \quad \vdots \\ \boxed{r_{h,h}} \cdots \boxed{r_{h,k}} \end{pmatrix}}_{\text{Cholesky factor } R}.$$

Here, matrices R as in lemma 2.2 are referred to as *Cholesky factors*. Section 2.4.2 recovers such factors directly from $\langle Y, Y \rangle$ via an implicit Gram-Schmidt orthogonalization.

As a corollary to lemma 2.3, there exists a Euclidean space V with inner product $\langle \cdot, \cdot \rangle_*$ and a spanning sequence $y_1, \dots, y_k \in V$ such that $\langle y_i, y_j \rangle = g_{i,j}$ corresponding to any positive semidefinite matrix $G \in \mathbb{S}^k$. In fact, the columns g_1, \dots, g_k of G span a subspace of \mathbb{R}^k , and the kernel equality required by lemma 2.3 holds trivially. Consequently, the factorization process in section 2.4.2 finishes successfully whenever it is applied to a (nonzero) symmetric and positive semidefinite matrix.

2.4.2. Cholesky factorization

This sections considers a nontrivial sequence y_1, \dots, y_k with $Y = [y_1 \cdots y_k]$ and Gramian $\langle Y, Y \rangle$. In this case, the Cholesky factorization comprises k major steps and a final reduction. The j -th of these steps parallels the j -th major step of a Gram-Schmidt orthogonalization. It transforms the j -th row of the Gramian $\langle Y, Y \rangle$ to the j -th row of a preliminary upper triangular matrix \bar{R} using the first $j - 1$ rows of the latter matrix. To this end, it employs a sequence of triangularization steps—paralleling the orthogonalization steps in the Gram-Schmidt orthogonalization—and a scaling step. If needed, the reduction extracts a row echelon matrix R from \bar{R} ; otherwise $R = \bar{R}$.

If $y_j \in \text{span}\{y_1, \dots, y_{j-1}\}$, thus $\tilde{y}_j^{(j-1)} = 0$, then the row in <2.8> equals zero, that is, the j -th row of a preliminary upper triangular matrix \bar{R} produced during a Gram-Schmidt orthogonalization. Consequently, the procedure may advance to the next major step. An alternative given in <2.9> multiplies this zero row by zero to endow every major step with a scaling operation. The case $\tilde{y}_j^{(j-1)} \neq 0$ allows one of the two possible choices $\bar{r}_{j,j} = \pm \|\tilde{y}_j^{(j-1)}\|$. Subsequently, scaling <2.8> by $1/\bar{r}_{j,j}$ to obtain $\bar{r}_{j,\ell} = \langle \bar{q}_j, \tilde{y}_\ell^{(j-1)} \rangle = \langle \tilde{y}_j^{(j-1)}, \tilde{y}_\ell^{(j-1)} \rangle / \bar{r}_{j,j}$ is meaningful and concludes the construction of the j -th row of \bar{R} .

An complete description is given in display <2.9>. Therein, major steps are indexed by j , triangularization steps by i , and elements of the current row—the final row in <2.6> at the start of the j -th major step—by ℓ . This indexing parallels the above discussion. If the equality $k = h = \text{rk} Y$ holds, that is, $\ker Y = \ker \langle Y, Y \rangle = \{0\}$, then \bar{R} is upper triangular with nonzero diagonal elements, thus, of row echelon form. Otherwise, dropping the zero rows of \bar{R} yields a Cholesky factor R as in lemma 2.2.

$$\begin{array}{l|l}
1 & \bar{r}_{i,j}^{(0)} = \langle y_i, y_j \rangle, \quad i, j \leq k \\
2 & \text{for } j = 1, \dots, k \\
3 & \quad \text{for } i = 1, \dots, j-1 \\
4 & \quad \quad \text{for } \ell = 1, \dots, k \\
5 & \quad \quad \quad \bar{r}_{j,\ell}^{(i)} = \bar{r}_{j,\ell}^{(i-1)} - \bar{r}_{i,j} \bar{r}_{i,\ell} \\
6 & \quad \text{if } \bar{r}_{j,j}^{(j-1)} \neq 0 & \text{<2.9>} \\
7 & \quad \quad s_j = \pm (\bar{r}_{j,j}^{(j-1)})^{-1/2} \\
8 & \quad \text{else} \\
9 & \quad \quad s_j = 0 \\
10 & \quad \text{for } \ell = 1, \dots, k \\
11 & \quad \quad \bar{r}_{j,\ell} = \bar{r}_{j,\ell}^{(j-1)} s_j
\end{array}$$

The factorization <2.9> applies the same operations to the rows of $\langle Y, Y \rangle$ to obtain R as a Gram-Schmidt orthogonalization with corresponding sign choices executes on the columns of $Y = [y_1 \cdots y_k]$ to obtain Q . The first part of <2.7> and the scaling applied to <2.8> exemplify this observation. Viewing the factorization as a sequence of pre-multiplications with suitable matrix factors yields a concise statement. For example,

$$L_3 = \begin{pmatrix} 1 & & & \\ & 1 & & \\ & -\bar{r}_{1,3}s_3 & -\bar{r}_{2,3}s_3 & s_3 \end{pmatrix} = \begin{pmatrix} 1 & & \\ & 1 & \\ & & s_3 \end{pmatrix} \begin{pmatrix} 1 & & \\ & 1 & \\ & -\bar{r}_{2,3} & 1 \end{pmatrix} \begin{pmatrix} 1 & & \\ & 1 & \\ & -\bar{r}_{1,3} & 1 \end{pmatrix}$$

implements the third major step of <2.9> for $k = 3$. The first and second major step exhibit analogous factors L_1 and L_2 , respectively, leading to $L_3 L_2 L_1 \langle Y, Y \rangle = \bar{R}$. The general case $\langle Y, Y \rangle \in \mathbb{R}^{k \times k}$ uses k lower triangular factors L_1, \dots, L_k such that $L_k L_{k-1} \cdots L_1 \langle Y, Y \rangle = \bar{R}$. The reduction step amounts to a factor $L_{k+1} \in \mathbb{R}^{h \times k}$ with rows e_{i_1}, \dots, e_{i_h} , wherein e_ℓ denotes the ℓ -th standard basis element of \mathbb{R}^k and $i_1 < i_2 <$

$\dots < i_h$ represent the indexes corresponding to nonzero rows in \bar{R} . Using this notation

$$\begin{aligned} [y_1 \ \dots \ y_k] L_1^\top L_2^\top \dots L_j^\top &= [\bar{q}_1 \ \bar{q}_2 \ \dots \ \bar{q}_j \ y_{j+1} \ \dots \ y_k] , \quad j \leq k , \quad \text{and} \\ [\bar{q}_1 \ \dots \ \bar{q}_k] L_{k+1}^\top &= [q_1 \ \dots \ q_h] \end{aligned}$$

is a restatement of the above Gram-Schmidt orthogonalization.

2.5. Singular values

2.5.1. Singular value decomposition

If $Y \in W^{\times k}$ has rank $h > 0$, then compactness of $(\ker Y)^\perp \cap \{\|\cdot\| = 1\} \neq \emptyset$ and continuity of $\|Y\cdot\|$ guarantee the existence of a unit length h -th *right singular vector* $v_h \in (\ker Y)^\perp$ such that $\sigma_h = \inf_{v \in (\ker Y)^\perp \cap \{\|\cdot\|=1\}} \|Yv\| = \|Yv_h\| > 0$. Scaling the image Yv_h to unit length generates the corresponding *left singular vector* $u_h = \sigma_h^{-1} Yv_h$ of Y . Then, the kernel of $Y' = Y - \sigma_h u_h \langle v_h, \cdot \rangle$ equals the sum $\ker Y + \text{span}\{v_h\} = \{v + v' \mid v \in \ker Y, v' \in \text{span}\{v_h\}\}$ of $\ker Y$ and $\text{span}\{v_h\}$ —a $k - \text{rk } Y + 1$ dimensional subspace of \mathbb{R}^k . As a consequence, the rank $\text{rk } Y'$ of Y' equals $\text{rk } Y - 1$.

*right singular
vector*
*left singular
vector*

If $\text{rk } Y = h > 1$, then the reduced map Y' features $(\ker Y')^\perp \neq \{0\}$. The latter allows to construct a further left singular vector u_{h-1} of the form $\sigma_{h-1}^{-1} Y'v_{h-1}$, wherein $\sigma_{h-1} = \inf_{v \in (\ker Y')^\perp \cap \{\|\cdot\|=1\}} \|Y'v\| = \|Y'v_{h-1}\| \geq \sigma_h > 0$. If several choices for v_h —except the trivial $-v_h$ —exist, then v_{h-1} and thereby $u_{h-1} = Y'v_{h-1} = Yv_{h-1}$ depend on the choice of v_h . In any case, orthogonality of v_h and v_{h-1} is guaranteed by the construction. Moreover, minimality of σ_h requires $u_h \in (\text{img } Y')^\perp$. In fact, the residual from orthogonally projecting $\sigma_h u_h = Yv_h$ onto $\text{span}\{Y'v\}$ for any $v \in (\ker Y')^\perp$ equals

$$Yc = Y \begin{bmatrix} v & v_h \end{bmatrix} \begin{pmatrix} -\langle Y'v, Yv_h \rangle / \|Y'v\|^2 \\ 1 \end{pmatrix}$$

and—by Pythagoras's theorem—has length not exceeding σ_h . In addition, the length of the coordinate vector c in the preceding display exceeds one unless $\langle Y'v, Yv_h \rangle = \sigma_h \langle Y'v, u_h \rangle = 0$. Hence, the converse contradicts minimality of σ_h .

This argument exemplifies the geometric principle underlying the present construction: the linear combinations $c_1x + c_2y$ of two linearly independent elements x and y with $\|(c_1, c_2)\| = 1$ form an ellipse. Furthermore, if $x(= Yv)$ and $y(= Yv_h)$ are not orthogonal, then neither of the two equals the shortest principal semi-axis of that ellipse. Thus, there exists a unit length $c \in \mathbb{R}^2$ with $\|c_1x + c_2y\| < \min\{\|x\|, \|y\|\}$. Figure 2.5 illustrates this observation and the above construction. That is, if $\|x\|$ exceeds $\|y\|$, then scaling the residual \tilde{y} from orthogonally projecting y onto $\text{span}\{x\}$ yields a point on the ellipse, whose length provides a close substitute to the length σ_2 of the shorter semi-axis.

If $h > 2$, then the process continues in analogy. The h -th step yields a zero reduced map of the form $Y - \sum_{i \leq h} \sigma_i u_i \langle v_i, \cdot \rangle$ and thereby a representation $Y = \sum_{i \leq h} \sigma_i u_i \langle v_i, \cdot \rangle$, wherein u_1, \dots, u_h and v_1, \dots, v_h form orthonormal bases of $\text{img } Y$ and $(\ker Y)^\perp$, respectively. The associated scaling constants $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_h \geq 0$ are called




Figure 2.5

The figure illustrates that the set of images $\{[x \ y] c \mid \|c\| = 1\}$ of unit length $c \in \mathbb{R}^2$ under $[x \ y]$ forms an ellipse $\{\sum_{i \leq 2} \langle u_i, \cdot \rangle^2 / \sigma_i^2 = 1\}$ with principal semi-axes $\sigma_1 u_1$ and $\sigma_2 u_2$. Moreover, the residual \tilde{y} from orthogonally projecting y onto $\text{span}\{x\}$ lies outside of that ellipse.

singular values. Accordingly, the representation $Y = \sum_{i \leq h} \sigma_i u_i \langle v_i, \cdot \rangle$ is referred to as a *singular value decomposition* of Y . The latter generalizes figure 2.5: the set of images $\{Yc \mid \|c\| = 1\}$ of unit length vectors $c \in \mathbb{R}^k$ forms an ellipse $\{\sum_{i \leq k} \langle u_i, \cdot \rangle^2 / \sigma_i^2 = 1\}$.

singular value decomposition

2.5.2. Unitarily invariant norms

Singular vectors of an element $Y \in W^{\times k}$ may not be unique. In contrast, the singular values of Y are uniquely determined via $\sigma_j = \inf_{X \in W^{\times k}, \text{rk } X < j} \sup_{\|u\|=1} \|(Y - X)u\|$, $j \leq h = \text{rk } Y$. Firstly, the latter expression cannot exceed σ_j as $X = \sum_{i < j} \sigma_i u_i \langle v_i, \cdot \rangle \in W^{\times k}$ has rank $j - 1$ and $\sup_{\|u\|=1} \|(Y - X)u\| = \sigma_j$. Secondly, if $X \in W^{\times k}$ has rank less than j , then Xv_1, \dots, Xv_j , wherein v_1, \dots, v_j provide (an initial stretch of a sequence of) right singular vectors of Y —as constructed in section 2.5.1, are j elements of the $j - 1$ dimensional space $\text{img } X$ and therefore exhibit linear dependence. Thus, there exists a unit length element $v \in \text{span}\{v_1, \dots, v_j\} \cap \ker X$ and consequently $\|(Y - X)v\| \geq \sigma_j$.

The notation $\sigma_1(Y), \dots, \sigma_h(Y)$ highlights the uniqueness implied by the approximation error characterization. In addition, the latter suggest setting $\sigma_{h+1}(Y) = \dots = \sigma_k(Y) = 0$ and implies the invariance of singular values to composition—from left or right—with suitable unitary maps. The most important elements of the singular value sequence $\sigma_1, \dots, \sigma_k$ are the maximal and least nonzero singular value and thereby deserve their own special symbols $\sigma_1 = \sigma_{\max}$ and $\sigma_h = \sigma_{\min, \neq 0}$, respectively.

Several norms on $W^{\times k}$ (and thus $\mathbb{R}^{m \times k}$) measure the length of an element $Y \in W^{\times k}$ in terms of its singular values. Such norms are invariant to composition with unitary maps and are therefore called *unitarily invariant*. The relevant examples include

unitarily invariant

(d₁) the *Frobenius norm* of $Y \in W^{\times k}$ defined by $\|Y\| = \langle Y, Y \rangle^{1/2}$ —see (d) in section 2.1.3—equals the square-root of $\sum_{i \leq k} \sigma_i^2(Y)$;

Frobenius norm

(d₂) the *operator norm* $\|Y\|_{\text{op}} = \sup_{\|c\|=1} \|Yc\|$ coincides with $\sigma_{\max}(Y)$; and,

operator norm

(d₃) lastly, the sum $\sum_{i \leq k} \sigma_i(Y)$ supplies the *nuclear norm* $\|Y\|_{\text{nuc}}$ of Y .

nuclear norm

The triangle inequality for the nuclear norm $\|\cdot\|_{\text{nuc}}$ follows from its characterization as the *dual norm* of $\|\cdot\|_{\text{op}}$. In general, the dual norm of an arbitrary norm $\|\cdot\|'$ on a Euclidean space is given by $\sup_{\|x\|'=1} \langle x, \cdot \rangle$. Section 2.1.3 shows that $\|x\| = \langle x, x \rangle^{1/2}$ equals its own dual, which implies the *Cauchy-Schwarz inequality* $|\langle x, y \rangle| \leq \|x\| \|y\|$. More generally, one has $|\langle x, y \rangle| \leq (\sup_{\|z\|'=1} \langle z, y \rangle) \|x\|'$. This inequality suggests a symmetric relation. In fact, the norm $\|\cdot\|'$ equals the dual norm of $\sup_{\|x\|'=1} \langle x, \cdot \rangle$.

dual norm

Cauchy-Schwarz inequality

The nuclear norm $\|\cdot\|_{\text{nuc}}$ in (d₃) satisfies $\langle Y, X \rangle = \sum_{i \leq h} \sigma_i \langle u_i, X v_i \rangle \leq \sum_{i \leq h} \sigma_i$ whenever $\|X\|_{\text{op}} = 1$ and $Y = \sum_{i \leq h} \sigma_i u_i \langle v_i, \cdot \rangle$ represents a singular value decomposition of Y . Equality holds for $X = \sum_{i \leq h} u_i \langle v_i, \cdot \rangle$. Thus, $\|\cdot\|_{\text{nuc}}$ and $\|\cdot\|_{\text{op}}$ are duals. The resulting inequality $|\langle Y, X \rangle| \leq \|Y\|_{\text{op}} \|X\|_{\text{nuc}}$ can be refined using the representation

$$W^{\times k} \ni Y = \sum_{i \leq \text{rk } Y} \sigma_i u_i \langle v_i, \cdot \rangle = \sum_{i \leq \text{rk } Y} (\sigma_i - \sigma_{i+1}) \sum_{j \leq i} u_j \langle v_j, \cdot \rangle = \sum_{i \leq \text{rk } Y} c_i Y_i,$$

wherein $\sigma_{\text{rk } Y + 1} = 0$, $c_i = \sigma_i - \sigma_{i+1} \geq 0$, and all singular values of $Y_i = \sum_{j \leq i} u_j \langle v_j, \cdot \rangle$ are unity. If $\sum_{i \leq \text{rk } X} \bar{c}_i X_i$ is in analogy with respect to $X \in W^{\times k}$, then

$$\begin{aligned} \langle X, Y \rangle &\leq \sum_{i \leq \text{rk } Y} \sum_{j \leq \text{rk } X} c_i \bar{c}_j |\langle Y_i, X_j \rangle| \leq \sum_{i \leq \text{rk } Y} \sum_{j \leq \text{rk } X} c_i \bar{c}_j \min\{i, j\} \\ &= \sum_{i \leq \text{rk } Y} \sum_{j \leq \text{rk } X} c_i \bar{c}_j \text{tr } B_i B_j = \sum_{i \leq k} \sigma_i(Y) \sigma_i(X), \quad \langle 2.10 \rangle \end{aligned}$$

wherein $B_i = \sum_{j \leq i} B_{j,j}$ with $B_{j,j}$ being the j, j -th element of the standard basis of $\mathbb{R}^{k \times k}$, and the second step follows from $\|\cdot\|_{\text{op}} / \|\cdot\|_{\text{nuc}}$ -duality. The comparison between the leftmost and rightmost term in $\langle 2.10 \rangle$ is known as the *von Neumann trace inequality*.

von Neumann trace inequality

Finally, the representations in (d₁), (d₂), and (d₃) imply the inequalities

$$\|Y\|_{\text{op}}^2 \leq \|Y\|^2 \leq \sum_{i, j \leq \text{rk } Y} \sigma_i \sigma_j = \|Y\|_{\text{nuc}}^2 = \left\langle \begin{pmatrix} \sigma_1 \\ \vdots \\ \sigma_h \end{pmatrix}, \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \right\rangle^2 \leq \text{rk } Y \|Y\|^2 \leq (\text{rk } Y)^2 \|Y\|_{\text{op}}^2$$

for $Y \in W^{\times k}$. The rank-related factors in the preceding display are for a given $Y \in W^{\times k}$. The respective upper subspace compatibility constants—as defined in section 2.1.2—are given by $(\min\{\dim W, k\})^{1/2}$ and its square—the maximum possible rank of $Y \in W^{\times k}$.

2.5.3. Singular space pairs

Left and right singular vectors are—at most—unique up to a sign choice. In particular, $\sigma_j u_j \langle v_j, \cdot \rangle$ remains unchanged if both u_j and v_j are multiplied by -1 . At the other extreme, a unitary map from \mathbb{R}^k to a Euclidean space W has all its singular values equal to unity, and every orthonormal basis of \mathbb{R}^k may serve as its right singular vectors.

This section considers $Y \in W^{\times k}$ with $\text{rk } Y = h > 0$ to address the general case. If the least nonzero singular value σ_h of Y is attained at two unit length elements v_h and $v'_h \notin \text{span}\{v_h\}$, then the residual \tilde{v}'_h from orthogonally projecting v'_h onto $\text{span}\{v_h\}$ —a linear

combination of v_h, v'_h —is a nonzero element of $(\ker Y')^\perp$, wherein $Y' = Y - \sigma_h u_h \langle v_h, \cdot \rangle$ and u_h symbolizes the left singular vector corresponding to v_h . Hence, this residual provides a valid choice for v_{h-1} in form of $\tilde{v}'_h / \|\tilde{v}'_h\|$. Consequently, $\langle Y \tilde{v}'_h, Y v_h \rangle = 0$ implies $\sigma_h^2 = \|Y v'_h\|^2 = \|Y \tilde{v}'_h\|^2 + \|Y \hat{v}'_h\|^2 = \|Y \tilde{v}'_h\|^2 + \sigma_h^2 \langle v_h, v'_h \rangle^2$, wherein $\hat{v}'_h = v_h - \tilde{v}'_h = v_h \langle v_h, v'_h \rangle$. Hence, one has $\|Y \tilde{v}'_h\|^2 = \sigma_h^2 (1 - \langle v'_h, v_h \rangle^2) = \sigma_h^2 \|\tilde{v}'_h\|^2$ and thereby $\sigma_h = \sigma_{h-1}$. Elements $c_1 v_h + c_2 v_{h-1}$ of $\text{span}\{v_h, v_{h-1}\}$, wherein $v_{h-1} = \tilde{v}'_h / \|\tilde{v}'_h\|$, satisfy $\|Y(c_1 v_h + c_2 v_{h-1})\|^2 = \sigma_h^2 (c_1^2 + c_2^2)$, thus, are elements of $\{\|Y \cdot\| = \sigma_h \|\cdot\|\}$.

Either $\text{span}\{v_h, v_{h-1}\}$ equals the latter set or that set contains an element $v'_{h-1} \notin \text{span}\{v_h, v_{h-1}\}$. In the latter case, the residual \tilde{v}'_{h-1} from orthogonally projecting v'_{h-1} onto the span of v_h and v_{h-1} is nonzero and leads to a candidate $\tilde{v}'_{h-1} / \|\tilde{v}'_{h-1}\|$ for v_{h-2} . Then it follows that $\langle Y \tilde{v}'_{h-1}, Y v \rangle = 0$ for all $v \in \text{span}\{v_h, v_{h-1}\}$ and thereby $\sigma_{h-2} = \sigma_h$. Hence, $\text{span}\{v_h, v_{h-1}, v_{h-2}\} \subset \{\|Y \cdot\| = \sigma_h \|\cdot\|\}$. If the latter sets differ, then a further iteration is possible. The recursion stops after $m_h \leq h$ steps. It identifies $\{\|Y \cdot\| = \sigma_h \|\cdot\|\}$ as a subspace and generates a corresponding orthonormal basis v_{h-m_h+1}, \dots, v_h .

This recursive argument is applicable to the reduced map $Y - \sigma_h \sum_{j=0}^{m_h-1} u_{h-j} \langle v_{h-j}, \cdot \rangle$, wherein u_j again represents the left singular vector corresponding to v_j , and so forth. Eventually, these arguments construct subspaces $V_j = \{\|Y \cdot\| = \bar{\sigma}_j \|\cdot\|\}$, $j \leq s$, corresponding to the *distinct (nonzero) singular values* $\bar{\sigma}_1 > \dots > \bar{\sigma}_s (= \sigma_h) > 0$ of Y .

The characterization $V_j = \{\|Y \cdot\| = \bar{\sigma}_j \|\cdot\|\}$ guarantees that these subspaces are uniquely determined by Y . Elements $v \in V_i, v' \in V_j, i \neq j$, are orthogonal, and the *sum* $\sum_{j \leq s} V_j = \{\sum_{j \leq s} v_j \mid v_j \in V_j, j \leq s\}$ of these subspaces equals $(\ker Y)^\perp$. The dimension $\dim V_j$ supplies the *multiplicity* m_j of $\bar{\sigma}_j$, and therefore one has $h = \text{rk } Y = \sum_{j \leq s} m_j$. Moreover, the images U_j of Y restricted to V_j satisfy $U_i \subset U_j^\perp, i \neq j$, and their sum equals $\text{img } Y$. The pairs (V_j, U_j) are called the *singular subspace pairs* for Y .

The valid selections of right and left singular values of Y consist of arbitrarily chosen orthonormal bases for $V_j, j \leq s$, and the corresponding scaled images, respectively.

2.5.4. Singular vectors of symmetric matrices

The symmetry of $A \in \mathbb{S}^m$ is reflected by its singular vectors.

Lemma 2.4. *If $A \in \mathbb{S}^m$ is of rank $h > 0$, then there exists a choice of right singular vectors v_1, \dots, v_h with corresponding left singular vectors u_1, \dots, u_h such that $u_i = \pm v_i$.*

A proof of lemma 2.4 starts on page 40 in appendix 2.b. This result guarantees that the approximation error characterization, the duality relations, and inequalities in section 2.5.2 are equally valid—and follow by the same arguments—if the symmetric matrices are considered in isolation. In particular, any two symmetric matrices $A, B \in \mathbb{S}^m$ satisfy $\langle A, B \rangle \leq \sum_{i \leq m} \sigma_i(A) \sigma_i(B) \leq \|A\|_{\text{op}} \|B\|_{\text{nuc}}$, wherein equality is possible.

A representation $A = \sum_{i \leq h} \pm \sigma_i v_i \langle v_i, \cdot \rangle$ as in lemma 2.4 is called a *spectral decomposition* of A . This expression reveals that the singular subspaces V_j are sums of the two subspaces $V_j^+ = \ker(\bar{\sigma}_j \text{id} - A)$ and $V_j^- = \ker(\bar{\sigma}_j \text{id} + A)$ with $V_j^+ \subset (V_j^-)^\perp$.

Thus, A is of the form $\sum_{i \leq s} \bar{\sigma}_i (P_{V_j^+} - P_{V_j^-})$, wherein $\bar{\sigma}_1, \dots, \bar{\sigma}_s, P_V$, and id represent the distinct nonzero singular values of A , the orthogonal projector onto a subspace $V \subset \mathbb{R}^m$, and the identity map on \mathbb{R}^m , respectively. Such (projector-based)

*distinct
(nonzero)
singular value*

sum

multiplicity

*singular
subspace*

*spectral
decomposition*

representations are unique and show that the linear space of symmetric matrices \mathbb{S}^m provides the (\subset -)smallest subspace of $\mathbb{R}^{m \times m}$ containing all orthogonal projectors onto subspaces of \mathbb{R}^m . By definition, at most one of the two subspaces V_j^+ and V_j^- , $j \leq s$, may equal $\{0\}$. In particular, if A is positive semidefinite, then $V_j^- = \{0\}$ for all $j \leq s$ and $\text{tr } A = \sum_{i \leq \text{rk } A} \sigma_i \text{tr}(u_i u_i^\top) = \|A\|_{\text{nuc}}$ as $\text{tr}(u_i u_i^\top) = \sum_{j \leq m} u_{j,i}^2 = \|u_i\|^2 = 1$.

The above terminology allows a complete characterization of $Y \in W^{\times k}$ such that $\langle Y, X \rangle = \|X\|_{\text{op}} \|Y\|_{\text{nuc}}$ for a given $X \in W^{\times k}$. A proof starts on page 40 in appendix 2.b.

Lemma 2.5. *Let X be a nonzero linear map from \mathbb{R}^k to a Euclidean space W and $X = \sum_{i \leq \text{rk } X} \sigma_i u_i \langle v_i, \cdot \rangle$ be any singular value decomposition of X , then the equality $\langle Y, X \rangle = \|X\|_{\text{op}}$ holds for a unit $\|\cdot\|_{\text{nuc}}$ -length $Y \in W^{\times k}$ if and only if there exists a positive semidefinite $S \in \mathbb{S}^{m_1}$ with $\text{tr } S = 1$ and*

$$Y = [u_1 \cdots u_{m_1}] S \langle [v_1 \cdots v_{m_1}], \cdot \rangle,$$

wherein m_1 denotes the multiplicity of the largest (distinct) singular value $\bar{\sigma}_1$ of X .

If $A \in \mathbb{S}^m$ and the columns $u_1^+, \dots, u_{m_1'}^+$ of U_1^+ and $u_1^-, \dots, u_{m_1''}^-$ of U_1^- form orthonormal bases of $V_1^\pm = \ker(A \mp \bar{\sigma}_1(A) \text{id})$, respectively, then $m_1' + m_1'' = m_1$, the multiplicity of $\bar{\sigma}_1(A)$, and (u_i^+, u_i^+) as well as $(u_j^-, -u_j^-)$ provide suitable singular vector pairs. Corollary 2.6—proved on page 41 in appendix 2.b—states the resulting representation.

Corollary 2.6. *For every symmetric $B \in \{\|\cdot\|_{\text{nuc}} = 1\}$ with $\langle A, B \rangle = \|A\|_{\text{op}}$, there exist positive semidefinite $S^+ \in \mathbb{S}^{m_1'}$, $S^- \in \mathbb{S}^{m_1''}$ such that*

$$B = U_1^+ S^+ \langle U_1^+, \cdot \rangle - U_1^- S^- \langle U_1^-, \cdot \rangle$$

and $\text{tr } S^+ + \text{tr } S^- = 1$. Moreover, there exists a selection of bases such that S^+ and S^- are diagonal matrices, that is, all non-diagonal entries of these matrices are zero.

*diagonal
matrices*

Comments and references

Section 2.1 Halmos (1974) covers most of the topics of section 2.1 in-depth. Note that the usage of the term *coordinates* in this text is nonstandard. A more detailed treatment of matrix norms may be found in Golub and Van Loan (2013, sec. 2.3). Vershynin (2012, sec. 5.2.2) defines ε -nets and covering numbers. Lemma 2.1 equals his lemma 5.2. At first sight, the covering of the $\frac{\varepsilon}{2}$ -balls by $\{\|\cdot\| \leq 1 + \varepsilon/2\}$ in the proof given in appendix 2.b may seem overly generous. However, more refined replacements for this set do not lead to more informative upper bounds on the covering number $N(\{\|\cdot\| = 1\}, \|\cdot\|, \varepsilon)$.

The presentation of Euclidean geometry in Strang (2005, chapter 3) exhibits a similar style as section 2.1.3. Specifically, his figures 3.1a, 3.6, and 3.7 closely resemble figure 2.1 [Panel (A)], 2.3 [Panel (A)], and figure 2.1 [Panel (B)], respectively. Kailath et al. (2000, appendix 4.A) supply the examples of section 2.1.1 except for the symmetric matrices. Borwein and Lewis (2010, sec. 1.2) fill this gap. The matrix-like notation for linear maps $\mathbb{R}^k \rightarrow V$ is borrowed from Morf and Kailath (1975, section IV).

Section 2.2 The motivation of the concept of a unitary map is taken from Halmos (1974, §73). The term *representation* is borrowed from Parzen (1961, def. 4A).

If $y_1, \dots, y_k \in \mathbb{R}^m$, then $Y = QR$ in lemma 2.2 is known as *QR decomposition* (Golub and Van Loan, 2013, sec. 5.2). Björck (1996, sec. 2.4.2) treats several variants of Gram-Schmidt orthogonalization. The formulation in <2.3> replicates his algorithm 2.4.4 and is therein termed *column oriented modified Gram-Schmidt process*. The present treatment of the case with linear dependence is nonstandard; the usual treatment (Björck, 1996, rem. 2.4.5) involves a rearrangement of the input sequence—called *pivoting*.

Section 2.3 Halmos (1974, §73, §41) provides a similar treatment of orthogonal and oblique projectors. Furthermore, his sections §18, §19 and §20 prove the assertions about complements. Figure 2 of Wedin (1983, p. 266) illustrates the corresponding decomposition into two projections in similar fashion as in panel (A) of figure 2.4. The notation for orthogonal and oblique projectors is close to Galántai (2008, sec. 2).

Section 2.4 The linear map $\langle Y, \bullet \rangle$ (on W) amounts to the *adjoint* of Y (Halmos, 1974, §44). Gramians are defined in Kailath et al. (2000, appendix 4.A, (4.A.2)). Doz et al. (2011, sec. 3) employ Gramian substitutes to derive (associated) oblique projections.

The corollary of lemma 2.3 on the existence of a Euclidean space supporting $\langle \bullet, \bullet \rangle_*$ is often proved by showing—by induction—that a Cholesky factorization with positive semidefinite input A finishes successfully (Trefethen and Bau, 1997, par. before thm. 23.1). The geometric approach taken here is an elementary version of property (4) of Aronszajn (1950, part I, sec. 2). The proof of the chosen inner product being well defined parallels Schölkopf and Smola (2002, sec 2.2, pp. 32–33).

Golub and Van Loan (2013, sec. 4.2) treat Cholesky factorization for linearly independent y_1, \dots, y_k , however, use the term *Cholesky factor* for the lower triangular R^T . Their equivalent to <2.9> in Golub and Van Loan (2013, algorithm 4.2.1)—called *gaxpy Cholesky*—differs accordingly. Therein, linear independence leads to a (unique) Cholesky factor with positive diagonal entries; this is the common usage of this term.

Kailath et al. (2000, prob. 12.3) mention the version of \bar{R} —the triangular matrix generated by <2.9>—with nonnegative diagonal entries but under a different name. They also discuss the Gram-Schmidt/Cholesky correspondence in their section 4.4. Eubank (2006, sec. 1.2.3) stresses the identical output of two very similar algorithms.

The representation of a Cholesky factorization as pre-multiplications with lower triangular matrices is from Trefethen and Bau (1997, lecture 23, pp.173–174, algorithm 23.1).

Section 2.5 Anderson (1958, sec. 11.2, thm. 11.2.1) constructs (left) singular vectors under the alternative label *principal components* but in opposite order. His construction treats singular values and vectors via eigen-theory; this approach is a widespread alternative to the topics of this sections (Stewart and Sun, 1990, I.3, I.4).

Trefethen and Bau (1997, lecture 4) motivate the notion of a (reduced) singular value decomposition by geometric arguments; in particular, figure 2.5 resembles their figure 4.1. Golub and Van Loan (2013, proof of thm. 2.4.1, thm. 2.4.8) justify the orthog-

onality of u_{h-1} and u_h in essentially the same way but in reverse direction. They refer to the approximation error characterization of singular values as the *Eckhart-Young theorem* and provide the corresponding argument.

Recht et al. (2010, sec. 2, prop. 2.1) represent the norms (d_1) , (d_2) , and (d_3) using singular values. Their section 2 also defines the concept of a dual norm and derives the nuclear norm/operator norm duality. The nuclear norm $\|\cdot\|_{\text{nuc}}$ is also known as the *trace norm* or *Schatten-1-norm*. The *Ky-Fan-h-norm* equals the sum of the h largest singular values; thus, both $\|\cdot\|_{\text{op}}$ and $\|\cdot\|_{\text{nuc}}$ are of this type. Alternative names for $\|\cdot\|$ in (d_1) include *Schatten-2-norm* and *Hilbert-Schmidt norm*. Unitarily invariant norms are the topic of Stewart and Sun (1990, II.3). The proof of von Neumann trace inequality is due to Grigorieff (1991). Stewart (1973, def. 6.1) contains a comparable, but less restrictive notion of singular space pairs. Halmos (1974, §79, thm. 1) presents the projector-based spectral decomposition. Lemma 2.5 amounts to theorem 4.3 of Ziętak (1988).

Appendixes Pollard (2002, ch. 2) contains the relevant results on L^2 -spaces. Therein, part (iii) of Pollard (2002, ch. 2, sec. 6, lem. 26) caters the techniques used to quantify the influence of the choice of basis element representatives.

The present approach yields a *reproducing kernel Hilbert space* (Schölkopf and Smola, 2002, def. 2.9). Aronszajn (1950, sec. 3) shows that for a given choice of representatives q_1, \dots, q_k the associated *reproducing kernel* is of the form $K(\omega, \omega') = \sum_{i \leq k} q_i(\omega)q_i(\omega')$. Its *reproducing property* yields the *evaluation functional* $f_\omega(y) = \langle \sum_{i \leq k} q_i(\omega)q_i, y \rangle$, which is implicitly used in section 2.4.1.

- Anderson, T. W. (1958). *An introduction to multivariate statistical analysis*. Wiley series in probability and mathematical statistics: Probability and mathematical statistics. New York: Wiley.
- Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American Mathematical Society* 68(3), 337–404.
- Björck, Å. (1996). *Numerical methods for least squares problems*. SIAM.
- Borwein, J. M. and A. S. Lewis (2010). *Convex analysis and nonlinear optimization: theory and examples* (2 ed.), Volume 3. Springer.
- Doz, C., D. Giannone, and L. Reichlin (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics* 164(1), 188–205.
- Eubank, R. L. (2006). *A Kalman filter primer*, Volume 186 of *Statistics*. Chapman & Hall/CRC.
- Galántai, A. (2008). Subspaces, angles and pairs of orthogonal projections. *Linear and Multilinear Algebra* 56(3), 227–260.
- Golub, G. H. and C. F. Van Loan (2013). *Matrix computations* (4 ed.). Johns Hopkins studies in the mathematical sciences. Baltimore, Md.: Johns Hopkins Univ. Press.
- Grigorieff, R. D. (1991). A note on von Neumanns trace inequality. *Math. Nachr* 151, 327–328.
- Halmos, P. R. (1974). *Finite-dimensional vector spaces* (Repr. of the 2. ed. ed.). New York: Springer.
- Kailath, T., A. H. Sayed, and B. Hassibi (2000). *Linear estimation*, Volume 1. Prentice Hall New Jersey.
- Morf, M. and T. Kailath (1975). Square-root algorithms for least-squares estimation. *Automatic Control, IEEE Transactions on* 20(4), 487–497.
- Parzen, E. (1961). An approach to Time Series Analysis. *The Annals of Mathematical Statistics* 32(4), pp. 951–989.
- Pollard, D. (2002). *A user’s guide to measure theoretic probability*, Volume 8 of *Cambridge series in statistical and probabilistic mathematics*. Cambridge: Cambridge University Press.

- Recht, B., M. Fazel, and P. A. Parrilo (2010). Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review* 52(3), 471–501.
- Schölkopf, B. and A. J. Smola (2002). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. Adaptive computation and machine learning. Cambridge, Mass.: MIT Press.
- Stewart, G. W. (1973). Error and perturbation bounds for subspaces associated with certain eigenvalue problems. *SIAM review* 15(4), 727–764.
- Stewart, G. W. and J. Sun (1990). *Matrix perturbation theory*. Computer science and scientific computing. Boston: Acad. Press.
- Strang, G. (2005). *Linear algebra and its applications* (4 ed.). Cengage Learning.
- Trefethen, L. N. and D. Bau (1997). *Numerical linear algebra*. Philadelphia: SIAM, Soc. for Industrial and Applied Math.
- Vershynin, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In Y. C. Eldar and G. Kutyniok (Eds.), *Compressed Sensing*, pp. 210–268. Cambridge University Press.
- Wedin, P. Å. (1983). On angles between subspaces of a finite dimensional inner product space. In *Matrix Pencils*, pp. 263–285. Springer.
- Zi̧tak, K. (1988). On the characterization of the extremal points of the unit sphere of matrices. *Linear Algebra and its Applications* 106, 57–75.

Appendix

2.a. Square integrable functions

The basic entities of the present section are μ -square integrable functions defined on a common measure space $(\Omega, \mathcal{F}, \mu)$ with finite measure μ . The geometry of the linear space spanned by such functions derives from an inner product, whose definition relies on μ -square integrability of the respective functions not finiteness of μ . In fact, if functions x and y are μ -square integrable, then $\int |x(\omega)y(\omega)|\mathbb{P}(d\omega) \leq \int \frac{1}{2}(x^2(\omega) + y^2(\omega))\mathbb{P}(d\omega)$ as $0 \leq (x(\omega) \pm y(\omega))^2 = 2[(x^2(\omega) + y^2(\omega))/2 \pm x(\omega)y(\omega)]$ for every $\omega \in \Omega$, which ensures that $\langle \cdot, \cdot \rangle$ is well defined as a real-valued map. However, finiteness of μ guarantees that the set $\mathcal{L}^2 = \mathcal{L}^2(\Omega, \mathcal{F}, \mu)$ of μ -square integrable functions on (Ω, \mathcal{F}) includes the $\mathcal{F}/\mathcal{R}^1$ -measurable functions— \mathcal{R}^1 being the Borel σ -field of the $|\cdot|$ -topology—with finite range. In fact, if μ is finite, then the indicators of elements of \mathcal{F} and thus linear combinations of the latter are μ -square integrable. In particular, if y is $\mathcal{F}/\mathcal{R}^1$ -measurable, then its composition with the *sign function*—given by $\text{sign}(x) = -1, 0, 1$ for, respectively, negative, zero, or positive $x \in \mathbb{R}$ —is μ -square integrable. Thus, if μ is finite, then $\int |y(\omega)|\mu(d\omega) = \langle y, \text{sign}(y) \rangle$ shows that μ -square integrability of y implies its μ -integrability.

sign function

If \mathcal{F} contains sets of μ -measure zero, then their indicators witness the existence of nonzero elements f of \mathcal{L}^2 with $\|y\| = (\int y^2(\omega)\mu(d\omega))^{1/2} = 0$. Here, these are called representatives of zero, and their existence degrades $\|\cdot\|$ to a seminorm. The conventional way out is to take $y = x$ if $y - x$ is a representative of zero by partitioning \mathcal{L}^2 into equivalence classes $[[y]] = \{\|\cdot - y\| = 0\}$. As the set $[[0]]$ of representatives of zero forms a subspace under pointwise linear operations, the set of equivalence classes $\{[[y]] \mid y \in \mathcal{L}^2\}$ forms a linear space L^2 —a so-called quotient space—with the linear operations $a[[y]] = [[ay]]$ and $[[y]] + [[x]] = [[y + x]]$. Adapting $\langle \cdot, \cdot \rangle$ to operate on L^2 via $\langle [[y]], [[x]] \rangle = \langle y, x \rangle$ yields a well defined inner product on that space. Thus, any finite dimensional subspace of L^2 provides a Euclidean space. The accustomed notation double uses y instead of the pair y and $[[y]]$, but supplements relations such = and \leq with “almost everywhere” qualifiers to warn against an unwarranted pointwise interpretation.

An alternative strategy—used herein—amounts to choosing a subspace of \mathcal{L}^2 containing a single point from each element—an equivalence class—of the L^2 -subspace under consideration. In the present setting this approach dispenses with the tiresome “almost everywhere” qualification and more importantly comes with the advantage that pointwise evaluation remains well defined. This construction builds on an initial choice of an $\mathcal{F}/\mathcal{R}^1$ -measurable element q_i from each equivalence class $[[q_i]]$ of an orthonormal basis $[[q_1]], \dots, [[q_m]]$ of the relevant L^2 -subspace. Then, a suitable element of $[[y]]$ follows from combining q_1, \dots, q_m with the coordinates $\langle [[y]], [[q_1]] \rangle, \dots, \langle [[y]], [[q_m]] \rangle$ of $[[y]]$ with respect to $[[q_1]], \dots, [[q_m]]$. In fact, the resulting \mathcal{L}^2 -subspace $W = \text{span}\{q_1, \dots, q_m\}$ has pointwise zero as its sole representative of zero. More generally, $y, x \in [[y]] \cap W$ implies $y - x \in [[0]] \cap W$, thus, the pointwise equality $y = x$. Hence, $\langle \cdot, \cdot \rangle$ and therefore $\|\cdot\|$ are an inner product and a norm on W , respectively.

The choice of the basis representatives does not affect the geometry—induced by $\langle \cdot, \cdot \rangle$ —of the resulting \mathcal{L}^2 -(sub)space W . Moreover, an alternative choice of a $\mathcal{F}/\mathcal{R}^1$ -measurable $q'_i \in [[q_i]]$ differs from q_i solely on the μ -measure zero set $N_i = \{q_i \neq q'_i\} \in \mathcal{F}$. Therefore, all elements of the two linear spaces $\text{span}\{q_1, \dots, q_k\}$ and $\text{span}\{q'_1, \dots, q'_k\}$ differ from their respective counterparts at most on the μ -measure zero set $\cup_{i \leq k} N_i$. In particular, the *image measure* $\mu \circ y^{-1}$, that is, $\mathcal{R}^1 \ni B \mapsto (\mu \circ y^{-1})B = \mu\{y \in B\}$, of a given linear combination $y = \sum_{i \leq k} c_i q_i$ does not depend on the choice of basis element representatives. In fact,

image measure

$$\mu\left\{\sum_{i \leq k} c_i q_i \in B\right\} = \mu\left(\left\{\sum_{i \leq k} c_i q_i \in B\right\} \cap \bigcap_{i \leq k} \{q_i = q'_i\}\right) = \mu\left\{\sum_{i \leq k} c_i q'_i \in B\right\}.$$

An analogous argument also applies to the image measure of $\omega \mapsto (y_1(\omega), \dots, y_\ell(\omega))$ —defined on the Borel σ -field of the norm topology on \mathbb{R}^k , wherein $y_i = \sum_{j \leq k} c_{j,i} q_j$ with $c_{j,i} \in \mathbb{R}$, $\ell \in \mathbb{N}$.

2.b. Proofs

Proof of lemma 2.1. A finite subset $\{x_1, \dots, x_q\}$ of $\{\|\cdot\| = 1\}$ is ε -separated if $d(x_i, x_j) = \|x_i - x_j\| > \varepsilon$ for all $i, j \leq q$ with $i \neq j$. The construction of a \subset -maximal element of the set \mathcal{S} of ε -separated subsets of $\{\|\cdot\| = 1\}$ succeeds by starting at an arbitrary unit length x_1 and recursively adding unit length x_n with $d(x_n, x_j) > \varepsilon$, $j < n$. Compactness of $\{\|\cdot\| = 1\}$ guarantees that the construction terminates after $q(\in \mathbb{N})$ steps. If $z \notin \cup_{i \leq q} \{d(x_i, \cdot) \leq \varepsilon\}$ for some unit length z , then $\{z, x_1, \dots, x_q\}$ is ε -separated, which contradicts \subset -maximality of $\{x_1, \dots, x_q\}$. Hence, $\{x_1, \dots, x_q\}$ provides an ε -net. The $\varepsilon/2$ -balls $\{d(x_i, \cdot) \leq \varepsilon/2\}$, $i \leq q$, are pairwise disjoint, and their union amounts to a subset of $\{\|\cdot\| \leq 1 + \varepsilon/2\}$. Hence, additivity, translation invariance, and the scaling property of the Lebesgue measure ν on \mathbb{R}^k imply

$$q \left(\frac{\varepsilon}{2}\right)^k \nu\{\|\cdot\| \leq 1\} \leq \left(1 + \frac{\varepsilon}{2}\right)^k \nu\{\|\cdot\| \leq 1\}.$$

Furthermore, $\nu\{\|\cdot\| \leq 1\} > 0$ together with the definition of a covering number implies $(1 + 2/\varepsilon)^k \geq q \geq N(\{\|\cdot\| = 1\}, \|\cdot\|, \varepsilon)$. \blacksquare

Proof of lemma 2.3. For any two element Ya and Yb of $\text{span}\{y_1, \dots, y_k\}$ define $\langle Ya, Yb \rangle_* = \langle a, Gb \rangle$. This expression inherits symmetry from G . Hence, the map $\langle \cdot, \cdot \rangle_* : V \times V \rightarrow \mathbb{R}$ is well defined as $Ya = Ya'$ implies $a - a' \in \ker Y = \ker G$. Bilinearity of $\langle \cdot, \cdot \rangle_*$ follows from its definition. Finally, positive semi-definiteness of G ensures $\langle Ya, Ya \rangle_* \geq 0$. Therein, equality

holds if and only if $a \in \ker G = \ker Y$, that is, $Ya = 0$. \blacksquare

Proof of lemma 2.4. If (V_j, U_j) symbolizes the j -th singular subspace pair for A , then $u \in U_j$ and symmetry of A imply the (in)equalities $\|Au\| = \sup_{\|v\|=1} \langle Au, v \rangle = \sup_{\|v\|=1} \langle u, P_{U_j} Av \rangle = \sup_{\|v\|=1} \langle u, AP_{V_j} v \rangle \leq \bar{\sigma}_j \|u\|$. Herein, P_L denotes the orthogonal projector onto a subspace L . Such maps are contractions; hence, $\|AP_{V_j} v\| = \bar{\sigma}_j \|P_{V_j} v\| \leq \bar{\sigma}_j$. The first equality expresses the self-duality of $\|\cdot\|$. The third equality is due to $\text{img } AP_{V_j^\perp} \subset U_j^\perp$. The subsequent inequality is an application of the Cauchy-Schwarz inequality. Consideration of $j = s$ —the number of distinct nonzero singular values of A —implies $V_s = U_s$ as $U_j \subset \sum_{i \leq s} U_i = \text{img } A = (\ker A)^\perp = \sum_{i \leq s} V_i$. In particular, if $s > 1$, then U_{s-1} is a subspace of V_s^\perp . Consequently, $\|Au\| \geq \bar{\sigma}_{s-1} \|u\|$ for $u \in U_{s-1}$ and so forth. The equalities $U_j = V_j$ and polarization imply that restricting $A/\bar{\sigma}_j$ to V_j provides a unitary map $V_j \rightarrow V_j$. If id denotes the identity map on V_j , $j \leq s$, then $\langle u, (\text{id} - A^2/\bar{\sigma}_j^2)v \rangle = 0$ for all $u, v \in V_j$ implies $\text{img}(\text{id} - A^2/\bar{\sigma}_j^2) = \{0\}$. Thus, one factor in $(\text{id} - A/\bar{\sigma}_j)(\text{id} + A/\bar{\sigma}_j) = (\text{id} - A^2/\bar{\sigma}_j^2)$ must have a nontrivial kernel, that is, $Av' \in \{\bar{\sigma}_j v', -\bar{\sigma}_j v'\}$ for some unit length $v' \in V_j$. If $V_j \cap (\text{span}\{v'\})^\perp$ is nontrivial, then the same argument applies. The recursion continues until all directions in V_j are exhausted. \blacksquare

Proof of lemma 2.5. Let $X = \sum_{i \leq \text{rk } X} \sigma_i u_i \langle v_i, \cdot \rangle$ and $Y = \sum_{i \leq \text{rk } Y} \sigma'_i u'_i \langle v'_i, \cdot \rangle$ be singular value decompositions of X and Y , respectively. If needed, then $\sigma_{\text{rk } X + p} = 0 = \sigma'_{\text{rk } Y + p}$ for all $p \geq 1$. The meaning of $u_{\text{rk } X + p}$, $v_{\text{rk } X + p}$, $u'_{\text{rk } Y + p}$, and $v'_{\text{rk } Y + p}$, $p \geq 1$, is immaterial. The distinct singular values of X are represented by $\bar{\sigma}_1, \dots, \bar{\sigma}_s$, wherein s provides the number of distinct singular values of X and $\bar{\sigma}_{s+p} = 0$ if $p \geq 1$. In addition, m_1 symbolizes the multiplicity of $\bar{\sigma}_1$. The equality $\|X\|_{\text{op}} = \bar{\sigma}_1 \sum_{i \leq \text{rk } Y} \sigma'_i = \langle Y, X \rangle \leq \bar{\sigma}_1 \sum_{i \leq m_1} \sigma'_i + \bar{\sigma}_2 \sum_{i > m_1} \sigma'_i$ reveals $\text{rk } Y \leq m_1$. Furthermore, the Cauchy-Schwarz inequality yields

$$\begin{aligned} \bar{\sigma}_1 \sum_{i \leq m_1} \sigma'_i &= \langle X, Y \rangle = \sum_{i \leq m_1} \sigma'_i \left[\bar{\sigma}_1 \sum_{j \leq m_1} \langle u'_i, u_j \rangle \langle v_j, v'_i \rangle \right] + \sum_{i \leq m_1} \sigma'_i \left[\sum_{m_1 < j \leq \text{rk } X} \sigma_j \langle u'_i, u_j \rangle \langle v_j, v'_i \rangle \right] \\ &\leq \bar{\sigma}_1 \sum_{i \leq m_1} \sigma'_i \sum_{j \leq m_1} |\langle u'_i, u_j \rangle| |\langle v_j, v'_i \rangle| + \bar{\sigma}_2 \sum_{i \leq m_1} \sigma'_i \sum_{m_1 < j \leq \text{rk } X} |\langle u'_i, u_j \rangle| |\langle v_j, v'_i \rangle| \\ &\leq \bar{\sigma}_1 \sum_{i \leq m_1} \sigma'_i \|a_i\| \|b_i\| + \bar{\sigma}_2 \sum_{i \leq m_1} \sigma'_i \sqrt{1 - \|a_i\|^2} \sqrt{1 - \|b_i\|^2} \leq \sigma_1 \sum_{i \leq m_1} \sigma'_i + 0, \quad \langle 2.11 \rangle \end{aligned}$$

wherein $a_i = (a_{1,i}, \dots, a_{m_1,i}) = (\langle u'_i, u_1 \rangle, \dots, \langle u'_i, u_{m_1} \rangle)$ as well as $b_i = (b_{1,i}, \dots, b_{m_1,i}) = (\langle v'_i, v_1 \rangle, \dots, \langle v'_i, v_{m_1} \rangle)$. The second inequality is due to the invariance of $\|\cdot\|$ to changes of the signs of the entries of its argument,

$$1 = \|u'_i\|^2 \geq \|P_{\text{img } X} u'_i\|^2 = \left\| \sum_{j \leq \text{rk } X} u_j \langle u_j, u'_i \rangle \right\|^2 = \sum_{j \leq m_1} a_{j,i}^2 + \sum_{m_1 < j \leq \text{rk } X} \langle u'_i, u_j \rangle^2,$$

and $1 - \|b_i\|^2 = \sum_{m_1 < j \leq \text{rk } X} \langle v'_i, v_j \rangle^2$. Moreover, the Cauchy-Schwarz inequality yields

$$\sqrt{\|a_i\|^2 \|b_i\|^2} + \sqrt{(1 - \|a_i\|^2)(1 - \|b_i\|^2)} = \left\langle \left(\frac{\sqrt{\|a_i\|^2}}{\sqrt{1 - \|a_i\|^2}} \right), \left(\frac{\sqrt{\|b_i\|^2}}{\sqrt{1 - \|b_i\|^2}} \right) \right\rangle \leq 1$$

and thereby $\bar{\sigma}_1 \|a_i\| \|b_i\| + \bar{\sigma}_2 \sqrt{1 - \|a_i\|^2} \sqrt{1 - \|b_i\|^2} \leq \bar{\sigma}_1$, which in turn generates the final inequality in $\langle 2.11 \rangle$. The resulting equalities in $\langle 2.11 \rangle$ require $\|a_i\| = 1 = \|b_i\|$ and $a_i = b_i$.

In fact, the first, second, and third inequality in <2.11> necessarily hold for each of the two main summands individually. Consequently, $u'_i = \sum_{j \leq m_1} a_{j,i} u_j = U_1 a_i$, $v'_i = \sum_{j \leq m_1} a_{j,i} v_j = V_1 a_i$, and $Y = \sum_{i \leq m_1} \sigma'_i u'_i \langle v'_i, \bullet \rangle = U_1 B \langle V_1, \bullet \rangle$, wherein $U_1 = [u_1 \cdots u_{m_1}]$, $V_1 = [v_1 \cdots v_{m_1}]$, and $B = \sum_{i \leq m_1} \sigma'_i a_i a_i^\top$. The matrix B is positive semi-definite and satisfies $\|B\|_{\text{nuc}} = \text{tr } B = \sum_{i \leq m_1} \sigma'_i \text{tr } a_i a_i^\top = \|Y\|_{\text{nuc}} = 1$ as $\text{tr } a_i a_i^\top = \langle a_i, a_i \rangle = \|a_i\|^2 = 1$.

Conversely, if $Y = U_1 B \langle V_1, \bullet \rangle$, then

$$\langle U_1 B \langle V_1, \bullet \rangle, \bar{\sigma}_1 U_1 \langle V_1, \bullet \rangle + \sum_{m_1 < j \leq \text{rk } X} \sigma_j u_j \langle v_j, \bullet \rangle \rangle = \bar{\sigma}_1 \text{tr} [\langle U_1, U_1 \rangle B \langle V_1, V_1 \rangle] = \bar{\sigma}_1 \text{tr } B$$

guarantees $\langle Y, X \rangle = \|X\|_{\text{op}} \|Y\|_{\text{nuc}} = \bar{\sigma}_1$. ■

Proof of corollary 2.6. Let U_1 consists of columns (in the given order) $u_1^+, \dots, u_{m'_1}^+, u_1^-, \dots, u_{m''_1}^-$ and likewise with V_1 and $u_1^+, \dots, u_{m'_1}^+, -u_1^-, \dots, -u_{m''_1}^-$. Then, lemma 2.5 ensures the existence of a positive semidefinite $S \in \mathbb{S}^{m_1}$ such that $B = U_1 S \langle V_1, \bullet \rangle$ and $\text{tr } S = 1$. Symmetry of B implies the equality of

$$\begin{aligned} \langle u_i^+, B u_j^- \rangle &= \langle e_i, -S e_{m'_1+j} \rangle = -s_{i,m'_1+j} \quad \text{and} \\ \langle B u_i^+, u_j^- \rangle &= \langle S e_i, e_{m'_1+j} \rangle = s_{m'_1+j,i} , \end{aligned}$$

wherein e_i symbolizes the i -th standard basis element of \mathbb{R}^{m_1} . Symmetry of S therefore guarantees that $s_{i,m'_1+j} = s_{m'_1+j,i} = 0$ for all $i \leq m'_1$, $j \leq m''_1$. The existence of a spectral decomposition validates the final claim. ■

3. Regularized least-squares estimation

3.1. Basic convex analysis

3.1.1. Geometry of convex sets

A subset C of a Euclidean space W is *convex* if it contains the line segment joining any two of its elements, that is, $x, y \in C$, $c \in [0, 1]$ implies $cx + (1 - c)y \in C$. The latter ensures $\sum_{i \leq k} c_i x_i \in C$ whenever $x_1, \dots, x_k \in C$ and $c_i \in [0, 1]$ with $\sum_{i \leq k} c_i = 1$. Such linear combinations are called *convex combinations*. The convex combinations of elements of an arbitrary subset $C' \subset W$ form a convex set—called the *convex hull* $\text{conv } C'$ of C' . Convex combinations are linear combinations and therefore compatible with linear maps L on W in the sense that $\text{conv}\{Lx \mid x \in D\} = \{Ly \mid y \in \text{conv } D\}$ for any $D \subset W$.

Examples of convex sets include *hyperplanes* $\{\langle a, \cdot \rangle = b\}$, wherein $a \neq 0$ and $b \in \mathbb{R}$, as well as the associated *halfspace* $\{\langle a, \cdot \rangle \leq b\}$. Embedding $a/\|a\|$ into an orthonormal basis of W identifies elements of the latter set as those extending into direction a by exactly $b/\|a\|$ and no more than $b/\|a\|$, respectively. Panel (A) of figure 3.1 illustrates this interpretation: it shows a point $a \in \mathbb{R}^2$ and a corresponding halfspace with threshold b . Herein, b is negative, and elements of the illustrated halfspace—those with coordinate with respect to $a/\|a\|$ no larger than $b/\|a\|$ —point in direction $-a$. Intersections of halfspaces are convex; more generally, convexity is preserved under arbitrary intersections.

Continuity of $\langle a, \cdot \rangle$ ensures that hyperplanes and halfspaces are closed. The same applies to the convex hull of a finite set $C' = \{x_1, \dots, x_k\}$. In fact, if a sequence $(\sum_{i \leq k} c_{i,n} x_i)$ in $\text{conv } C'$ converges as $n \rightarrow \infty$, then sequential compactness of $[0, 1]^k$ guarantees the existence of a subsequence $(\sum_{i \leq k} c_{i,n_\ell} x_i)$ with converging coefficients. The (common) limit of both sequences lies in $\text{conv } C'$ by virtue of continuity of $c \mapsto \sum_{i \leq k} c_i$. Herein, the coefficients $c_{i,n}$ need not be uniquely determined by the sequence elements; then, different choices lead to different subsequences but the same conclusion. The open ball $\{\|x - \cdot\| < 1\}$ shows that convex sets need not be closed.

If a convex set C is closed, then $x \notin C$ exhibits a unique closest substitute x' in C . Its existence is a topological feature. More specifically, the compact set $C_n = \{\|x - \cdot\| \leq n\} \cap C$ is nonempty for some $n_* \in \mathbb{N}$, and the continuous function $\|x - \cdot\|$ attains its minimum $\|x - x'\| \leq n_*$ on C_{n_*} at some $x' \neq x$. Hence, if $x'' \in C$, then either $\|x - x''\| > n_* \geq \|x - x'\|$ or $x'' \in C_{n_*}$. Convexity of C characterizes this closest substitute x' by the inequalities $\langle x - x', x'' - x' \rangle \leq 0$, $x'' \in C$. In fact, if the converse equality holds for some $x'' \in C$, then $(1 - c)x' + cx'' = x' + c(x'' - x') \in C$ for all $c \in [0, 1]$, and the latter element is a closer substitute to x for small $c > 0$. If $x', x'' \in C$ both satisfy this criterion, then the inequality $0 \geq \langle x - x', x'' - x' \rangle = -\langle x - x'', x' - x'' \rangle + \|x'' - x'\|^2 \geq \|x'' - x'\|^2$ implies $x' = x''$. Thus, the intersection $\cap_{x'' \in C} \{\langle x - \cdot, x'' - \cdot \rangle \leq 0\}$ contains single element—the closest substitute, and therefore the criterion is sufficient.

This criterion, namely $C \subset \{\langle x - x', \cdot - x' \rangle \leq 0\}$, shows that the angle between the approximation error $x - x'$ and the difference $x'' - x'$ is no less than $\pi/2$ whenever $x'' \in C$. Panel (B) of figure 2.2 illustrates this relation for an element $x \in \mathbb{R}^2$. This angular inequality generalizes the orthogonality condition required of closest substitutes in case

convex
convex combination
convex hull
hyperplane
halfspace




Figure 3.1

The figure illustrates the notions of (supporting) hyperplanes and normal cones. Panel (A) shows the hyperplane $\{\langle a, \cdot \rangle = b\}$ and the halfspace $\{\langle a, \cdot \rangle \leq b\}$ corresponding to a nonzero point $a \in \mathbb{R}^2$ and $b < 0$. Panel (B) considers $x \in \mathbb{R}^2$ outside the closed and convex set C . The point x exhibits a closest substitute x' in C . The difference $d = x - x'$ lies in the normal cone of C at x' , that is, $\langle d, x'' - x' \rangle \leq 0$ for all $x'' \in C$. Equivalently, the angle θ between d and $x'' - x'$ exceeds $\pi/2$, and the hyperplane $\{\langle d, \cdot \rangle = \langle d, x' \rangle\}$ supports C at x' .

of subspaces. In fact, subspaces are convex, and orthogonal projections provide the respective closest substitutes. The *normal cone* $\text{ncone}(C, x') = \bigcap_{x'' \in C} \{\langle \cdot, x'' - x' \rangle \leq 0\}$ of C at x' comprises all elements of W , which satisfy the same angular inequalities as the approximation error. In particular, $0 \in \text{ncone}(C, x')$ for all $x' \in C$. Moreover, if $d \in \text{ncone}(C, x')$ and $t \geq 0$, then $td \in \text{ncone}(C, x')$. A subset $K \subset W$ with this property, that is, $t \geq 0, x \in K$ imply $tx \in K$, is called a *cone*, and $\text{ncone}(C, x')$ is of this type.

If d is any nonzero element of $\text{ncone}(C, x')$, then x' lies in the hyperplane $\{\langle d, \cdot \rangle = \langle d, x' \rangle\}$, whose associated halfspace $\{\langle d, \cdot \rangle \leq \langle d, x' \rangle\}$ contains C . Such hyperplanes $\{\langle a, \cdot \rangle = b\}$, $a \neq 0$, that is, with $\{\langle a, \cdot \rangle = b\} \cap C \neq \emptyset$ and $\{\langle a, \cdot \rangle \leq b\} \supset C$, provide *supporting hyperplanes* of C at $x \in \{\langle a, \cdot \rangle = b\} \cap C$ and generally satisfy $a \in \text{ncone}(C, x)$.

If C is a convex superset of some nonempty open set, then $\text{ncone}(C, x) \neq \{0\}$ whenever $x \in C$ is a *boundary point* of C , that is, lies in all closed supersets of C but outside its *interior* $\text{int} C$ —the union of all its open subsets. Convexity guarantees that boundary points x are limit points of $\text{int} C$: if $y \in \text{int} C$ and $c \in (0, 1]$ with $\{\|y - \cdot\| < \delta\} \subset \text{int} C$ for some $\delta > 0$, then $(1 - c)x + cy + e = (1 - c)x + c(y + e/c) \in C$, whenever $\|e\| < c\delta$, thus, $(1 - c)x + cy \in \text{int} C$. The same argument shows that $\text{int} C$ is itself convex and therefore contains the convex hull of any finite sequence x_1, \dots, x_k of its elements—a closed set. A boundary point $x \in C$ exhibits a closest substitute $x' \neq x$ in $\text{conv}\{x_1, \dots, x_k\}$. In particular, $\bigcap_{i \leq k} \{\langle \cdot, x_i - x \rangle \leq 0\} \cap \{\|\cdot\| = 1\}$ is nonempty: the scaled error $d = (x - x')/\|x - x'\|$ satisfies $\langle d, x_i - x \rangle = \langle d, x_i - x' \rangle - \|x' - x\| < 0$. Hence, the closed subsets $\{\langle \cdot, x'' - x \rangle \leq 0\} \cap \{\|\cdot\| = 1\}$, indexed by elements x'' of the interior of C , of the compact set $\{\|\cdot\| = 1\}$ exhibit the finite intersection property. Thus, there

normal cone

cone

supporting hyperplane

boundary point

interior




Figure 3.2

The figure shows the closed nuclear ball $\{\|\cdot\|_{\text{nuc}} \leq \ell\}$, $\ell > 0$, in \mathbb{S}^2 . Points are represented by their coordinates with respect to the orthonormal basis $\bar{B}_{1,1}, \bar{B}_{1,2}, \bar{B}_{2,2}$ (example (c) in section 2.1.1). The boundaries of the two ellipses consist of matrices of the form $\ell u u^T$ and $-\ell u u^T$ with $\|u\| = 1$, respectively. A set containing only one of these matrices forms an (exposed) face of $\{\|\cdot\|_{\text{nuc}} \leq \ell\}$. The normal cones of the ball at its boundary points x and x' are given by all multiples of d (the scaled identity matrix) and d' with nonnegative coefficient, respectively.

exists a unit length d with $\text{int } C \subset \{\langle d, \cdot - x \rangle \leq 0\}$. The latter superset is a halfspace, hence closed, and consequently contains all elements of C . As a consequence, d lies in $\text{ncone}(C, x)$, and the associated hyperplane $\{\langle d, \cdot \rangle \leq \langle d, x \rangle\}$ supports C at x .

Conversely, if $x \in \text{int } C$, then $\{\|\cdot - x\| < \delta\} \subset C$ for some $\delta > 0$. Thus, $x + td'$ is an element of C whenever $d' \in W$ is of unit length and $|t| < \delta$. Consequently, if $d \in \text{ncone}(C, x)$, then suitable scaling ensures $\|d\| \leq 1$. Furthermore, $0 \geq \langle d, x + (\delta/2)d - x \rangle = \delta\|d\|^2/2 \geq 0$, and therefore $d = 0$. It follows that $\text{ncone}(C, x) = \{0\}$.

If C is an arbitrary convex subset of W , then the subset $\{d \in \text{ncone}(C, \cdot)\}$ of C is convex: two of its elements x, y satisfy $\langle d, x - y \rangle = 0$; thus, $\langle d, x' - cx - (1 - c)y \rangle = \langle d, x' - y \rangle \leq 0$ for any $c \in [0, 1]$, $x' \in C$. Moreover, if $c \in (0, 1)$, $x, y \in C$, and $cx + (1 - c)y \in \{d \in \text{ncone}(C, \cdot)\}$, then $\langle d, x - cx - (1 - c)y \rangle \leq 0$ and $\langle d, y - cx - (1 - c)y \rangle \leq 0$ imply $\langle d, y - x \rangle = 0$. In addition, if $x' \in C$, then $\langle d, x' - x \rangle = \langle d, x' - x + (1 - c)(x - y) \rangle \leq 0$ and likewise for y , which reveals $x, y \in F$. More generally, a convex subset F of C provides a *face* of C if $x, y \in C$, $c \in (0, 1)$, and $cx + (1 - c)y \in F$ imply $x, y \in F$. Both \emptyset and C are faces of C . The other faces of C are termed *proper*. If $\text{int } C \neq \emptyset$ and $\{d \in \text{ncone}(C, \cdot)\}$ is nonempty with nonzero d , then the latter set amounts to a proper face of C containing only boundary points. Such a face, that is, the intersection of C with one of its supporting hyperplanes, is called an *exposed face* of C with *exposing element* d .

Figure 3.2 illustrates these concepts with C equal to the nuclear norm ball $\{\|\cdot\|_{\text{nuc}} \leq \ell\}$, $\ell > 0$, in \mathbb{S}^2 . The spectral decomposition of symmetric matrices implies that this ball equals the convex hull of $\{\pm \ell u \langle u, \cdot \rangle \mid \|u\| = 1\}$. All elements x on the inside of and




Figure 3.3

The figure illustrates the interior and the boundary points of a convex set as well as the epigraph and subgradients of a convex function. Panel (A) shows (parts of) the epigraph of a convex function f defined on a convex subset $C \subset \mathbb{R}$. A point $x \in \text{int } C$ leads to an interior point $(x, f(x) + \delta)$, $\delta > 0$, of $\text{epi } f$. This conclusion does not hold for boundary points x' . Panel (B) considers the same function f and illustrates that elements d of the subdifferential $\partial f(x'')$ of f at x'' are in one-to-one correspondence with elements $(d, -1)$ of the normal cone of $\text{epi } f$ at $(x'', f(x''))$. The latter induce halfspaces which contain $\text{epi } f$.

on the upper ellipse contain the 2×2 identity matrix $I = [e_1 \ e_2] = \begin{pmatrix} 1 & \\ & 1 \end{pmatrix}$ in their normal cones, thus, form an exposed face. Figure 3.2 suggests that all proper and closed faces of $\{\|\cdot\|_{\text{muc}} \leq \ell\}$ are exposed. This is correct even for $m > 2$, however, is not needed here.

identity matrix

3.1.2. Convex functions

A real-valued function f defined on a convex subset C of a Euclidean space W is called *convex* if its *epigraph* $\text{epi } f = \{(x, t) \mid x \in C, f(x) \leq t\}$ is a convex subset of $W \times \mathbb{R}$. Herein, the latter (product) space is endowed with coordinate-wise linear operations and the inner product $\langle (x, t), (x', t') \rangle = \langle x, x' \rangle + tt'$. Convexity of C ensures that f is convex if and only if $cf(x) + (1 - c)f(y) \geq f(cx + (1 - c)y)$ for all $x, y \in C$, $c \in [0, 1]$. Panel (A) of figure 3.3 illustrates this case. Panel (B) of figure 3.4 shows a nonconvex function g .

convex epigraph

Suitable combinations of convex functions inherit convexity from their ingredients. If the pointwise supremum $\bar{f}(x) = \sup_{f \in F} f(x)$ over a set F of convex functions defined on a common set C is finite on a (necessarily) convex subset $C' \subset C$, then the epigraph of $C' \ni x \mapsto \bar{f}(x)$ equals the intersection $\bigcap_{f \in F} \text{epi } f$, thus, is convex. The second characterization of convexity ensures that linear combinations $\sum_{i \leq n} c_i f_i$ of convex functions $f_1, \dots, f_n \in F$ are convex provided that the coefficients c_1, \dots, c_n are nonnegative.

Lemma 3.1. *If f is a real-valued and convex function on a convex subset C of a Euclidean space W and x lies in the interior $\text{int } C$, then $(x, f(x) + \delta) \in \text{int epi } f$ for all $\delta > 0$.*

Panel (A) of figure 3.3 visualizes this assertion, that is, the subset $\{f < f(x) + \delta - \varepsilon'\}$ of C contains an open superset of $\{x\}$, and the open ball $\{\|(x, f(x) + \delta) - \bullet\| < \varepsilon'\}$ is therefore contained in the interior of $\text{epi } f$ if its radius $\varepsilon' > 0$ is sufficiently small. The argument on page 79 in appendix 3.b considers the general case.

In summary, $(x, f(x))$ is a boundary point of $\text{int epi } f \neq \emptyset$ if $x \in \text{int } C$. Hence, there exists a hyperplane $\{\langle a, \bullet \rangle \leq b\}$, $a \neq 0$, supporting $\text{epi } f$ at $(x, f(x))$. Writing $a = (a_1, a_2) \in W \times \mathbb{R}$ leads to the inequalities $a_2 t \leq b - \langle a_1, x' \rangle$, $(x', t) \in \text{epi } f$. If $x' \in C$ is fixed, then the righthand side is real number while $t \geq f(x')$ can be arbitrary large, thus, $a_2 \leq 0$. Moreover, $a_2 = 0$ implies $\langle a_1, x' \rangle \leq b$ for all $x' \in C$, wherein equality holds if $x' = x$. This scenario cannot occur since $a_2 = 0$ necessitates $a_1 \neq 0$ and $x + \delta a_1 \in C$ for small $\delta > 0$. Consequently, $f(x') \geq b/a_2 - \langle a_1/a_2, x' \rangle = f(x) + \langle -a_1/a_2, x' - x \rangle$ for all $x' \in C$. The implied inequality between the leftmost and rightmost term—uniformly across $x' \in C$ —identifies $-a_1/a_2 \in W$ as a *subgradient* of f at x . That is, $d \in W$ is a subgradient of f at x if and only if $(d, -1)$ lies in the normal cone of $\text{epi } f$ at $(x, f(x))$. Panel (B) of figure 3.3 illustrates this one-to-one correspondence. Therein, $d \in \mathbb{R}$ provides a subgradient of the convex function f at x'' , that is, $f \geq f(x'') + d(\bullet - x'')$; $(d, -1)$ lies in the normal cone of $\text{epi } f$ at $(x'', f(x''))$; and therefore the halfspace $\{\langle (d, -1), \bullet \rangle \leq \langle (d, -1), (x'', f(x'')) \rangle\}$ contains $\text{epi } f$. Panel (A) and panel (B) of figure 3.1 consider the concept of a halfspace and a normal cone, respectively.

In particular, the set $\partial f(x)$ of subgradients of f at x —called the *subdifferential* of f at x —is nonempty for all $x \in \text{int } C$. The existence of a subgradient $d \in \partial f(x)$ guarantees the inequality $\inf_{\|y-x\| \leq c\varepsilon} f(y) \geq f(x) - c\varepsilon \|d\|$. The latter and lemma 3.1 imply that a convex function f on the convex set C is continuous at all $x \in \text{int } C$. Panel (A) of figure 3.3 illustrates that this conclusion does not (necessarily) hold for boundary points x' of C . In fact, f may jump upwards at x' , which rules out $\partial f(x') \neq \emptyset$.

The notion of a subgradient generalizes the concept of a gradient. More specifically, the function f is called *differentiable* at $x \in \text{int } C$ if there exists a (necessarily unique) element $\nabla f(x) \in W$ —called the *gradient* of f at x —such that $\Delta(h) = (f(x+h) - f(x) - \langle \nabla f(x), h \rangle) / \|h\|$, wherein $h \in \{h' \in W \mid h' \neq 0, x+h' \in C\}$, converges to zero as $\|h\| \rightarrow 0$. If, in addition, f is convex and $d \in \partial f(x)$, then $h_n = (d - \nabla f(x)) / n \neq 0$ implies the inequalities $0 < \|d - \nabla f(x)\| = \langle d - \nabla f(x), h_n \rangle / \|h_n\| \leq \Delta(h_n)$, $n \in \mathbb{N}$. The latter contradict the definition of $\nabla f(x)$; thus, the equality $\partial f(x) = \{\nabla f(x)\}$ holds.

A scaled norm $\lambda \|\bullet\|'$, $\lambda > 0$, defined on W is convex due to homogeneity and the triangle inequality. Thus, $\partial \lambda \|x\|' \neq \emptyset$ for all $x \in W$. Lemma 3.2 contains a geometric characterization of these subdifferentials; it relies on the dual norm $\sup_{\|y\|'=1} \langle y, \bullet \rangle$ of $\|\bullet\|'$ introduced in section 2.5.2 and is proved on page 79 in appendix 3.b.

Lemma 3.2. *The subdifferential of a scaled norm $\lambda \|\bullet\|'$, $\lambda > 0$, on W is given by*

$$\begin{aligned} \partial \lambda \|0\|' &= \{\|\bullet\|'_d \leq \lambda\} \quad \text{and} \\ \partial \lambda \|x\|' &= \text{ncone}(\{\|\bullet\|' \leq \|x\|'\}, x) \cap \{\|\bullet\|'_d = \lambda\} \\ &= \{\langle \bullet, x \rangle = \|\bullet\|'_d \|x\|'\} \cap \{\|\bullet\|'_d = \lambda\}, \quad x \neq 0, \end{aligned}$$

wherein $\|\bullet\|'_d = \sup_{\|y\|'=1} \langle y, \bullet \rangle$ symbolizes the dual norm of $\|\bullet\|'$.




Figure 3.4

The figure visualizes the first equality in lemma 3.2 for $\lambda = 1$ and $\|\cdot\|' = \|\cdot\|_{\text{nuc}}$ as well as the subgradients and conjugate function of a nonconvex function. Panel (A) shows the (partial) graph of $\|\cdot\|_{\text{nuc}}$ restricted to the linear space $\text{span}\{\bar{B}_{1,1}, \bar{B}_{2,2}\}$ of diagonal matrices—a subspace of \mathbb{S}^2 —alongside the set $\{\|\cdot\|_{\text{op}} = 1\} \cap \text{span}\{\bar{B}_{1,1}, \bar{B}_{2,2}\}$. Herein, diagonal matrices are represented by their coordinate vectors with respect to $\bar{B}_{1,1}, \bar{B}_{2,2}$. By lemma 3.2, $D \in \{\|\cdot\|_{\text{op}} \leq 1\}$ provide the subgradients of (the restricted) $\|\cdot\|_{\text{nuc}}$ at 0. Panel (B) considers the epigraph of a nonconvex function g : convex combinations of $(x', g(x'))$ and $(x, g(x))$ with coefficients $c, 1 - c \notin \{0, 1\}$ lie outside epi g . Moreover, the conjugate g_* satisfies $-g_*(d) \leq g(x) - \langle d, x \rangle$, wherein equality holds at x' with $d \in \partial g(x')$.

Panel (A) of figure 3.4 illustrates the first equality in lemma 3.2 for $\lambda = 1$ and the restriction of $\|\cdot\|' = \|\cdot\|_{\text{nuc}}$ to the linear space $\text{span}\{\bar{B}_{1,1}, \bar{B}_{2,2}\} \subset \mathbb{S}^2$ of diagonal 2×2 matrices. Herein, the relevant dual norm equals $\|\cdot\|'_d = \|\cdot\|_{\text{op}}$ (restricted to the diagonal 2×2 matrices). Every diagonal matrix D with $\|D\|_{\text{op}} \leq 1$ provides a subgradient (of the restriction) of $\|\cdot\|_{\text{nuc}}$ at 0. Then, the pair $(D, -1) \in \text{span}\{\bar{B}_{1,1}, \bar{B}_{2,2}\} \times \mathbb{R}$ lies in the normal cone of the epigraph (of the restriction) of $\|\cdot\|_{\text{nuc}}$ at 0, and the resulting inequalities $\langle D, D' \rangle \leq \|D'\|_{\text{nuc}}$, $D' \in \text{span}\{\bar{B}_{1,1}, \bar{B}_{2,2}\}$, reflect the $\|\cdot\|_{\text{op}}/\|\cdot\|_{\text{nuc}}$ -duality.

Lemma 3.2 shows that $G \in \mathbb{S}^m$ lies in the subdifferential of $\lambda\|\cdot\|_{\text{nuc}}$ at some nonzero $\Theta \in \{\|\cdot\|_{\text{nuc}} = \ell\} \subset \mathbb{S}^m$, $\ell > 0$, if and only if $\|G\|_{\text{op}} = \lambda$ and Θ lies in the exposed face $\{G \in \text{ncone}(B_{\text{nuc}}, \cdot)\}$ of the nuclear norm ball $B_{\text{nuc}} = \{\|\cdot\|_{\text{nuc}} \leq \ell\}$. Either of these cases is equivalent to $\|\Theta\|_{\text{nuc}} = \ell > 0$ and G being an element of the exposed face $\{\Theta \in \text{ncone}(B_{\text{op}}, \cdot)\}$ of $B_{\text{op}} = \{\|\cdot\|_{\text{op}} \leq \lambda\}$. Finally, if $\|\Theta\|_{\text{nuc}} = \ell > 0$, then lemma 3.2 leads to an expression for Θ/ℓ in terms of a spectral decomposition of G via corollary 2.6.

If f is bounded below, then $\{0 \in \partial f(\cdot)\}$ gathers the *minimizers* $x' \in \text{argmin}_{x \in C} f(x)$ of f , that is, elements of C with $f(x') = \inf_{x \in C} f(x)$. More generally, if $d \in \partial f(y)$ with $y \in C$, then the nonempty set $\{d \in \partial f(\cdot)\}$ consist of the orthogonal projections

minimizer

onto W —a subspace of $W \times \mathbb{R}$ —of the elements of $\{(x, f(x)) \mid d \in \partial f(x)\} = \{(d, -1) \in \text{ncone}(\text{epi } f, \bullet)\}$. Thus, linearity of projectors ensures convexity of $\{d \in \partial f(\bullet)\}$, and f equals the *affine function* $f(y) + \langle d, \bullet - y \rangle$ on this set. Panel (B) of figure 3.4 reveals that this argument does not apply to nonconvex functions. Furthermore, if $\partial f(x) \neq \emptyset$ and $\partial f(x) \cap \partial f(y) = \emptyset$ for all $x, y \in C$ with $x \neq y$, then f is *strictly convex*, that is, the strict inequality $f(cx + (1-c)y) < cf(x) + (1-c)f(y)$ holds whenever $x \neq y$ and $c \in (0, 1)$.

Strict convexity guarantees that f exhibits at most one minimizer. However, the set of ε -almost minimizers $\{f < \inf_{x \in C} f(x) + \varepsilon\}$, $\varepsilon > 0$, may be arbitrary large. A quantitative analog of strict convexity bounds this set. A differentiable function f defined on an open and convex set $C \subset W$ exhibits *strong convexity* with *curvature constant* $\kappa > 0$ if $f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\kappa}{2} \|y - x\|^2$ holds for all $x, y \in C$. Then $\{f < \inf_{x \in C} f(x) + \varepsilon\} \subset \{\|x' - \bullet\| < \sqrt{2\varepsilon/\kappa}\}$, wherein x' denotes the unique minimizer of f .

3.1.3. Conjugates and convex envelopes

The inequality $g \geq g(x) + \langle d, \bullet - x \rangle$ identifies $d \in W$ as a *subgradient* at x of a real-valued function g on a subset D of a Euclidean space W . The *subdifferential* $\partial g(x)$ of g at $x \in D$ gathers the subgradients of g at x . If $d \in \partial g(x)$, then the halfspace $\{\langle (d, -1), \bullet \rangle \leq b\}$, $b = \langle (d, -1), (x, g(x)) \rangle = \langle d, x \rangle - g(x)$, contains the *epigraph* $\text{epi } g = \{(x, t) \mid x \in D, g(x) \leq t\}$ of g —see panel (B) of figure 3.4, that is, $\langle d, x \rangle - g(x) = \sup_{x' \in D} (\langle d, x' \rangle - g(x')) = g_*(d)$. The quantity $g_*(d)$ is finite on a convex superset D_* of $\cup_{x \in D} \partial g(x)$. If D_* is nonempty, then $D_* \ni d \mapsto g_*(d)$ is referred to as the *conjugate* (function) of g .

By definition of the conjugate, the *Fenchel-Young inequality* $g(x) + g_*(d) \geq \langle x, d \rangle$ holds for all $x \in D$ and $d \in D_*$. The inequality becomes an equality if and only if $d \in \partial g(x)$ since $g_*(d) = \langle d, x \rangle - g(x)$ implies $\langle d, x' \rangle - g(x') \leq \langle d, x \rangle - g(x)$ for all $x' \in D$. This result implies that $-g_*(d) + \langle d, \bullet \rangle \leq g$ if $d \in D_*$. Conversely, if $b + \langle d, \bullet \rangle \leq g$ for some $b \in \mathbb{R}$, then $\langle d, \bullet \rangle - g \leq -b$, thus $g_*(d)$ is finite, and $b \leq -g_*(d)$. Thus, $D_* = \{d \in W \mid \text{there exists } b \in \mathbb{R} \text{ such that } b + \langle d, \bullet \rangle \leq g\}$, and if $d \in D_*$, then $-g_*(d)$ provides the largest intercept $b \in \mathbb{R}$ such that $b + \langle d, \bullet \rangle$ is an affine minorant of g . Panel (B) of figure 3.4 exemplifies this interpretation for $d \in \partial g(x')$: the difference $g - \langle d, \bullet \rangle \geq 0$ becomes minimal at x' ; hence, the inequality $b + \langle d, \bullet \rangle \leq g$ holds if and only if $b \leq -g_*(d)$.

The conjugate g_* of g is defined as the (pointwise) supremum of the affine functions $D_* \ni d \mapsto a_x(d) = \langle d, x \rangle - g(x)$, $x \in D$. Hence, $\text{epi } g_*$ equals $\cap_{x \in D} \text{epi } a_x$, wherein $\text{epi } a_x = \{(d, t) \mid d \in D_*, t \geq a_x(d) = \langle d, x \rangle - g(x)\} = \{\langle \bullet, (x, -1) \rangle \leq g(x)\} \cap D_* \times \mathbb{R}$. The intersection equals $\cap_{x \in D} \{\langle \bullet, (x, -1) \rangle \leq g(x)\}$ as the right hand side of $t \geq \langle d, x \rangle - g(x)$ is unbounded unless $d \in D_*$. Consequently, the epigraph $\text{epi } g_*$ is closed and convex.

The Fenchel-Young inequality implies $g(x) \geq \sup_{d \in D_*} (\langle x, d \rangle - g_*(d)) = g_{**}(x)$. Thus, the rightmost term is finite for elements x of a convex superset D_{**} of D . The *biconjugate* (function) g_{**} of g given by $D_{**} \ni x \mapsto g_{**}(x)$ amounts to the conjugate of g_* , hence, is a convex function with closed epigraph $\text{epi } g_{**} \supset \text{epi } g$. Moreover, its definition $g_{**} = \sup_{d \in D_*} (-g_*(d) + \langle d, \bullet \rangle)$ characterizes g_{**} as the (pointwise) supremum over all affine minorants of g . Panel (B) of figure 3.4 indicates this construction. If $f \leq g$ is a convex function on (a convex superset of) D_{**} , then $x \in \text{int } D_{**}$ implies $g_{**}(x) \geq f(x)$, which identifies g_{**} as the largest convex minorant—called *convex envelope*—of g on $\text{int } D_{**}$.

Finally, if $d \in \partial g(x) \subset D_*$ for $x \in D$, then $g_*(d') \geq \langle d', x \rangle - g(x) = \langle d', x \rangle - \langle d, x \rangle + g_*(d)$ for all $d' \in D_*$, that is, $x \in \partial g_*(d)$. Moreover, $g(x) + g_*(d) = \langle x, d \rangle \leq g_*(d) + g_{**}(x)$ proves $g_{**}(x) = g(x)$. Conversely, the equality $g(x) = g_{**}(x)$ together with $x \in \partial g_*(d)$ ensures $g(x) + g_*(x) = g_*(x) + g_{**}(x) = \langle x, d \rangle$, thus, $d \in \partial g(x)$.

3.2. Regularized least-squares

3.2.1. Primal formulation

The least-squares problem considered here amounts to the minimization of the criterion

$$l_\lambda(\Theta) = \frac{1}{2\bar{\mu}} \|Y - X\Theta\|^2 + \lambda \|\Theta\|_{\text{nuc}}, \quad \langle 3.1 \rangle$$

wherein $Y, X \in W^{\times m}$ are given linear maps from \mathbb{R}^m into a Euclidean space W ; Θ ranges over the $m \times m$ symmetric matrices in \mathbb{S}^m ; $\bar{\mu} > 0$ adjusts the scaling of the first summand; and $\lambda > 0$ controls the relative importance of the two summands of l_λ .

Usually W is spanned by finitely many real-valued and μ -square integrable functions defined on a finite measure space $(\Omega, \mathcal{F}, \mu)$. Then $\bar{\mu} = \mu\Omega$ is a natural choice, but not required by the results of this section, which hold irrespective of the particular value $\bar{\mu} > 0$. As an example, if W equals \mathbb{R}^n , that is, $Y, X \in \mathbb{R}^{n \times m}$, then the total mass equals $\bar{\mu} = n$. Section 2.1.3 presents further instances of this construct.

Both summands of l_λ , namely, $g = \lambda \|\cdot\|_{\text{nuc}} : \mathbb{S}^m \rightarrow \mathbb{R}$ and the composition $f \circ X$, wherein $f = \|Y - \cdot\|^2 / (2\bar{\mu})$ is convex and $X : \mathbb{S}^m \rightarrow W^{\times m}$ is linear, are convex. Therefore, the criterion function $\mathbb{S}^m \ni \Theta \mapsto l_\lambda(\Theta)$ is, too. In particular, $l_\lambda : \mathbb{S}^m \rightarrow \mathbb{R}$ is continuous, and therefore its *sublevel sets* $\{l_\lambda \leq t\}$, $t > 0$, are closed. The second summand $g = \lambda \|\cdot\|_{\text{nuc}}$ of l_λ implies that these sublevel sets are also bounded and therefore compact. Continuity and compactness guarantee the existence of a minimizer $\hat{\Theta} \in \mathbb{S}^m$.

The summands of l_λ mirror the twofold goal behind the minimization of $\langle 3.1 \rangle$. The first term ensures that a minimizer $\hat{\Theta}$ of l_λ yields a close substitute—in terms of $\|\cdot\|$ —to Y in form of $X\hat{\Theta}$. In addition, lemma 3.3 suggests that the second term promotes a low rank of minimizers $\hat{\Theta} \in \mathbb{S}^m$. Hence, minimizers $\hat{\Theta}$ of the criterion l_λ trade off “fidelity to the data” X, Y against their own complexity—expressed as the dimension of their image. A proof of lemma 3.3 follows on page 80 in appendix 3.b.

Lemma 3.3. *The restriction g' of $g = \lambda \|\cdot\|_{\text{nuc}}$ to the $\|\cdot\|_{\text{op}}$ -ball $H = \{\|\cdot\|_{\text{op}} \leq 1\} \subset \mathbb{S}^m$, wherein $\|\cdot\|_{\text{op}} : \mathbb{S}^m \rightarrow \mathbb{R}$, equals the convex envelope of $H \ni B \mapsto h'(B) = \lambda \text{rk } B$.*

The restriction on the $\|\cdot\|_{\text{op}}$ -length in lemma 3.3 is essential: if $\lambda \text{rk } B \geq a + \langle S, B \rangle$ for all $B \in \mathbb{S}^m$, then $\lambda m \geq \lambda \text{rk}(tS) \geq a + t\|S\|^2$ for all $t > 0$. Consequently, the conjugate of λrk (on \mathbb{S}^m) is merely defined at zero, and its biconjugate equals zero everywhere.

Lemma 3.3 is related to the equality $\text{conv}\{\pm uu^T \mid \|u\| = 1\} = \{\|\cdot\|_{\text{nuc}} \leq 1\}$. In fact, if convex functions f, f' satisfy $f(B) \leq f'(B) \leq \text{rk}(B)$ for all $B \in \{\|\cdot\|_{\text{op}} \leq 1\}$, then $\{\pm uu^T \mid \|u\| = 1\} \subset \{\text{rk} \leq 1\} \cap \{\|\cdot\|_{\text{op}} \leq 1\} \subset \{f' \leq 1\} \subset \{f \leq 1\}$. The latter two subsets are convex and thus contain $H' = \text{conv}\{\pm uu^T \mid \|u\| = 1\}$. Hence, lemma 3.3 implies $H' \subset \{\|\cdot\|_{\text{nuc}} \leq 1\} \subset \{f \leq 1\}$ for every convex f with $f \leq \text{rk}$ on $\{\|\cdot\|_{\text{op}} \leq 1\}$.

3.2.2. Dual formulation

The minimization of the criterion function l_λ in <3.1> over \mathbb{S}^m can be cast in alternative form. This so-called dual problem amounts to the maximization of the dual criterion shown in <3.3>. Its form derives from analyzing the sensitivity of the optimal value $\inf_{\Theta \in \mathbb{S}^m} l_\lambda(\Theta)$ to perturbations. Lemma 3.4 contains the essential ingredients of this approach. A proof of this result starts on page 80 in appendix 3.b.

Lemma 3.4. *The function $W^{\times m} \ni Z \mapsto v(Z) \in \mathbb{R}$, wherein*

$$v(Z) = \inf_{\Theta \in \mathbb{S}^m} \left[\frac{1}{2\bar{\mu}} \|Y - (X\Theta + Z)\|^2 + \lambda \|\Theta\|_{\text{nuc}} \right] = \inf_{\Theta \in \mathbb{S}^m} \left[f(X\Theta + Z) + g(\Theta) \right]$$

is convex. Its conjugate is defined for all $D \in W^{\times m}$ with $\|\langle X, D \rangle + \langle D, X \rangle\|_{\text{op}} \leq 2\lambda$ by

$$v_\star(D) = f_\star(D) + g_\star \left(-\frac{\langle X, D \rangle + \langle D, X \rangle}{2} \right), \quad \text{wherein}$$

$$f_\star : W^{\times m} \rightarrow \mathbb{R}, \quad D \mapsto \frac{\bar{\mu}}{2} \left[\|D + Y/\bar{\mu}\|^2 - \|Y/\bar{\mu}\|^2 \right], \quad g_\star : \{\|\cdot\|_{\text{op}} \leq \lambda\} \rightarrow \mathbb{R}, \quad M \mapsto 0.$$

Lemma 3.4 considers the scaled nuclear norm $g = \lambda \|\cdot\|_{\text{nuc}}$ as a real-valued function on \mathbb{S}^m —instead of $H = \{\|\cdot\|_{\text{op}} \leq \lambda^{-1}\} \subset \mathbb{S}^m$, which explains the difference of g_\star and the intermediate result, in particular, <A3.1> on page 80, used to prove lemma 3.3.

The Fenchel-Young inequality together with lemma 3.4 implies

$$l_\lambda(\Theta) \geq v(0) \geq \sup_{\|\langle X, M \rangle + \langle M, X \rangle\|_{\text{op}} \leq 2\lambda} -f_\star(M) \geq -f_\star(D) \quad <3.2>$$

for every symmetric Θ and $D \in \{\|\langle X, \cdot \rangle + \langle \cdot, X \rangle\|_{\text{op}} \leq 2\lambda\}$. Equality holds if and only if D is a subgradient of v at 0—the neutral element of the additive group $(W^{\times m}, +)$. Convexity of v on the open set $W^{\times m}$ guarantees the existence of such a subgradient \hat{D} . Consequently, the minimal value $\inf_{\Theta \in \mathbb{S}^m} l_\lambda(\Theta) = v(0)$ equals the supremum of

$$-f_\star(D) = \frac{\bar{\mu}}{2} \left[\|Y/\bar{\mu}\|^2 - \|D + Y/\bar{\mu}\|^2 \right], \quad \left\| \frac{\langle X, D \rangle + \langle D, X \rangle}{2} \right\|_{\text{op}} \leq \lambda, \quad <3.3>$$

which is attained at every element \hat{D} of $\partial v(0) \neq \emptyset$. This maximization exercise provides the *dual problem* to the minimization of the (primal) criterion l_λ in <3.1>.

Figure 3.5 (re-)interprets l_λ as $-(-f \circ X) + g$. The upper part of its panel (A) visualizes this difference for the case $W = \mathbb{R}$, $m = 1$ with $X = 1$ and in terms of the (partial) graphs of f and g . Its lower part contains the corresponding part of the graph of l_λ . The primal objective function l_λ is not smaller than the sum of the intercepts $-f_\star(D)$ and $-g_\star(-D)$ for every choice of (Θ, D) , wherein D and $-D$ correspond to affine minorants of f and g , respectively. The upper part also sheds light on the form of g_\star : in fact, there exists an affine minorant of $g = \lambda|\cdot|$ with slope $-D$ if and only if $|D| \leq \lambda$; in that case, the

dual problem




Figure 3.5

The figure illustrates (the relation between) the primal and dual formulation of the least-squares problem <3.1> for the case $W = \mathbb{R}$, $m = 1$, and $X = 1$. The upper part of panel (A) rephrases l_λ as the difference $-(-f) + g$; the lower part shows parts of its graph. The primal objective value $l_\lambda(\Theta)$ exceeds the sum of the intercepts $-f_*(D)$ and $-g_*(-D)$ for any Θ and D such that D and $-D$ correspond to affine minorants of f and g , respectively. Panel (B) reproduces the setting of panel (A) and shows that primal and dual objectives coincide for a minimizing/maximizing pair $(\hat{\Theta}, \hat{D})$, which satisfies $\hat{D} \in \partial f(\hat{\Theta})$ and $-\hat{D} \in \partial g(\hat{\Theta})$.

maximal intercept equals zero. Panel (B) of figure 3.5 concerns the identical setting. It illustrates that the equality $v(0) = l_\lambda(\hat{\Theta}) = -v_*(\hat{D})$ is possible.

If $(\hat{\Theta}, \hat{D}) \in \mathbb{S}^m \times \{\|\langle X, \cdot \rangle + \langle \cdot, X \rangle\|_{\text{op}} \leq 2\lambda\}$ is a minimizing/maximizing pair, then

$$\begin{aligned} 0 &= [f(X\hat{\Theta}) + g(\hat{\Theta})] + [f_*(\hat{D}) + g_*(\hat{G})] \\ &= [f(X\hat{\Theta}) + f_*(\hat{D}) - \langle X\hat{\Theta}, \hat{D} \rangle] + [g(\hat{\Theta}) + g_*(\hat{G}) - \langle X\hat{\Theta}, -\hat{D} \rangle], \end{aligned}$$

wherein $\hat{G} = -(\langle X, \hat{D} \rangle + \langle \hat{D}, X \rangle)/2$. The Fenchel-Young inequality together with $\langle X\hat{\Theta}, -\hat{D} \rangle = \langle \hat{\Theta}, \hat{G} \rangle$ implies that the latter two summands are generally nonnegative. Consequently, both summands equal zero, and the two pairs of optimality conditions

$$\begin{aligned} \hat{D} \in \partial f(X\hat{\Theta}) &= \left\{ \frac{X\hat{\Theta} - Y}{\bar{\mu}} \right\}, & X\hat{\Theta} \in \partial f_*(\hat{D}) &= \{\bar{\mu}\hat{D} + Y\}, \\ \hat{G} = -\frac{\langle X, \hat{D} \rangle + \langle \hat{D}, X \rangle}{2} \in \partial g(\hat{\Theta}) & \text{ and } & \hat{\Theta} \in \partial g_*(\hat{G}) &= \text{ncone}(\{\|\cdot\|_{\text{op}} \leq \lambda\}, \hat{G}) \end{aligned} \quad <3.4>$$

hold. The second twin in <3.4> follows from f and g being convex functions on the open sets $W^{\times m}$ and \mathbb{S}^m , respectively. More specifically, these properties guarantee $f = f_{**}$

and $g = g_{**}$, wherein f_{**} and g_{**} equal the biconjugate functions of f and g , respectively. Panel (B) of figure 3.5 reflects the first pair of necessary conditions: if $W = \mathbb{R}$ and X equals the identity, then an optimal \hat{D} provides a subgradient of f at the corresponding minimizer $\hat{\Theta}$ and at the same time $-\hat{D}$ lies in the subdifferential $\partial g(\hat{\Theta})$.

The conditions in <3.4> are also sufficient in the sense that if $(\hat{\Theta}, \hat{D}) \in \mathbb{S}^m \times W^{\times m}$ satisfies either set of conditions, then this pair is minimizing/maximizing. More specifically, either of its lower parts implies $\|\hat{G}\|_{\text{op}} \leq \lambda$, wherein $\hat{G} = -(\langle X, \hat{D} \rangle + \langle \hat{D}, X \rangle)/2$. In fact, if the lower part on the lefthand side of <3.4> holds, then $\lambda \|\Theta\|_{\text{nuc}} \geq \lambda \|\hat{\Theta}\|_{\text{nuc}} + \langle \hat{G}, \Theta - \hat{\Theta} \rangle$ for all $\Theta \in \mathbb{S}^m$. Corollary 2.6 guarantees the existence of $\Theta' \in \{\|\cdot\|_{\text{nuc}} = 1\}$ with $\langle \Theta', \hat{G} \rangle = \|\hat{G}\|_{\text{op}}$ and thus $\lambda \|\hat{\Theta}\|_{\text{nuc}} + \lambda \geq \lambda \|\hat{\Theta} + \Theta'\|_{\text{nuc}} \geq \lambda \|\hat{\Theta}\|_{\text{nuc}} + \langle \hat{G}, \Theta' \rangle$. In case of the righthand conditions, this inequality holds since subgradients exists only at points where g_* is defined. Furthermore, either pair in <3.4> implies the first equality in the display above <3.4>. Hence, the (in)equalities $0 = l_\lambda(\hat{\Theta}) + v_*(\hat{D}) \geq v(0) + v_*(\hat{D}) \geq 0$ ensure via <3.2> that $\hat{\Theta}$ and \hat{D} are optimal in <3.1> and <3.3>, respectively.

3.2.3. The least-squares solution set

Section 3.2.1 proofs the existence of a minimizer of the objective function l_λ in <3.1>, that is, an element $\hat{\Theta} \in \mathbb{S}^m$ with $l_\lambda(\hat{\Theta}) \leq l_\lambda(\Theta)$ for all $\Theta \in \mathbb{S}^m$. This section characterizes the set of minimizers of l_λ , which is denoted by $\text{argmin}_{\Theta \in \mathbb{S}^m} l_\lambda(\Theta)$ as in section 3.1.2.

The conjugate $f_* : W^{\times m} \rightarrow \mathbb{R}$ in lemma 3.4 is convex by construction. Its form implies differentiability with $\partial f_*(D) = \{\bar{\mu}D + Y\}$ and thereby strict convexity. The restriction of f_* to the set $\{\|\langle X, \cdot \rangle + \langle \cdot, X \rangle\|_{\text{op}} \leq 2\lambda\}$ inherits this property, and therefore exhibits at most one minimizer. Consequently, the second component of any minimizing/maximizing pair $(\hat{\Theta}, \hat{D})$ is uniquely determined. The upper parts of <3.4> ensure that $X\hat{\Theta} = \bar{\mu}\hat{D} + Y$ is, too. In particular, any two minimizers $\hat{\Theta}, \hat{\Theta}'$ of <3.1> provide boundary points of the nuclear norm ball $\{\|\cdot\|_{\text{nuc}} \leq \hat{\ell}\}$ with radius

$$\begin{aligned} \hat{\ell} = \|\hat{\Theta}\|_{\text{nuc}} = \|\hat{\Theta}'\|_{\text{nuc}} &= \frac{1}{\lambda} \left[\inf_{\Theta \in \mathbb{S}^m} l_\lambda(\Theta) - \frac{1}{2\bar{\mu}} \|Y - X\hat{\Theta}\|^2 \right] \\ &= \frac{1}{\lambda} \left[\inf_{\Theta \in \mathbb{S}^m} l_\lambda(\Theta) - \frac{1}{2\bar{\mu}} \|Y - X\hat{\Theta}'\|^2 \right]. \end{aligned}$$

If $\hat{\ell} = 0$, then the unique minimizer $\hat{\Theta}$ amounts to the $m \times m$ zero matrix. Otherwise, the lefthand side of <3.4> together with lemma 3.2 implies that $\hat{G} = -(\langle X, \hat{D} \rangle + \langle \hat{D}, X \rangle)/2$ exhibits $\|\cdot\|_{\text{op}}$ -length λ and that $\langle \hat{G}, \hat{\Theta} \rangle = \lambda \hat{\ell}$ for every minimizer $\hat{\Theta}$ of l_λ . The latter together with the requirement $X\hat{\Theta} = \bar{\mu}\hat{D} + Y$ leads to the assertion of lemma 3.5. Its proof starts on page 81 in appendix 3.b.

Lemma 3.5. *The least-squares criterion l_λ exhibits a unique minimizer $\hat{\Theta}$ if $\ker(\hat{G} \mp \lambda \text{id}) \cap \ker X = \{0\}$, wherein id symbolizes the identity map on \mathbb{R}^m .*

The condition of lemma 3.5 is generally satisfied. In fact, if $u \in \ker(\hat{G} - \lambda \text{id}) \cap \ker X$, then $\lambda u = \hat{G}u = -\langle X, \hat{D} \rangle u/2 \in \text{img} \langle X, \cdot \rangle = (\ker X)^\perp$, thus, $u = 0$. The case $u \in$




Figure 3.6

The figure shows a selection of rank one boundary points of $\{\|\cdot\|_{\text{nuc}} \leq \ell\} \subset \mathbb{S}^2$, $\ell > 0$, and the differences of elements of the same exposed face of $\{\|\cdot\|_{\text{nuc}} \leq \ell\}$. The two ellipses in panel (A) consist of the rank one boundary points. The panel also contains the spans of a selection of these matrices (represented by dots). Panel (B) shows two pairs $(\hat{\Theta}, \hat{\Theta}')$ and $(\hat{\Theta}'', \hat{\Theta}''')$ —each lying in an exposed face—alongside their component differences. In addition, the panel indicates the solution set of a linear equation. Both panels show coordinates with respect to $\bar{B}_{1,1}, \bar{B}_{1,2}, \bar{B}_{2,2}$; see (c) (section 2.1.1). Panel (A) omits the coordinate axes for visual clarity.

$\ker(\hat{G} + \lambda \text{id}) \cap \ker X$ is in analogy. Proposition 3.6 summarizes the preceding discussion. A proof of this assertion starts on page 81 of appendix 3.b.

Proposition 3.6. *The least-squares criterion l_λ in <3.1> exhibits a unique minimizer $\hat{\Theta}$. The latter equals zero if and only if $\|\langle X, Y \rangle + \langle Y, X \rangle\|_{\text{op}} \leq 2\bar{\mu}\lambda$.*

From a geometric perspective, the lefthand side of <3.4> shows that $\text{argmin}_{\Theta \in \mathbb{S}^m} l_\lambda(\Theta)$ equals the intersection of the exposed face $\{\hat{G} \in \text{ncone}(B_{\text{nuc}}, \cdot)\}$ of the nuclear norm ball $B_{\text{nuc}} = \{\|\cdot\|_{\text{nuc}} \leq \hat{\ell}\}$ and the set of solutions $\{X \cdot = \bar{\mu}\hat{D} + Y\}$.

If $m = 2$, then panel (A) of figure 3.6 indicates that the difference between two distinct elements $\hat{\Theta}$ and $\hat{\Theta}'$ of the same exposed face has rank two. In fact, all rank one matrices of $\|\cdot\|$ -length $\ell > 0$, that is, matrices of the form $\pm \ell uu^T$, $\|u\| = 1$, lie in the upper and lower ellipse shown in that panel. This panel also shows the spans of a selection—represented by dots—of these matrices. The neighboring panel (B) verifies this observation for two pairs $\hat{\Theta} \neq \hat{\Theta}'$ and $\hat{\Theta}'' \neq \hat{\Theta}'''$. Each of these pairs lies in an exposed face of the ball $\{\|\cdot\|_{\text{nuc}} \leq \ell\}$: the gray line connecting the two ellipses and the area circumscribed by the upper ellipse. This observations is, however, incompatible with $\Delta = \hat{\Theta} - \hat{\Theta}' \in \ker X$ unless X equals zero. After all, the equality $\text{rk } \Delta = 2$ ensures that its columns form a basis of \mathbb{R}^2 . Panel (B) illustrates this point by showing part of $\{X \cdot = \bar{\mu}\hat{D} + Y\}$ with $X = (0, 1)$ and $\bar{\mu}\hat{D} + Y = (0, 1)$.

3.3. A gradient descent algorithm

An ε -almost minimizer, $\varepsilon > 0$, of l_λ in <3.1> can be calculated via an iterative search procedure. Each iteration of the latter consists of two steps focusing on the first summand $f \circ X = \|Y - X \cdot\|^2/(2\bar{\mu})$ of l_λ and its second summand $g = \lambda \|\cdot\|_{\text{nuc}}$, respectively. In the $k(\geq 1)$ -th step, the former turns the search point $\widehat{\Theta}^{(k-1)}$ generated by the previous step into an intermediate $\widetilde{\Theta}^{(k)} \in \mathbb{S}^m$ with $(f \circ X)(\widetilde{\Theta}^{(k)}) \leq (f \circ X)(\widehat{\Theta}^{(k-1)})$. The latter step adjusts $\widetilde{\Theta}^{(k)}$ to the presence of the second summand g and thereby generates a new search point $\widehat{\Theta}^{(k)}$ such that $l_\lambda(\widehat{\Theta}^{(k)}) \leq l_\lambda(\widehat{\Theta}^{(k-1)})$. Proposition 3.7 asserts that the resulting recipe <3.7> generates ε -almost minimizers of l_λ for any given $\varepsilon > 0$. If $\|X\|_{\text{op}} = 0$, then the minimizer $\widehat{\Theta}$ is given by the $m \times m$ zero matrix. Consequently, no iterative search is needed. The subsequent discussion assumes $\|X\|_{\text{op}} > 0$.

The update $\widetilde{\Theta}^{(k)} = \widehat{\Theta}^{(k-1)} + \Delta$ by $\Delta \in \mathbb{S}^m$ in the first step derives from

$$\begin{aligned} (f \circ X)(\widehat{\Theta}^{(k-1)} + \Delta) &= \|Y - X(\widehat{\Theta}^{(k-1)} + \Delta)\|^2/(2\bar{\mu}) \\ &= (f \circ X)(\widehat{\Theta}^{(k-1)}) + \langle -\widetilde{G}^{(k)}, \Delta \rangle + \|X\Delta\|^2/(2\bar{\mu}), \end{aligned} \quad <3.5>$$

wherein $\widetilde{G}^{(k)} = -(\langle X, \widetilde{D}^{(k)} \rangle + \langle \widetilde{D}^{(k)}, X \rangle)/2$ and $\widetilde{D}^{(k)} = (X\widehat{\Theta}^{(k-1)} - Y)/\bar{\mu}$ resemble <3.4>. The equality $\|X\Delta\|^2 \leq \|X\|_{\text{op}}^2 \|\Delta\|^2$ identifies $-\widetilde{G}^{(k)}$ as the gradient $\nabla(f \circ X)(\widehat{\Theta}^{(k-1)})$ of $f \circ X$ at $\widehat{\Theta}^{(k-1)}$. Convexity of $f \circ X$ ensures $\partial(f \circ X)(\widehat{\Theta}^{(k-1)}) = \{-\widetilde{G}^{(k)}\}$. The resulting inequality $(f \circ X)(\widehat{\Theta}^{(k-1)} + \Delta) \geq (f \circ X)(\widehat{\Theta}^{(k-1)}) - \langle \widetilde{G}^{(k)}, \Delta \rangle$ reveals that an update of the form $\alpha\widetilde{G}^{(k)}$, $\alpha > 0$, allows for the possibility of a decrease of $f \circ X$. In fact, a suitable choice of $\alpha > 0$ ensures that an update of the form $\widetilde{\Theta}^{(k)} = \widehat{\Theta}^{(k-1)} + \Delta = \widehat{\Theta}^{(k-1)} + \alpha\widetilde{G}^{(k)}$ accords with the above objective. More specifically, <3.5> implies $(f \circ X)(\widehat{\Theta}^{(k-1)} + \alpha\widetilde{G}^{(k)}) \leq (f \circ X)(\widehat{\Theta}^{(k-1)}) + \alpha[\alpha\|X\|_{\text{op}}^2/(2\bar{\mu}) - 1]\|\widetilde{G}^{(k)}\|^2$, and therefore $(f \circ X)(\widetilde{\Theta}^{(k)}) \leq (f \circ X)(\widehat{\Theta}^{(k-1)})$ whenever $0 < \alpha \leq 2\bar{\mu}/\|X\|_{\text{op}}^2$.

The second step considers the special case $l_\lambda^{(k)}$ of the criterion l_λ in <3.1> with $W = \mathbb{R}^m$, X being the canonical embedding of \mathbb{S}^m into $\mathbb{R}^{m \times m}$, $Y = \widetilde{\Theta}^{(k)} \in \mathbb{S}^m \subset \mathbb{R}^{m \times m}$, and $\bar{\mu} = \alpha$. That is, $l_\lambda^{(k)}(\Theta) = \|\widetilde{\Theta}^{(k)} - \Theta\|^2/(2\alpha) + \lambda\|\Theta\|_{\text{nuc}}$. The unique $\widehat{\Theta}^{(k)} \in \mathbb{S}^m$ minimizing $l_\lambda^{(k)}$ exhibits a traceable form. In fact, the corresponding dual solution $\widehat{D}^{(k)} = (\widehat{\Theta}^{(k)} - \widetilde{\Theta}^{(k)})/\alpha$ is symmetric and therefore $\widehat{G}^{(k)} = -\widehat{D}^{(k)}$. By proposition 3.6 and <3.4>, one of the following two cases applies. If $\|\widetilde{\Theta}^{(k)}\|_{\text{op}} \leq \alpha\lambda$, then $\widehat{\Theta}^{(k)} = 0$. Otherwise, the unique minimizer is characterized by $\|\widetilde{\Theta}^{(k)} - \widehat{\Theta}^{(k)}\|_{\text{op}} = \alpha\lambda$, that is, $\widehat{G}^{(k)}$ is a boundary point of $\{\|\cdot\|_{\text{op}} \leq \lambda\}$, and $\langle \widehat{\Theta}^{(k)}, \widetilde{\Theta}^{(k)} - \widehat{\Theta}^{(k)} \rangle = \alpha\lambda\|\widehat{\Theta}^{(k)}\|_{\text{nuc}}$, that is, $\widehat{\Theta}^{(k)}$ lies in the normal cone of $\{\|\cdot\|_{\text{op}} \leq \lambda\}$ at $\widehat{G}^{(k)}$. If $\sum_{i \leq \text{rk } \widetilde{\Theta}^{(k)}} s_i \sigma_i v_i \langle v_i, \cdot \rangle$, $s_i \in \{\pm 1\}$, denotes a spectral decomposition of $\widetilde{\Theta}^{(k)}$, then $\bar{\Theta} = \sum_{i \leq \text{rk } \widetilde{\Theta}^{(k)}} s_i \max\{\sigma_i - \alpha\lambda, 0\} v_i \langle v_i, \cdot \rangle$ satisfies

$$\begin{aligned} \|\widetilde{\Theta}^{(k)} - \bar{\Theta}\|_{\text{op}} &= \left\| \sum_{i \leq \text{rk } \widetilde{\Theta}^{(k)}} s_i [\sigma_i - \max\{\sigma_i - \alpha\lambda, 0\}] v_i \langle v_i, \cdot \rangle \right\|_{\text{op}} = \alpha\lambda \quad \text{and} \\ \langle \bar{\Theta}, \widetilde{\Theta}^{(k)} - \bar{\Theta} \rangle &= \sum_{i, j \leq \text{rk } \widetilde{\Theta}^{(k)}} s_i s_j \max\{\sigma_j - \alpha\lambda, 0\} [\sigma_i - \max\{\sigma_i - \alpha\lambda, 0\}] \langle v_i, v_j \rangle^2 \\ &= \alpha\lambda \sum_{i \leq \text{rk } \widetilde{\Theta}^{(k)}} \max\{\sigma_i - \alpha\lambda, 0\} = \alpha\lambda \|\bar{\Theta}\|_{\text{nuc}} \end{aligned}$$

and therefore $\widehat{\Theta}^{(k)} = \bar{\Theta}$. The latter equality is also true if $\|\widetilde{\Theta}^{(k)}\|_{\text{op}} \leq \alpha\lambda$.

The (optimality) properties of both steps ensure that any $\Theta \in \mathbb{S}^m$ satisfies

$$\begin{aligned}
l_\lambda(\Theta) &\geq [(f \circ X)(\widehat{\Theta}^{(k-1)}) + \langle -\widetilde{G}^{(k)}, \Theta - \widehat{\Theta}^{(k-1)} \rangle] + [g(\widehat{\Theta}^{(k)}) + \langle \widehat{G}^{(k)}, \Theta - \widehat{\Theta}^{(k)} \rangle] \\
&= l_\lambda(\widehat{\Theta}^{(k)}) - \frac{\|X(\widehat{\Theta}^{(k)} - \widehat{\Theta}^{(k-1)})\|^2}{2\bar{\mu}} + \langle -\widetilde{G}^{(k)}, \Theta - \widehat{\Theta}^{(k)} \rangle + \langle \widehat{G}^{(k)}, \Theta - \widehat{\Theta}^{(k)} \rangle \\
&\geq l_\lambda(\widehat{\Theta}^{(k)}) - \frac{\|X\|_{\text{op}}^2}{2\bar{\mu}} \|\widehat{\Theta}^{(k)} - \widehat{\Theta}^{(k-1)}\|^2 + \frac{1}{\alpha} \langle \widehat{\Theta}^{(k-1)} - \widehat{\Theta}^{(k)}, \Theta - \widehat{\Theta}^{(k)} \rangle \quad \langle 3.6 \rangle
\end{aligned}$$

wherein the first equality amounts to an application of [<3.5>](#) with $\Delta = \widehat{\Theta}^{(k)} - \widehat{\Theta}^{(k-1)}$. The second inequality utilizes $\widetilde{\Theta}^{(k)} = \widehat{\Theta}^{(k-1)} + \alpha\widetilde{G}^{(k)}$ and $\widehat{G}^{(k)} = (\widetilde{\Theta}^{(k)} - \widehat{\Theta}^{(k)})/\alpha$, which is due to the choice of the scaling parameter in second part. Using [<3.6>](#) together with the choice $\Theta = \widehat{\Theta}^{(k-1)}$ implies $l_\lambda(\widehat{\Theta}^{(k)}) \leq l_\lambda(\widehat{\Theta}^{(k-1)})$ whenever $0 < \alpha \leq 2\bar{\mu}/\|X\|_{\text{op}}^2$.

The recipe [<3.7>](#) considers the choice $\alpha = \bar{\mu}/\|X\|_{\text{op}}^2$ and starts at the zero matrix. The k -th iteration calculates the primal criterion $p_{k-1} = l_\lambda(\widehat{\Theta}^{(k-1)})$ at its starting point $\widehat{\Theta}^{(k-1)}$ as well as a corresponding dual value d_{k-1} . If $\widetilde{D}^{(k)}$ lies in $\{\|\langle X, \cdot \rangle + \langle \cdot, X \rangle\|_{\text{op}} \leq 2\lambda\}$, then d_{k-1} equals the dual objective $-f_\star$ in [<3.3>](#) at $\widetilde{D}^{(k)}$. Otherwise, $-f_\star$ is evaluated at the surrogate $\lambda\widetilde{D}^{(k)}/\|\widetilde{G}^{(k)}\|_{\text{op}}$. This choice implies $0 \leq l_\lambda(\widehat{\Theta}^{(k-1)}) - l_\lambda(\widehat{\Theta}) \leq p_{k-1} - d_{k-1}$, wherein $\widehat{\Theta}$ denotes the unique minimizer of l_λ . If $p_{k-1} - d_{k-1} < \varepsilon$ holds, then no improvement beyond ε is possible, and the search terminates. Otherwise, the next search point $\widehat{\Theta}^{(k)}$ derives from a spectral decomposition $\sum_{i \leq \text{rk} \widetilde{\Theta}^{(k)}} s_i \sigma_i v_i \langle v_i, \cdot \rangle$ of $\widetilde{\Theta}^{(k)}$.

$$\begin{array}{l}
1 \quad \widehat{\Theta}^{(0)} = 0, \quad \alpha = \bar{\mu}/\|X\|_{\text{op}}^2, \quad \eta_0 = 0 \\
2 \quad \text{for } k \geq 1 \\
3 \quad \quad \widetilde{D}^{(k)} = (X\widehat{\Theta}^{(k-1)} - Y)/\bar{\mu} \\
4 \quad \quad \widetilde{G}^{(k)} = -(\langle X, \widetilde{D}^{(k)} \rangle + \langle \widetilde{D}^{(k)}, X \rangle)/2 \\
5 \quad \quad p_{k-1} = \bar{\mu}\|\widetilde{D}^{(k)}\|^2/2 + \lambda\eta_{k-1} \\
6 \quad \quad \text{if } \|\widetilde{G}^{(k)}\|_{\text{op}} > \lambda \\
7 \quad \quad \quad d_{k-1} = -f_\star(\lambda\widetilde{D}^{(k)}/\|\widetilde{G}^{(k)}\|_{\text{op}}) \\
8 \quad \quad \text{else} \\
9 \quad \quad \quad d_{k-1} = -f_\star(\widetilde{D}^{(k)}) \\
10 \quad \quad \text{if } p_{k-1} - d_{k-1} < \varepsilon \\
11 \quad \quad \quad \text{break} \\
12 \quad \quad \widetilde{\Theta}^{(k)} = \widehat{\Theta}^{(k-1)} + \alpha\widetilde{G}^{(k)} = \sum_{i \leq \text{rk} \widetilde{\Theta}^{(k)}} s_i \sigma_i v_i \langle v_i, \cdot \rangle \\
13 \quad \quad \widehat{\Theta}^{(k)} = \sum_{i \leq \text{rk} \widetilde{\Theta}^{(k)}} s_i \max\{\sigma_i - \alpha\lambda, 0\} v_i \langle v_i, \cdot \rangle \\
14 \quad \quad \eta_k = \sum_{i \leq \text{rk} \widetilde{\Theta}^{(k)}} \max\{\sigma_i - \alpha\lambda, 0\}
\end{array} \quad \langle 3.7 \rangle$$

In case $\|\langle X, Y \rangle + \langle Y, X \rangle\|_{\text{op}} \leq 2\bar{\mu}\lambda$, that is, $\widehat{\Theta} = 0$, then [<3.7>](#) terminates during the first iteration. In general, proposition [3.7](#) guarantees that the criterion in line 10 of [<3.7>](#)

is met after finitely many iterations. Its proof starts on page 81 in appendix 3.b.

Proposition 3.7. *Let $(\widehat{\Theta}^{(k)})_{k \in \mathbb{N}}$ denote a sequence generated by <3.7>. Then*

$$l_\lambda(\widehat{\Theta}^{(k)}) - l_\lambda(\widehat{\Theta}) \leq \frac{\|X\|_{\text{op}}^2 \|\widehat{\Theta}\|^2}{2\bar{\mu}k}.$$

Hence, $\widehat{\Theta}^{(k)} \rightarrow \widehat{\Theta}$ as $k \rightarrow \infty$ and $p_k - d_k < \varepsilon$ is eventually met for any $\varepsilon > 0$.

Quantitative statements on the speed of convergence of $\widehat{\Theta}^{(k)}$ to $\widehat{\Theta}$ require some information on the curvature of l_λ such as strong convexity of its first summand.

3.4. A poor man's factor model

3.4.1. Temporal dependence

This section considers the span W of a finite sequence of \mathbb{P} -square integrable random variables $v_{t,j}$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and with the index (t, j) ranging over a subset I_v of $\mathbb{N} \times \mathbb{N}$. The inner product $\langle \cdot, \cdot \rangle$ is given by the \mathbb{P} -expectation $\mathbb{E}xy = \int x(\omega)y(\omega)\mathbb{P}(d\omega) = \langle x, y \rangle$ of the product xy —as in example (e) of sections 2.1 and 2.4.1— and equips this linear space with a Euclidean geometry such that $v_{t,j}, (t, j) \in I_v$, form an orthonormal basis of W . Additional random variables $x_{t,j}$ with (t, j) ranging over a subset I_x of $\mathbb{N} \times \mathbb{N}$ result as linear combinations of the basis elements $v_{t,j}, (t, j) \in I_v$. This setup allows a formal representation as at the end of appendix 2.a. For now, the vector $\alpha \in \mathbb{R}^k, k \in \mathbb{N}$, gathers all coordinates of $x_{t,j}, (t, j) \in I_x$, with respect to $v_{t,j}, (t, j) \in I_v$. Section 3.5 focuses on the task of estimating a transformation $\Theta_* \in \mathbb{S}^m$ of α using a single realization—the data— $x_{t,j}(\omega), (t, j) \in I_x$, and knowledge of the overall structure including that α lies in a subset $\mathcal{M} \subset \mathbb{R}^k$. In this framework, a successful (relative to \mathcal{M}) estimation strategy approximately recovers the respective transformation $\Theta_*(\alpha')$ from data generated using $\alpha' \in \mathcal{M}$ irrespective of the particular value of $\alpha' \in \mathcal{M}$.

The space W is spanned by random variables $v_{t,j}$. Herein, the first index t ranges from $1-l$ to n for $j \leq h$ and from one to n for $h+1 \leq j \leq m$ for some $m, n \in \mathbb{N}, l \geq 0$, and $0 \leq h \leq m$. These random variables are independent and with zero mean $\mathbb{E}v_{t,j} = 0$ as well as $\mathbb{E}v_{t,j}^2 = 1$, thus, form an orthonormal basis of W . The data used in the following sections equal one realization of the columns of $X_t = [x_{t,1} \ \cdots \ x_{t,m}]$ given by

$$\begin{aligned} X_t &= F_t U_1^\top + \rho V_{t,2} U_2^\top \\ F_t &= [f_{t,1} \ \cdots \ f_{t,h}] = V_{t,1} A_0 + \sum_{i \leq l} V_{t-i,1} A_i \quad , \quad 1 \leq t \leq n. \end{aligned} \quad <3.8>$$

Herein, $V_{t,1} = [v_{t,1} \ \cdots \ v_{t,h}]$, $1-l \leq t \leq n$, and $V_{t,2} = [v_{t,h+1} \ \cdots \ v_{t,m}]$, $1 \leq t \leq n$. In addition, $A_0, \dots, A_l \in \mathbb{R}^{h \times h}$ are diagonal matrices, $\rho > 0$, and the columns of $U = [U_1 \ U_2] = [u_1 \ \dots \ u_h \ u_{h+1} \ \dots \ u_m]$ form an orthonormal basis of \mathbb{R}^m . If $h > 0$, then all diagonal entries of A_0 as well as at least one diagonal entry of A_l are nonzero. The former requirement guarantees that $x_{t,j}$ are linearly independent; the latter gives meaning to l . Finally, if $h = 0$, then all quantities related to the first summand of X_t and, in particular,

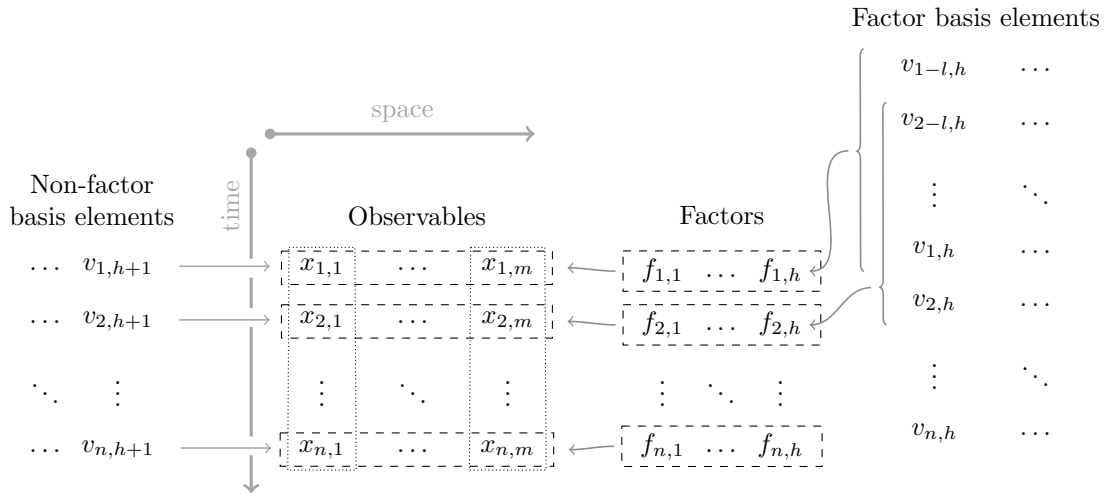


Figure 3.7

The figure illustrates the construction of $x_{t,1}, \dots, x_{t,m}$ in <3.8> as linear combinations of the basis elements $v_{t,h+1}, \dots, v_{t,m}$ and the factors $f_{t,1}, \dots, f_{t,h}$. Herein, the case $h = 0$ (no factors) is allowed, but the figure concerns the case $1 < h < m$. The factors equal linear combinations of a “rolling window” of the basis elements $v_{t,j}$, $j \leq h$, with identical second index j . Dashed lines surround the columns of X_t and F_t , respectively.

the second equation of the specification <3.8> disappear. Then, $x_{t,1}, \dots, x_{t,m}$ are even mutually orthogonal. Below, this extreme case usually receives—but also requires—no explicit mention in order to simplify the exposition. The same applies to the case $h = m$ which eliminates all quantities related to second summand of X_t such as U_2 and ρ .

The random variables $x_{t,j}$, $(t, j) \in I_x$, represent a (numerical) characteristic—referred to as x —of m spatial entities at n points in time. In particular, the first index t indicates the respective time point; the second index j points to the location in space. This interpretation suggests calling the subspaces $\text{img } X_{t-1}$ and $\text{img } [X_{t-1} \dots X_1]$ the recent past and the past of x at $t > 1$, respectively. Elements of the innovation space $\text{span}\{v_{t,1}, \dots, v_{t,m}\}$ at t lie in the orthogonal complement of the past of x at t . Their part in $\text{span}\{v_{t,h+1}, \dots, v_{t,m}\}$ exerts only momentary influence. In contrast, $v_{t,1}, \dots, v_{t,h}$ enter in the construction of the *factors* $f_{t,1}, \dots, f_{t,h}$ at t and thereby impact the columns of X_{t+1}, \dots, X_n . These factors $f_{t,1} = X_t u_1, \dots, f_{t,h} = X_t u_h$ lie in $\text{img } X_t$ by virtue of the (pairwise) orthogonality of the columns of $U = [U_1 \ U_2] = [u_1 \ \dots \ u_h \ u_{h+1} \ \dots \ u_m]$.

Each factor sequence $f_{1,j}, \dots, f_{n,j}$, $j \leq h$, embodies one of a small number— h is thought to be “much smaller” than m —of underlying determinants of x . The elements $f_{1,j}, \dots, f_{n,j}$ of the j -th factor sequence equal linear combinations of overlapping subsets of the basis elements $v_{1-l,j}, \dots, v_{n,j}$. Thus, the factor variables $f_{t,j}$ are generally independent across $j \leq h$ but dependent across the time index $t \leq n$ unless $l = 0$. Figure 3.7 contains a visual summary of the construction in <3.8> for the case $1 < h < m$, and, in particular, highlights the overlap of the subsets of basis elements $v_{1-l,h}, \dots, v_{n,h}$ needed to construct the members of the h -th factor sequence $f_{1,h}, \dots, f_{n,h}$.

The coefficient matrices U_1 and U_2 govern the dependence among the columns of X_t

and are discussed in more detail in section 3.4.2. The equality $A_0 = \rho I$ generates a notable special case. Herein, $I = [e_1 \cdots e_h]$ denotes the $h \times h$ identity matrix, and (thus) e_1, \dots, e_h symbolizes the standard basis of \mathbb{R}^h . Then, the specification <3.8> becomes

$$X_t = \left[\sum_{i \leq l} V_{t-i,1} A_i \right] U_1^\top + \rho [V_{t,1} \ V_{t,2}] \begin{pmatrix} U_1^\top \\ U_2^\top \end{pmatrix}, \quad 1 \leq t \leq n, \quad \langle 3.9 \rangle$$

wherein the first term disappears if $l = 0$. Moreover, the columns of the final term, which equals a scaled composition of unitary maps, amount to m pairwise orthogonal elements of the innovation space (at t) of length ρ . These columns represent idiosyncratic innovations to the individual $x_{t,j}$. In particular, U_1 controls the entire spatial—across j —Euclidean space dependence between the observables $x_{t,1}, \dots, x_{t,m}$.

Many properties of the setting in <3.8> are reflected by the implied (unordered) spectral decompositions of the symmetric inner product matrices $\langle X_t, X_{t-s} \rangle$ given by

$$\langle X_t, X_{t-s} \rangle = \begin{cases} U \begin{pmatrix} \sum_{i=0}^l A_i^2 & \\ & \rho^2 I \end{pmatrix} U^\top & , \quad s = 0, \\ U_1 \left(\sum_{i=0}^{l-s} A_i A_{i+s} \right) U_1^\top & , \quad 0 < s \leq \min\{l, t-1\}, \\ m \times m \text{ zero matrix} & , \quad s > \min\{l, t-1\}. \end{cases} \quad \langle 3.10 \rangle$$

Firstly, time invariance of the coefficients in <3.8> ensures the absence of t on the right-hand side of <3.10>. Secondly, the inner product matrices $\langle X_t, X_{t-s} \rangle$ are symmetric due to the specific separation of time and space dependence. Thirdly, if $s > 0$, then $\text{rk} \langle X_t, X_{t-s} \rangle$ does not exceed the number of factors h , which provide the sole link across t .

These properties become evident when projecting the elements $x_{t,1}, \dots, x_{t,m}$, $1 < t \leq n$, onto the recent past $\text{img} X_{t-1}$ of x . In fact, the coordinate matrix Θ_* with respect to X_{t-1} of the composition $P_{\text{img} X_{t-1}} X_t = X_{t-1} \Theta_*$ coincides for all $t \geq 2$. It is uniquely determined by the condition $\langle X_{t-1}, X_{t-1} \rangle \Theta_* = \langle X_{t-1}, X_t \rangle$, thus, equals

$$\Theta_* = U_1 \Gamma_* U_1^\top = U_1 \left[\left(\sum_{i=0}^l A_i^2 \right)^{-1} \sum_{i=0}^{l-1} A_i A_{i+1} \right] U_1^\top, \quad \langle 3.11 \rangle$$

wherein the superscript $^{-1}$ marks the inverse of (the bijective linear map) $\sum_{i=0}^l A_i^2$. The (bracketed) diagonal matrix $\Gamma_* \in \mathbb{R}^{h \times h}$ provides the coordinates in $P_{\text{img} F_{t-1}} F_t = F_{t-1} \Gamma_*$. If either $h = 0$ or $h > 0$ together with $l = 0$, then the Θ_* equals the $m \times m$ zero matrix.

If $h \geq 1$ and $l \geq 1$, then these considerations lead to the alternative representation

$$\begin{aligned} X_t &= X_{t-1} \Theta_* + [H_t + R_t] = X_{t-1} \Theta_* + \bar{E}_t, & \langle 3.12 \rangle \\ H_t &= \left(\sum_{i \leq l} V_{t-i,1} (A_i - A_{i-1} \Gamma_*) - V_{t-l-1,1} A_l \Gamma_* \right) U_1^\top, \quad 2 \leq t \leq n, \\ R_t &= [V_{t,1} \ V_{t,2}] \begin{pmatrix} A_0 & \\ & \rho I \end{pmatrix} \begin{pmatrix} U_1^\top \\ U_2^\top \end{pmatrix}. \end{aligned}$$

The inner product matrix $\langle X_{t-s}, R_t \rangle$ has all its entries equal to zero if $s \geq 1$. If $l > 0$,

then the same applies to $\langle X_{t-s}, H_t \rangle$ for $s = 1$ but generally fails for $t \geq 3$ and $2 \leq s \leq \min\{l + 1, t - 1\}$ as elements of $\text{img } H_t$ are not contained in the innovation space at t . However, if $A_0 = \rho I$, $A_i = \rho D^i$, $1 \leq i \leq l$, for some diagonal matrix $D \in \mathbb{R}^{h \times h}$ with diagonal entries $|d_{i,i}| < 1$, then elements of $\text{img } \bar{E}_t$ approach the subspace $\text{img } R_t$ of the innovation space as $l \rightarrow \infty$. More specifically, one may consider a sequence of Euclidean spaces of the above type—indexed by $k \in \mathbb{N}$ —such that $l = l_k$ increases in parallel with the sequence index k . No further definition is required as m is shared across these spaces, and $A_i = \rho D^i$ is valid for all $i \in \mathbb{N}$. Then all of these spaces come with a measure of distance $\sup_{x \in \text{img } \bar{E}_t \cap \{\|\bullet\|=1\}} \|P_{(\text{img } R_t)^\perp} x\|$, and the sequence of these distances approaches zero. Moreover, this case features the equality $A_0 = \rho I$, thus, is a special case of <3.9> and therefore exhibits $\langle R_t, R_t \rangle = \rho^2 I$. In the above “asymptotic” sense, the symmetric matrix Θ_* controls the transition from the recent past to the present and is therefore called the *transition matrix*. If $l = 0$, then $X_t = R_t$. Thus, the transition matrix Θ_* is zero, and these considerations are meaningless.

*transition
matrix*

3.4.2. Spatial dependence

Figure 3.7 proposes two views on the observables: firstly, as m *time series* $x_{1,j}, \dots, x_{n,j}$ (dotted lines), that is, sequences of random variables indexed by time, and, secondly, as n *random fields* $x_{t,1}, \dots, x_{t,m}$ (dashed lines)—sequences of random variables indexed by space. From a constructional point of view, the presentation in <3.8> stresses the first of these interpretations: the observable time series result as linear combinations of the factor time series $f_{1,j}, \dots, f_{n,j}$, $j \leq h$, and the non-factor time series $v_{1,j}, \dots, v_{n,j}$, $h + 1 \leq j \leq m$. The random vectors $x_t = (x_{t,1}, \dots, x_{t,m})$, $t \leq n$, facilitate a presentation stressing the second interpretation. More specifically, expressing the relations <3.8> in terms of these random vectors and the similarly defined random vectors $f_t = (f_{t,1}, \dots, f_{t,h})$, $v_t^{(1)} = (v_{t,1}, \dots, v_{t,h})$, and $v_t^{(2)} = (v_{t,h+1}, \dots, v_{t,m})$ leads to

*time series
random fields*

$$\begin{aligned} x_t &= U_1 f_t + \rho U_2 v_t^{(2)} \\ f_t &= A_0 v_t^{(1)} + \sum_{i \leq l} A_i v_{t-i}^{(1)} \quad , \quad t \leq n \quad , \end{aligned} \tag{3.13}$$

wherein the second summand of the second equation is present only if $l > 0$. In particular, the formulation in <3.13> emphasizes that realizations $x_t(\omega) \in \mathbb{R}^m$, given by $(x_{t,1}(\omega), \dots, x_{t,m}(\omega))$, of the random vectors x_t consist of two mutually orthogonal parts. The first part $U_1 f_t(\omega) = \sum_{j \leq h} f_{t,j}(\omega) u_j$ reflects the influence of the factors. The second part $\rho U_2 v_t^{(2)}(\omega) = \sum_{j=h+1}^m \rho v_{t,j}(\omega) u_j$ captures deviations associated with the specific time point t . In particular, the columns u_1, \dots, u_h of U_1 may be understood as h “spatial patterns” whose strengths at time t is determined by $f_{t,1}, \dots, f_{t,h}$, respectively.

These patterns u_1, \dots, u_h amount to functions—as explained in example (a) of section 2.1.1—on the space index set $\{1, \dots, m\}$. Herein, some form of smoothness of the “spatial patterns” u_j is expected. Squared difference quotients of the form $(u_j(i') - u_j(i))^2 / \text{dist}(i', i)^2 = w_{i',i} (u_{i',j} - u_{i,j})^2$, $i' \neq i$, measure their roughness, wherein $\text{dist}(i', i) = \text{dist}(i, i')$ and $w_{i',i} \geq 0$ denote a symmetric notion of distance between locations i' and i

and the square of its reciprocal, respectively. The subsequent discussion refers to dist only through $w_{i,i'} = w_{i',i}$, $i \neq i'$. In fact, the role of dist is to facilitate the interpretation, and using $w_{i,i'} = 0$ to represent “infinite distance” introduces no technical complications.

If one sets $w_{i,i} = 0$ for all $i \leq m$, then the integral of the difference quotients corresponding to a fixed $j \leq h$ with respect to the product (counting) measure on $\{1, \dots, m\} \times \{1, \dots, m\}$ may be expressed in the form $\sum_{i',i \leq m} w_{i',i} (u_{i',j} - u_{i,j})^2 = 2\langle u_j, \Lambda u_j \rangle$, wherein the matrix Λ is defined in the following display. This equality implies that

$$\Lambda = \begin{pmatrix} \sum_{i' \leq m} w_{i',1} & & & \\ & \ddots & & \\ & & \ddots & \\ & & & \sum_{i' \leq m} w_{i',m} \end{pmatrix} - \begin{pmatrix} 0 & w_{1,2} & \dots & w_{1,m} \\ w_{1,2} & 0 & \dots & w_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ w_{1,m} & w_{2,m} & \dots & 0 \end{pmatrix} \quad \langle 3.14 \rangle$$

is positive semidefinite and is subsequently assumed to be nonzero, that is, at least one pair $i, j \leq m$ exhibits finite distance. The form of Λ implies $(1, \dots, 1) \in \ker \Lambda$, which fits the role of $u \mapsto \langle u, \Lambda u \rangle$ as a measure of roughness and reveals $\text{rk } \Lambda < m$. More precisely, one has $\text{rk } \Lambda = \inf \{m - k \mid \text{there exists a partition } C_1, \dots, C_k \text{ of } \{1, \dots, m\} \text{ with } i \in C_s \not\equiv i' \Rightarrow w_{i,i'} = 0\}$. In fact, the infimum $m - k_*$ is attained due to the well-ordering principle. If C_1, \dots, C_{k_*} form a corresponding partition, $a_j = \sum_{i \in C_j} e_i$, $j \leq k_*$, and R provides a Cholesky factor of Λ , then $\|Ra_j\|^2 = \langle a_j, \Lambda a_j \rangle = \frac{1}{2} \sum_{i,i' \leq m} w_{i,i'} (a_{i,j} - a_{i',j})^2 = 0$ as $i \in C_s, i' \in C_t$ with either $s = t$ and therefore $a_{i,j} = a_{i',j}$ or $s \neq t$ and therefore $w_{i,i'} = 0$. Conversely, if $a \in \mathbb{R}^m$ exhibits entries $a_i \neq a_{i'}$ with $w_{i,i'} \neq 0$, then $\langle a, \Lambda a \rangle > 0$.

Due to its symmetry, Λ exhibits a spectral decomposition $\Lambda = \sum_{i \leq \text{rk } \Lambda} \sigma_i(\Lambda) o_i \langle o_i, \bullet \rangle$, wherein $o_1, \dots, o_{\text{rk } \Lambda}$ represents an orthonormal sequence of singular vectors of the form given in lemma 2.4. In this notation, the suggested measure of roughness of u_j equals $\|\Lambda^{1/2} u_j\|^2$, wherein $\Lambda^{1/2} = \sum_{i \leq \text{rk } \Lambda} \sigma_i^{1/2}(\Lambda) o_i \langle o_i, \bullet \rangle$ does not depend on the particular choice of singular vectors. The same applies to the alternative roughness matrix $\Lambda^{q/2} = \sum_{i \leq \text{rk } \Lambda} \sigma_i^{q/2}(\Lambda) o_i \langle o_i, \bullet \rangle$, wherein $q > 0$ allows adjustment of the weights $\sigma_i^{q/2}(\Lambda)$ for a given distance. More specifically, $q < 1$ downplays differences in the singular values; $q > 1$ amplifies these differences. In addition, symmetry and $\text{img } \Lambda = \text{span}\{o_1, \dots, o_{\text{rk } \Lambda}\} = \text{img } \Lambda^{q/2}$ ensure that $\ker \Lambda = \ker \Lambda^{q/2}$ and $\text{rk } \Lambda^{q/2} < m$, for all $q > 1$. The sum $\|\Lambda^{q/2} U_1\|^2 = \sum_{j \leq h} \|\Lambda^{q/2} u_j\|^2$ measures the total roughness of the (spatial) patterns u_j . The alternative quantity $\|\Lambda^{q/2} \Theta_*\|^2$ amounts to a weighted sum—with weights equal to the squared diagonal entries of Γ_* —of the individual roughness terms $\|\Lambda^{q/2} u_j\|^2$.

Any valid choice for the above sequence $o_1, \dots, o_{\text{rk } \Lambda}$ of singular vectors for Λ can be extended to an orthonormal basis o_1, \dots, o_m of \mathbb{R}^m . If $\text{rk } \Lambda < m - 1$ or if $\dim \ker(\Lambda \pm \bar{\sigma}_j(\Lambda) \text{id}) > 1$ for some j , wherein $\bar{\sigma}_j(\Lambda)$ and id denote the j -th distinct singular value of Λ and the identity map on \mathbb{R}^m , respectively, then—according to section 2.5.4—the choice of singular vectors and $o_{\text{rk } \Lambda + 1}, \dots, o_m$ involves some ambiguity beyond sign choices. However, these arbitrary choices are practically immaterial to the subsequent discussion as they do not affect the key quantities derived from the chosen basis. Two observations are essential in this regard. Firstly, one has $\text{span}\{o_1, \dots, o_{\text{rk } \Lambda}\} = (\ker \Lambda)^\perp = \text{img } \Lambda$. Secondly, positive semidefiniteness of Λ implies $\ker(\Lambda + \bar{\sigma}_j(\Lambda) \text{id}) = \{0\}$ for all distinct

singular vectors. Hence, $L_k^\perp = \text{span}\{o_1, \dots, o_{k-1}\}$, $1 < k \leq \text{rk } \Lambda + 1$, is unequivocal whenever either $k = \text{rk } \Lambda + 1$ or $1 < k \leq \text{rk } \Lambda$ together with $\sigma_{k-1}(\Lambda) > \sigma_k(\Lambda)$.

Every orthonormal basis o_1, \dots, o_m of \mathbb{R}^m induces—comparable to e_i and $\bar{B}_{i,j}$ in examples (a) and (c) of section 2.1.1—an orthonormal basis $\bar{O}_{i,j}$, $i \leq j \leq m$, of \mathbb{S}^m , which is given by $\bar{O}_{i,i} = o_i o_i^\top$ and $\bar{O}_{i,j} = (o_i o_j^\top + o_j o_i^\top) / \sqrt{2}$ for $i < j$. In terms of the latter, a “small”—relative to the other parameters such as A_0, \dots, A_l , and ρ —value of $\|\Lambda^{1/2} \Theta_*\|^2$ corresponds to the transition matrix Θ_* being close to k -model space

k-model space

$$V_k = \text{span}\{\bar{O}_{i,j} \mid j \geq i \geq k\} = \{A \in \mathbb{S}^m \mid \text{img } A \subset L_k\}, \quad L_k = \text{span}\{o_k, \dots, o_m\},$$

for some “large” $k \in \mathbb{N}$. The latter is herein restricted to $k \leq \text{rk } \Lambda + 1 \leq m$ with $\sigma_{k-1}(\Lambda) > \sigma_k(\Lambda)$ if $1 < k \leq \text{rk } \Lambda$ to ensure an unambiguous definition. In general, the proximity of Θ_* to V_k may be expressed in terms of the residual length $\|P_{V_k^\perp} \Theta_*\|^2 = \|\Theta_* - \sum_{j \geq i \geq k} \langle \Theta_*, \bar{O}_{i,j} \rangle \bar{O}_{i,j}\|$, which should be “small” relative to $\|\Theta_*\|$.

3.5. Transition matrix estimation

3.5.1. Estimation strategy

This section considers a single realization $x_{t,j}(\omega)$, $t \leq n$, $j \leq m$, henceforth called the data, of corresponding random variables $x_{t,j}$. These random variables are linear combinations as in <3.8> of an orthonormal basis $v_{t,j}$, $(t, j) \in I_v$, consisting of \mathbb{P} -square integrable random variables defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Hence, the strategy of appendix 2.a may be used to build a formal model. Furthermore, the assertions of this and the next section concern only the linear relations of elements of $\text{span}\{v_{t,j} \mid (t, j) \in I_v\}$ and are therefore valid for all choices of the basis element representatives. The inner product has the form $\langle x, y \rangle = \mathbb{E}xy = \int x(\omega)y(\omega)\mathbb{P}(d\omega)$; however, this Euclidean space structure has no explicit role in the upcoming investigation. Therein, the coordinates A_0, \dots, A_l , ρ , and U_1, U_2 as well as $h, l \geq 0$ are considered unknown, but satisfy the restriction that all diagonal entries of A_0 as well as at least one diagonal entry of A_l are nonzero. The case $h = 0$ is possible, and then the factor related quantities disappear.

The ultimate goal is to estimate the transition matrix Θ_* in <3.11> using the above data. The estimate $\hat{\Theta}$ used here takes the form of the unique minimizer of

$$\begin{aligned} \mathbb{S}^m \ni \Theta &\mapsto \frac{1}{2(n-1)} \|Y - X\Theta\|^2 + \lambda \|\Theta\|_{\text{nuc}} + \xi \|\Lambda^{q/2} \Theta\|^2 \\ &= \frac{1}{2(n-1)} \left\| \begin{pmatrix} Y \\ \end{pmatrix} - \begin{pmatrix} X \\ \sqrt{2(n-1)\xi} \Lambda^{q/2} \end{pmatrix} \Theta \right\|^2 + \lambda \|\Theta\|_{\text{nuc}}, \end{aligned} \tag{3.15}$$

wherein $\lambda, \xi > 0$ control the relative importance of the second and third summand,

$$Y = \begin{pmatrix} x_{2,1}(\omega) & \dots & x_{2,m}(\omega) \\ \vdots & \ddots & \vdots \\ x_{n,1}(\omega) & \dots & x_{n,m}(\omega) \end{pmatrix}, \quad X = \begin{pmatrix} x_{1,1}(\omega) & \dots & x_{1,m}(\omega) \\ \vdots & \ddots & \vdots \\ x_{n-1,1}(\omega) & \dots & x_{n-1,m}(\omega) \end{pmatrix},$$

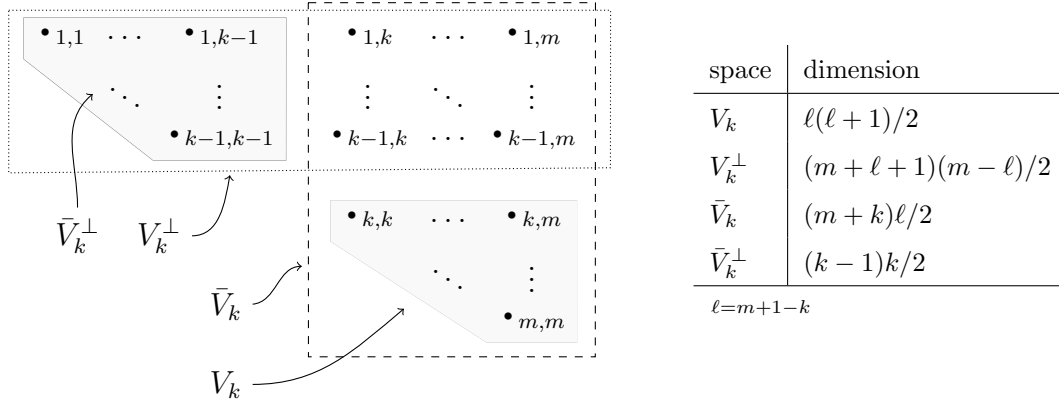


Figure 3.8

The figure shows an abstract set of coordinates $\bullet_{i,j}$ with respect to the orthonormal basis $\bar{O}_{i,j}$ of \mathbb{S}^m —defined in section 3.4.2—arranged in an upper triangular scheme and for $1 < k < m$. Solid gray lines and a gray background highlight coordinates associated with the k -model space V_k as well as those associated with the orthogonal complement \bar{V}_k^\perp of the extended model space \bar{V}_k . Coordinates associated with \bar{V}_k and V_k^\perp are encircled by dashed and dotted lines, respectively. The table on the righthand side lists the dimensions of the four subspaces of \mathbb{S}^m .

and $\Lambda^{q/2}$ is as explained below <3.14>, thus $1 \leq \text{rk } \Lambda < m$, for some given symmetric notion dist of distance on $\{1, \dots, m\}$ and $q > 0$. The final term in <3.15> identifies the present criterion as a special case of <3.1> in section 3.2.1. Thus, section 3.2.3 deals with the above uniqueness assertion. Section 3.3 shows that this strategy is practicable.

The connection of the objective function <3.15> with the modeling of the previous section is threefold. Firstly, the considerations surrounding <3.12> suggests that $X\Theta_*$ should—at least in special cases and then for all ω in a \mathbb{P} -large set—be a close substitute to Y . In fact, the t -th row of X carries a realization of the columns of X_t , while the t -th row of Y consists of the corresponding realization of the columns of X_{t+1} . Secondly, the number of factors h being “small” relative to m implies that the transition matrix Θ_* exhibits “low” rank. The second component $\lambda \|\cdot\|_{\text{nuc}}$ encourages this property for the estimate $\hat{\Theta}$. Lastly, section 3.4.2 shows that $\|\Lambda^{q/2} \cdot\|^2$ provides a measure of smoothness of singular vectors of a (symmetric) matrix—viewed as functions on $(\{1, \dots, m\}, \text{dist})$, which herein have the interpretation of basic “spatial patterns”. At a higher level, the objective function <3.15> amounts to the sum of a data-based term—the first summand—and a structure-based term consisting of the second and third summand.

Section 3.5.2 derives conditions on $X, Y \in \mathbb{R}^{n-1 \times m}$ as well as $\lambda, \xi > 0$, which ensure that $\|\Theta_* - \hat{\Theta}\|$ is “small”. The discussion is in terms of a specific data set, that is, point-wise with respect to ω . Section 3.5.3 shows that these conditions hold for all $\omega \in S \in \mathcal{F}$, wherein the probability $\mathbb{P}S$ is controlled by the number of time points n amongst others.

Section 3.4.2 observes that the structural assumptions—“low” rank and “smooth” singular vectors—on Θ_* roughly correspond to Θ_* being close to the k -model space

$$V_k = \{A \in \mathbb{S}^m \mid \text{img } A \subset L_k\} = \text{span}\{\bar{O}_{i,j} \mid j \geq i \geq k\}, \quad L_k = \text{span}\{o_k, \dots, o_m\},$$

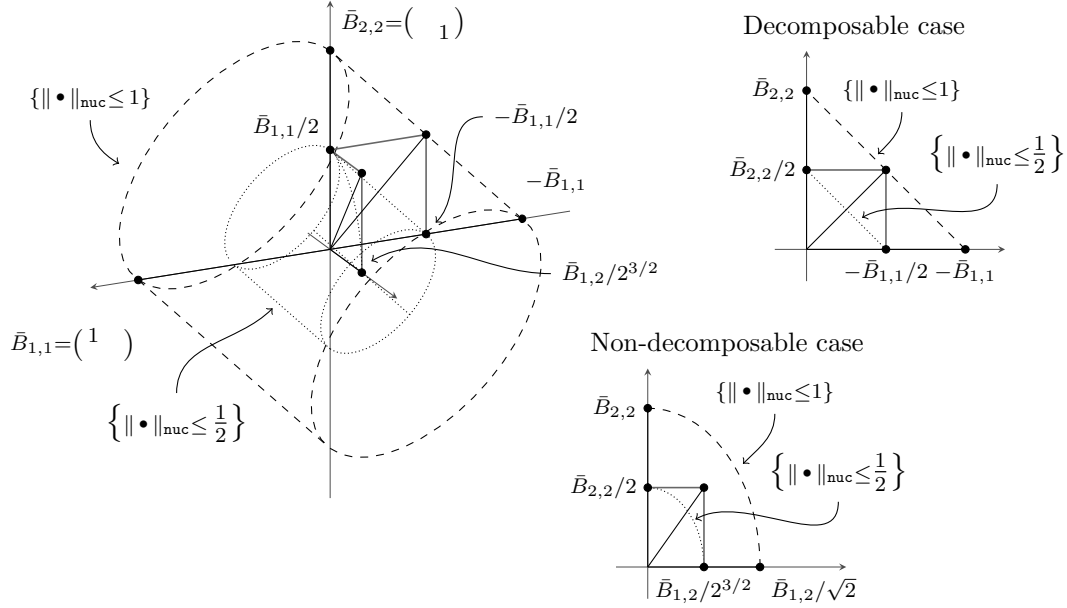


Figure 3.9

The figure visualizes the decomposability $\|\Theta_* + \Delta\|_{\text{nuc}} = \|\Theta_*\|_{\text{nuc}} + \|\Delta\|_{\text{nuc}}$ for elements $\Theta_* \in V_k$ and $\Delta \in \bar{V}_k^\perp$ with $m = 2$ as well as the possibility of non-decomposability in case $\Delta \in V_k^\perp$. Dashed and dotted lines indicate the $\|\cdot\|_{\text{nuc}}$ -unit- and $\|\cdot\|_{\text{nuc}} - \frac{1}{2}$ -ball, respectively. Here, $o_1 = e_1$, $o_2 = e_2$, and $k = 2$, thus, $V_k = \text{span}\{\bar{B}_{2,2}\}$ and $\bar{V}_k^\perp = \text{span}\{\bar{B}_{1,1}\}$, wherein e_i and $\bar{B}_{i,j}$ denote standard basis elements of \mathbb{R}^2 and \mathbb{S}^2 , respectively, as defined in section 2.1.1. The righthand part shows a part of the relevant two dimensional cross-sections of the lefthand side.

for “large” k , wherein the orthonormal basis o_1, \dots, o_m amounts to an extension of a singular vector sequence $o_1, \dots, o_{\text{rk}\Lambda}$ for Λ , and $\bar{O}_{i,j}$ equals $o_i o_i^\top$ if $i = j$ or $(o_i o_j^\top + o_j o_i^\top)/\sqrt{2}$ if not, respectively. To avoid ambiguity, the k -model space is defined only for $k \leq \text{rk}\Lambda + 1 \leq m$, wherein the second inequality follows from the definition of Λ which implies $\text{rk}\Lambda < m$, and with $\sigma_{k-1}(\Lambda) > \sigma_k(\Lambda)$ if $1 < k \leq \text{rk}\Lambda$. The same applies to the *extended k -model space* $\bar{V}_k = \text{span}\{\bar{O}_{i,j} \mid j \geq \max\{i, k\}\}$. Figure 3.8 arranges an abstract coordinate sequence $\bar{o}_{i,j}, i \leq j$, with respect to the orthonormal basis $\bar{O}_{i,j}, i \leq j$, in a triangular scheme and highlights the coordinates associated with the two types of model spaces V_k and \bar{V}_k as well as their orthogonal complements for the case $1 < k < m$; the neighboring table lists the dimensions of the four subspaces $V_k, V_k^\perp, \bar{V}_k$, and \bar{V}_k^\perp .

*extended
k-model space*

The extended k -model space further clarifies how <3.15> encourages $\hat{\Theta}$ to assume the structure expected in the (unknown) transition matrix Θ_* . More specifically, the minimization of this criterion function may be rephrased as the minimization of

$$\bar{l}_{\lambda,\xi}(\Delta) = \frac{1}{2(n-1)} \|\bar{E} - X\Delta\|^2 + \lambda \|\Theta_* + \Delta\|_{\text{nuc}} + \xi \|\Lambda^{q/2}(\Theta_* + \Delta)\|^2 \quad <3.16>$$

over $\Delta \in \mathbb{S}^m$, which represents the deviation $\Delta = \Theta - \Theta_*$ of Θ from Θ_* , and consequently $\bar{E} = Y - X\Theta_*$. This approach is practically infeasible as Θ_* and thereby \bar{E} are not available. However, it is helpful in the present—purely theoretical—discussion.

The orthogonal complement $\bar{V}_k^\perp = \text{span}\{\bar{O}_{i,j} \mid i \leq j < k\}$ of the extended k -model space gathers the directions which suffer the strongest opposition from the structure-based part in <3.16> if $\Theta_* \in V_k$. Then, for every $\Delta \in \bar{V}_k^\perp$ the inequality

$$\lambda\|\Theta_* + \Delta\|_{\text{nuc}} + \xi\|\Lambda^{q/2}(\Theta_* + \Delta)\|^2 \leq [\lambda\|\Theta_*\|_{\text{nuc}} + \xi\|\Lambda^{q/2}\Theta_*\|^2] + [\lambda\|\Delta\|_{\text{nuc}} + \xi\|\Lambda^{q/2}\Delta\|^2]$$

becomes an equality. Figure 3.9 visualizes this property for the first summand $\lambda\|\Theta_* + \Delta\|_{\text{nuc}}$ and the case $m = 2$, $\Theta_* \in \text{span}\{\bar{B}_{2,2}\}$, $\bar{B}_{2,2} = \begin{pmatrix} & \\ & 1 \end{pmatrix}$. It also highlights the connection with the facial structure of $\{\|\cdot\|_{\text{nuc}} \leq 1\}$ discussed in section 3.1.1. In addition, this figure exemplifies the possibility of strict inequality for $\Delta \in V_k^\perp$. Consequently, the consideration of $\bar{V}_k \supset V_k$ and thereby $\bar{V}_k^\perp \subset V_k^\perp$ is essential in this regard.

The relation between the orthonormal basis elements o_j and $\bar{O}_{i,j}$ translates into a relation between the orthogonal projector P_{L_k} with $L_k = \text{span}\{o_k, \dots, o_m\}$ and P_{V_k} as well as $P_{\bar{V}_k}$, respectively. More specifically, if $A \in \mathbb{S}^m$, then

$$A = (P_{L_k} + P_{L_k^\perp})A(P_{L_k} + P_{L_k^\perp}) = \underbrace{P_{L_k}AP_{L_k} + P_{L_k}AP_{L_k^\perp} + P_{L_k^\perp}AP_{L_k} + P_{L_k^\perp}AP_{L_k^\perp}}_{P_{\bar{V}_k}A} . \quad \text{<3.17>}$$

Therein, terms of the type $P_{L_k}AP_{L_k}$ embody compositions of linear maps $\mathbb{R}^m \rightarrow \mathbb{R}^m$. In contrast, terms of the type $P_{\bar{V}_k}A$ denote a projection in \mathbb{S}^m . Herein, the projection $P_{\bar{V}_k}A$ —considered as an element of $\mathbb{R}^{m \times m}$ —equals the sum $P_{L_k}A + P_{L_k^\perp}AP_{L_k}$ of two matrices of with rank no exceeding $\dim L_k = m + 1 - k = \ell$, thus, $\text{rk } P_{\bar{V}_k}A \leq 2\ell$.

3.5.2. Recovery conditions

This section derives an upper bound on the norm of the estimation error $\|\hat{\Theta} - \Theta_*\|$ in terms of a given realization $x_{t,j}(\omega)$, $t \leq n$, $j \leq m$, of corresponding random variables $x_{t,j}$ —the observables. Herein, the estimate $\hat{\Theta}$ equals the unique minimum of the criterion function <3.15> for given $\xi, \lambda > 0$, which are specified in the course of the analysis. Section 3.5.3 generalizes these bounds to hold for all ω in some $S \in \mathcal{F}$.

Section 3.5.1 justifies the presence of the second term $\lambda\|\cdot\|_{\text{nuc}}$ as well as the third term $\xi\|\Lambda^{q/2}\cdot\|^2$ in <3.15> by the vague idea of $\text{rk } \Theta_* (\leq h)$ and $\xi\|\Lambda^{q/2}\Theta_*\|^2$ being “small”. These conditions roughly translate to a low number h of factors and the singular vectors u_1, \dots, u_h of Θ_* being “smooth” functions on $(\{1, \dots, m\}, \text{dist})$, respectively. In this regard, the present section assumes the following two conditions. Firstly, the diagonal entries of the (diagonal) matrix $\sum_{j=0}^{l-1} A_j A_{j+1}$ are nonzero (if $h \neq 0$). Thus, one has $l \geq 1$, and the rank of $\Theta_* = U_1 \Gamma_* U_1^\top$ equals the number h of underlying factors. Moreover, this condition ensures that u_1, \dots, u_h are singular vectors of Θ_* of the form considered in lemma 2.4, that is, $u_i \in \ker(\Theta_* \pm \bar{\sigma}_j(\Theta_*) \text{id})$, $i \leq h$, with $\bar{\sigma}_j(\Theta_*) > 0$ being a distinct singular value of Θ_* . The possible ambiguity due to two or more diagonal entries of Γ_* being identical has no practical consequences for the present investigation as only Θ_* and $\hat{\Theta} - \Theta_*$ are of concern. Secondly, the columns u_1, \dots, u_h of U_1 and Θ_* are *perfectly aligned* with Λ , that is, either $h = \text{rk } \Theta_* = 0$ or there

perfectly aligned

exists $k \leq \text{rk } \Lambda + 1 (\leq m \text{ as } \text{rk } \Lambda < 0)$ with $\sigma_{k-1}(\Lambda) > \sigma_k(\Lambda)$ if $1 < k \leq \text{rk } \Lambda$ such that $\text{img } \Theta_* = \text{span}\{u_1, \dots, u_h\} = \text{span}\{o_k, \dots, o_m\} = L_k$. In the latter case, this requirement implies the inclusion $\Theta_* \in V_k = \text{span}\{\bar{O}_{i,j} \mid k \leq i \leq j\}$. Thus, the availability of the roughness matrix Λ amounts to a considerable understanding on u_1, \dots, u_h . In addition, if $\Theta_* \neq 0$, then the equalities $h = \text{rk } \Theta_* = m - k + 1 = \ell \geq 1$ hold, wherein the second equality provides the link to the notation used in section 3.5.1.

In terms of the alternative objective function $\bar{l}_{\lambda,\xi}$ in <3.16>, the definition of $\hat{\Theta}$ ensures that $\bar{l}_{\lambda,\xi}(\hat{\Delta}) \leq \bar{l}_{\lambda,\xi}(0)$, wherein $\hat{\Delta} = \hat{\Theta} - \Theta_*$ equals the estimation error and 0 symbolizes the $m \times m$ zero matrix. Rephrasing this inequality leads to the main result of this section (proposition 3.9). Lemma 3.8 summarizes the first part of its proof. The details of this derivation may be found on page 82 in appendix 3.b.

Lemma 3.8. *If Θ_* is perfectly aligned with Λ , $\text{rk } \Theta_* = h$, and $\lambda \geq \|G\|_{\text{op}}$ with $G = (X^\top \bar{E} + \bar{E}^\top X)/(2(n-1))$, $\bar{E} = Y - X\Theta_*$, then the minimizer $\hat{\Delta}$ of $\bar{l}_{\lambda,\xi}$ in <3.16> satisfies*

$$\left\| \frac{X\hat{\Delta}}{\sqrt{n-1}} \right\|^2 + \xi \frac{\sigma_{k-1}^q(\Lambda)}{2} \|P_{V_k^\perp} \hat{\Delta}\|^2 \leq 5\sqrt{h}\lambda \|\hat{\Delta}\| + 4\xi \|\Lambda^{q/2} \Theta_*\|^2, \quad \langle 3.18 \rangle$$

wherein $h = m + 1 - k$ and either $k \leq m$ with $V_k = \text{span}\{\bar{O}_{i,j} \mid k \leq i \leq j\}$ or $k = m + 1$ with $V_k = \{0\}$. In the latter case, one has $\sigma_{k-1}(\Lambda) = \sigma_m(\Lambda) = 0$ as $\text{rk } \Lambda < m - 1$.

The lower bound for λ in lemma 3.8 is a valid choice in the sense that it does not depend on the outcome $\hat{\Theta}$ of the optimization process; however, it cannot provide guidance in practical situations when Θ_* and thereby $\bar{E} = Y - X\Theta_*$ is unknown.

The requirement $\text{rk } \Theta_* = h$ implies that a zero transition matrix Θ_* occurs if and only if $h = 0$. In this extreme case, the righthand side of <3.18> equals zero. Proposition 3.6 explains this observation. In particular, comparing the final term in <3.15> with <3.1> reveals that the (unique) minimizer $\hat{\Theta}$ of the former equals the $m \times m$ zero matrix if and only if $\|(X^\top Y + Y^\top X)/(2(n-1))\|_{\text{op}} \leq \lambda$. Moreover, the equality $h = 0$ implies $\bar{E} = Y - X\Theta_* = Y$. Therefore, the requirement $\lambda \geq \|G\|_{\text{op}}$ ensures that in this special case the minimizer $\hat{\Theta}$ and thereby $\hat{\Delta}$ equals the $m \times m$ zero matrix, which verifies <3.18>. Finally, the inequality <3.18> is valid if $k = 1$, that is, $h = m + 1 - k = m$ due to perfect alignment. Then, $V_k = \mathbb{S}^m$, $V_k^\perp = \{0\}$, and the second summand on the lefthand side of <3.18> is absent. In particular, there is no need to ponder the meaning of a zeroth singular value. However, the below analysis is geared towards “large” k .

The second part of the analysis leading to proposition 3.9 takes the model structure presented in section 3.4.1 into account. This requires the definition of the matrices $F \in \mathbb{R}^{(n-1) \times h}$ and $V_2 \in \mathbb{R}^{(n-1) \times (m-h)}$ in analogy to X and Y , that is,

$$F = \begin{pmatrix} f_{1,1}(\omega) & \cdots & f_{1,h}(\omega) \\ \vdots & \ddots & \vdots \\ f_{n-1,1}(\omega) & \cdots & f_{n-1,h}(\omega) \end{pmatrix} \quad \text{and} \quad V_2 = \begin{pmatrix} v_{1,h+1}(\omega) & \cdots & v_{1,m}(\omega) \\ \vdots & \ddots & \vdots \\ v_{n-1,h+1}(\omega) & \cdots & v_{n-1,m}(\omega) \end{pmatrix}.$$

These definitions imply—by virtue of <3.8>—the two equalities

$$X = FU_1^\top + \rho V_2 U_2^\top \quad \text{and} \quad \|X\widehat{\Delta}\|^2 = \|FU_1^\top \widehat{\Delta} + \rho V_2 U_2^\top \widehat{\Delta}\|^2. \quad <3.19>$$

If $k = 1$, then $h = m$ and the quantities ρ , V_2 , and U_2 are absent. The same applies to the factor related quantities F and U_1 if $h = 0$. In this case, the remark following lemma 3.8 reveals that the equality $\|\widehat{\Delta}\| = \|\widehat{\Theta} - \Theta_*\| = 0$ holds whenever $\lambda \geq \|G\|_{\text{op}}$; thus, no further investigation is needed. In case $h > 0$, proposition 3.9 requires that the least singular value $\sigma_h(F^\top F/(n-1))$ of the symmetric and positive semidefinite $h \times h$ matrix $F^\top F/(n-1)$ exceeds a positive number $\kappa > 0$, which plays the role of a curvature constant as defined in section 3.1.2. This requirement amounts to $\text{rk } F^\top F = h$ or equivalently linear independence of the columns of F , which in turn necessitates $n - 1 \geq h$. A proof of proposition 3.9 follows on page 83 in appendix 3.b.

Proposition 3.9. *If Θ_* is perfectly aligned with Λ in the above sense, $\text{rk } \Theta_* = h$, and if $h > 0$, then $\sigma_h(F^\top F/(n-1)) \geq \kappa > 0$ for some $\kappa > 0$, as well as*

$$\begin{aligned} \lambda \geq \hat{\lambda} &= \|G\|_{\text{op}}, \quad G = \frac{X^\top \bar{E} + \bar{E}^\top X}{2(n-1)}, \quad \bar{E} = Y - X\Theta_*, \\ \xi \geq \hat{\xi} &= \begin{cases} \frac{1}{\sigma_{k-1}^q(\Lambda)} \left[\sigma_h \left(\frac{F^\top F}{n-1} \right) + 4\rho^2 \frac{\|V_2^\top F/(n-1)\|_{\text{op}}^2}{\sigma_h(F^\top F/(n-1))} \right], & h > 0, k > 1, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

then the minimizer $\widehat{\Delta} = \widehat{\Theta} - \Theta_*$ of $\bar{l}_{\lambda, \xi}$ in <3.16> satisfies

$$\|\widehat{\Delta}\| \leq \max \left\{ 20 \frac{\lambda}{\kappa} \sqrt{h}, 4 \sqrt{\frac{\xi}{\kappa}} \|\Lambda^{q/2} \Theta_*\| \right\} \quad \text{with } \kappa = 1 \text{ if } h = 0. \quad <3.20>$$

The lower bounds for λ and ξ are valid as both can—in principle—be calculated prior to the minimization process. In particular, if $h > 0$ and $k > 1$, then the requirements $k \leq \text{rk } \Lambda + 1$ and $\text{rk } \Lambda \geq 1$ guarantee the inequality $\sigma_{k-1}(\Lambda) \geq \sigma_{\min, \neq 0}(\Lambda) > 0$.

The (literal) numbers appearing in lemma 3.8 and proposition 3.9 are arbitrary to the degree that they reflect one of a range of possible choices used in the proofs. These proofs justify the form of $\hat{\lambda}$ and $\hat{\xi}$; however, a supplementary comment is in order. Section 3.4.1 defines the transition matrix Θ_* as the unique minimizer of the t -invariant objective $\mathbb{S}^m \ni \Theta \mapsto \|X_{t+1} - X_t \Theta\|^2/2 = \mathbb{E} \|x_{t+1} - \Theta x_t\|^2/2$. The latter expectation amounts to an integral over \mathbb{R}^{2m} with respect to the t -invariant *distribution* $\mu_{(x_t, x_{t+1})}$ of the random vector $(x_t, x_{t+1}) = (x_{t,1}, \dots, x_{t,m}, x_{t+1,1}, \dots, x_{t+1,m})$, that is, the image measure $\mathbb{P} \circ (x_t, x_{t+1})^{-1}$ on $(\mathbb{R}^{2m}, \mathcal{R}^{2m})$ with \mathcal{R}^{2m} symbolizing the Borel σ -field of the norm topology on \mathbb{R}^{2m} . The data-based term in <3.15> has the form of a similar integral but with respect to the *empirical distribution* $\hat{\mu}_{(x_t, x_{t+1})}$ given by $\hat{\mu}_{(x_t, x_{t+1})} B = \frac{1}{n-1} \sum_{t \leq n-1} 1_B(x_t(\omega), x_{t+1}(\omega))$, wherein 1_B symbolizes the indicator function of $B \in$

\mathcal{R}^{2m} . These two integrals differ to the extent that

$$\mathbb{E}\|x_{t+1} - \Theta x_t\|^2 = \mathbb{E}\|x_{t+1} - \Theta_* x_t\|^2 + \mathbb{E}\|(\Theta - \Theta_*)x_t\|^2, \quad \text{whereas}$$

$$\sum_{t \leq n-1} \frac{\|x_{t+1}(\omega) - \Theta x_t(\omega)\|^2}{n-1} = \frac{\|Y - X\Theta_*\|^2}{n-1} + \frac{\|X(\Theta - \Theta_*)\|^2}{n-1} + \frac{2\langle X^\top \bar{E}, \Theta - \Theta_* \rangle}{n-1}$$

contains an additional term. Therein, $X^\top \bar{E}/(n-1)$ can be replaced by its projection G onto \mathbb{S}^m due to symmetry of Θ and Θ_* . Hence, this (final) term is upper bounded by $2\|G\|_{\text{op}}\|\Theta - \Theta_*\|_{\text{nuc}}$, which shows that the given $\hat{\lambda}$ allows the $\|\cdot\|_{\text{nuc}}$ -part of $\langle 3.15 \rangle$ to counter the additional term. A similar remark applies to the second summand of $\hat{\xi}$. Its first summand serves a different purpose. Proposition 3.9 is geared towards the case $h < n-1 < m$ —although this is not explicitly stated, wherein the differences $n-1-h$ and $m-n+1$ are thought to be “substantial”. In case $n-1 > m-h$, a modified argument dispenses with the first summand of $\hat{\xi}$ and leads to a comparable upper bound.

3.5.3. Probabilistic guarantees

This section derives an expression for the probability that the upper bound $\langle 3.20 \rangle$ on the estimation error length $\|\hat{\Theta} - \Theta_*\|$ holds when estimating the transition matrix Θ_* in $\langle 3.11 \rangle$ via the unique minimizer $\hat{\Theta}$ of the objective function in $\langle 3.15 \rangle$. More specifically, the main result (proposition 3.13) provides positive numbers $\bar{\lambda}$, $\bar{\xi}$, and $\bar{\kappa}$ —depending on the matrices A_0, \dots, A_l , and ρ as well as m and the number of observations n —such that there exists a subset S of Ω contained in the σ -field \mathcal{F} with

$$S \subset \left\{ \omega \in \Omega \mid \sigma_h \left(\frac{F(\omega)^\top F(\omega)}{n-1} \right) \geq \bar{\kappa} \right\} \cap \left\{ \omega \in \Omega \mid \bar{\lambda} \geq \hat{\lambda}(\omega) \right\} \cap \left\{ \omega \in \Omega \mid \bar{\xi} \geq \hat{\xi}(\omega) \right\},$$

whose probability depends on the just mentioned model quantities. Hence, the minimizers $\hat{\Theta}$ (of $\langle 3.15 \rangle$ with $\lambda \geq \bar{\lambda}$, $\xi \geq \bar{\xi}$) and $\hat{\Delta}$ (of $\bar{l}_{\lambda, \xi}$, $\lambda \geq \bar{\lambda}$, $\xi \geq \bar{\xi}$) satisfy the inequality

$$\|\hat{\Theta} - \Theta_*\| = \|\hat{\Delta}\| \leq \max \left\{ 20\sqrt{h}\frac{\lambda}{\bar{\kappa}}, 4\sqrt{\frac{\xi}{\bar{\kappa}}}\|\Lambda^{q/2}\Theta_*\| \right\} \quad \langle 3.21 \rangle$$

for all $\omega \in S$, which is abbreviated as $\langle 3.21 \rangle$ being true *with probability at least* $\mathbb{P}S$. The question whether the set of ω satisfying $\langle 3.21 \rangle$ or the above superset of S are measurable, that is, elements of \mathcal{F} , is not addressed and has merely aesthetic value. The formal framework conforms with the construct in appendix 2.a. In particular, the above sets depend on the choice of basis element representatives; however, their probabilities are invariant to this choice due to the invariance of the underlying distributions. Finally, the present analysis focuses on $1 \leq h < m$ and considers a fixed choice of model quantities satisfying the restrictions of section 3.4. The conclusions apply generally but depend on these quantities. The case $h \in \{0, m\}$ receives only minimal attention.

In light of proposition 3.9, the present investigation amounts to a study of singular values—defined pointwise with respect to ω —of random matrices. Generally, if A

*with probability
at least*

symbolizes a $d_1 \times d_2$ random matrix, then $A(\omega)$ denotes the image of $\omega \in \Omega$ under A , that is, the element of $\mathbb{R}^{d_1 \times d_2}$ with i, j -th entry $a_{i,j}(\omega)$. Sections 3.5.1 and 3.5.2 omit the argument to simplify the notation, which is justified as these sections never refer to $\omega \mapsto A(\omega)$. This section considers both A and $A(\omega)$, which requires a more careful notation. In particular, the symbol A refers to a random matrix unless $A = A(\omega)$ is explicitly indicated. This comment also applies to random vectors and random variables.

The (transposed) rows of the random matrices considered here are given by

$$v_t^{(1)} = \begin{pmatrix} v_{t,1} \\ \vdots \\ v_{t,h} \end{pmatrix}, v_t^{(2)} = \begin{pmatrix} v_{t,h+1} \\ \vdots \\ v_{t,m} \end{pmatrix}, f_t = \begin{pmatrix} f_{t,1} \\ \vdots \\ f_{t,h} \end{pmatrix}, x_t = \begin{pmatrix} x_{t,1} \\ \vdots \\ x_{t,m} \end{pmatrix}, \bar{e}_t = x_t - \Theta_* x_{t-1}$$

Therein, the random variables $v_{t,j}$ with (t, j) ranging over a subset $I_v \subset \mathbb{N} \times \mathbb{N}$ are independent with zero mean and $\mathbb{E}v_{t,j}^2 = 1$. Section 3.4 presents the complete specification. Proposition 3.13 necessitates—on top of the specification in section 3.4—that the *distribution* of each random variable $v_{t,j}$, that is, the image measure $\mathbb{P} \circ v_{t,j}^{-1}$, is *subgaussian*. The latter requirement amounts to the existence of some $s_{t,j} > 0$ such that the inequality $\mathbb{P}\{|v_{t,j}| > t\} \leq \exp(1 - t^2/s_{t,j}^2)$ holds for all $t > 0$. Appendix 3.a contains a brief treatment of such distributions. Two facts are essential: firstly, a “large” subgaussian norm $\|v_{t,j}\|_{\psi_2} = \inf\{s > 0 \mid \mathbb{E} \exp((v_{t,j}/s)^2) \leq 2\}$ corresponds to a “slow” decay of the probabilities $\mathbb{P}\{|v_{t,j}| > w\}$ for $0 < w \rightarrow \infty$; and, secondly, $\|v_{t,j}\|_{\psi_2} \geq 1$ as $\mathbb{E}v_{t,j}^2 = 1$.

An analysis of the singular values of the symmetric and positive semidefinite matrix $F^\top F/(n-1)$ leads to an appropriate value $\bar{\kappa} > 0$ and showcases all steps involved in following investigations. The first step of the argument is pointwise with respect to ω , that is, $F = F(\omega)$ as in section 3.5.2. Sections 2.5.2 and 2.5.4 express the extreme singular values $\hat{\sigma}_1 = \hat{\sigma}_1(\omega) = \sigma_1(\frac{F^\top F}{n-1})$ and $\hat{\sigma}_h = \hat{\sigma}_h(\omega) = \sigma_h(\frac{F^\top F}{n-1})$ in the form $\hat{\sigma}_1 = \sup_{\|c\|=1} \langle \frac{F^\top F}{n-1} c, c \rangle$ and $\hat{\sigma}_h = \inf_{\|c\|=1} \langle \frac{F^\top F}{n-1} c, c \rangle$. Therein, the map $c \mapsto \langle \frac{F^\top F}{n-1} c, c \rangle$ is Lipschitz continuous on the unit sphere $\{\|\cdot\| = 1\}$ of \mathbb{R}^h . More specifically,

$$\left| \left\langle \frac{F^\top F}{n-1} c, c \right\rangle - \left\langle \frac{F^\top F}{n-1} c', c' \right\rangle \right| \leq \left| \left\langle \frac{F^\top F}{n-1} c', c - c' \right\rangle \right| + \left| \left\langle \frac{F^\top F}{n-1} c, c - c' \right\rangle \right| \leq 2\hat{\sigma}_1 \|c - c'\|$$

provides the upper bound $2\hat{\sigma}_1$ on its $(\|\cdot\|)$ -Lipschitz constant. Thus, <2.1> implies

$$\min_{i \leq q} \left\langle \frac{F^\top F}{n-1} c_i, c_i \right\rangle - 2\hat{\sigma}_1 \varepsilon \leq \hat{\sigma}_h \leq \hat{\sigma}_1 \leq \max_{i \leq q} \left\langle \frac{F^\top F}{n-1} c_i, c_i \right\rangle + 2\hat{\sigma}_1 \varepsilon, \quad <3.22>$$

wherein c_1, \dots, c_q provides an ε -net (section 2.1.2) of $\{\|\cdot\| = 1\}$ with $\varepsilon \in (0, 1)$.

Subsequently, the symbol F refers to the random matrix $\omega \mapsto F(\omega)$, whose (transposed) rows are given by the random vectors f_1, \dots, f_{n-1} . Consequently, the summands of $\langle F^\top F c, c \rangle = \sum_{t \leq n-1} \langle f_t, c \rangle^2$ with $c \in \{\|\cdot\| = 1\}$ equal $\langle f_t, c \rangle^2 = \langle B_t v, c \rangle^2 = v^\top B_t^\top c c^\top B_t v$, wherein $B_t \in \mathbb{R}^{h \times m(n+l)}$, $t \leq n-1$, has the form

$$\left[\underbrace{\begin{matrix} 0 & \dots & 0 \end{matrix}}_{n-t \text{ zero matrices in } \mathbb{R}^{h \times m}} \quad \bar{A}_0 \quad \dots \quad \bar{A}_l \quad \underbrace{\begin{matrix} 0 & \dots & 0 \end{matrix}}_{t-1 \text{ zero matrices in } \mathbb{R}^{h \times m}} \right], \quad \bar{A}_j = [A_j \quad 0] \in \mathbb{R}^{h \times m},$$

and the random vector v consists of $v_n^{(1)}, v_n^{(2)}, v_{n-1}^{(1)}, \dots, v_{1-l}^{(1)}, v_{1-l}^{(2)}$ —in that order from top to bottom—with $v_j^{(2)}$ equal to the zero vector in \mathbb{R}^{m-h} for $1-l \leq j \leq 0$. Hence, the entries of v are independent and exhibit subgaussian distributions. In total, this representation implies the equality $\langle F^\top F c, c \rangle = v^\top A_c v$ with $A_c = \sum_{t \leq n-1} B_t^\top c c^\top B_t$.

The expectation of the summands $\langle f_t, c \rangle^2 = \sum_{i \leq h} \sum_{j \leq h} c_i c_j f_{t,i} f_{t,j}$ is given by $\langle V_f c, c \rangle$ for all $t \leq n-1$, wherein $V_f = \sum_{i=0}^l A_i^2$ equals the t -invariant Gramian $\langle F_t, F_t \rangle$ of the linear map $F_t = [f_{t,1} \cdots f_{t,h}]$. Moreover, the examples (d₁) and (d₂) in section 2.5.2 together with $\text{rk } A_c \leq n-1$ imply the (in)equalities $\|A_c\|^2 = \sum_{j \leq \text{rk } A_c} \sigma_j^2(A_c) \leq (n-1)\|A_c\|_{\text{op}}^2$, wherein the inequality for the rank follows from the inclusion $\text{img } A_c \subset \text{span}\{B_t^\top c \mid t \leq n-1\}$. As a consequence, every unit length $c \in \mathbb{R}^h$ satisfies

$$\frac{(n-1)^2 \langle V_f c, c \rangle^2}{4C^4 \|A_c\|^2} \geq \frac{(n-1)^2 \langle V_f c, c \rangle^2}{4C^4 (n-1) \|A_c\|_{\text{op}}^2} = (n-1) \left(\frac{\langle V_f c, c \rangle}{2C^2 \|A_c\|_{\text{op}}} \right)^2.$$

Thus, the Hanson-Wright inequality (lemma 3.14 in appendix 3.a) yields

$$\begin{aligned} \mathbb{P}\left\{ \left| \left\langle \frac{F^\top F}{n-1} c, c \right\rangle - \langle V_f c, c \rangle \right| > \frac{1}{2} \langle V_f c, c \rangle \right\} &= \mathbb{P}\left\{ |v^\top A_c v - \mathbb{E} v^\top A_c v| > \frac{n-1}{2} \langle V_f c, c \rangle \right\} \\ &\leq 2 \exp\left(-\bar{C}(n-1) \min\{\zeta_c, \zeta_c^2\} \right), \quad \text{wherein } \zeta_c = \frac{\langle V_f c, c \rangle}{2C^2 \|A_c\|_{\text{op}}} \end{aligned}$$

and $C \geq 1, \bar{C} > 0$ equal an upper bound on the subgaussian norms $\|v_{t,j}\|_{\psi_2}$, $(t, j) \in I_v$, and the (unspecified) constant in the Hanson-Wright inequality, respectively.

This inequality holds for every unit length $c \in \mathbb{R}^h$. In particular, it applies to all elements c_1, \dots, c_q of a \subset -minimal ε' -net of $\{\|\cdot\| = 1\}$, wherein $\varepsilon' = \sigma_h(V_f)/(20\sigma_1(V_f))$. The choice of ε' is tailored to the below derivations and ensures $\varepsilon' < 1/2$, that is, $1 - 2\varepsilon' > 0$. Next, an application of the *union bound* $\mathbb{P} \cup_{i \leq q} A_i \leq \sum_{i \leq q} \mathbb{P} A_i$, which holds for arbitrary \mathcal{F} -measurable sets A_1, \dots, A_q , leads to the inequality

$$\mathbb{P} \cap_{i \leq q} \left\{ \frac{1}{2} \langle V_f c_i, c_i \rangle \leq \left\langle \frac{F^\top F}{n-1} c_i, c_i \right\rangle \leq \frac{3}{2} \langle V_f c_i, c_i \rangle \right\} \geq 1 - 2 \sum_{i \leq q} \exp(-\bar{C}(n-1)\eta_i), \quad \langle 3.23 \rangle$$

wherein $\eta_i = \min\{\zeta_{c_i}, \zeta_{c_i}^2\}$ with $\zeta_{c_i} > 0$ whenever $h > 0$ due to the above requirements. Lemma 2.1 and \subset -minimality of the chosen ε' -net imply that q equals the covering number $N(\{\|\cdot\| = 1\}, \|\cdot\|, \varepsilon') \leq (1 + 2/\varepsilon')^h \leq \exp(h \log[41\sigma_1(V_f)/\sigma_h(V_f)])$.

If ω lies in the intersection on the lefthand side of $\langle 3.23 \rangle$, then the inequality $\varepsilon' < 1/2$ together with the final inequality in $\langle 3.22 \rangle$ imply that

$$\hat{\sigma}_1(\omega) \leq \frac{1}{1 - 2\varepsilon'} \max_{i \leq q} \frac{\langle F(\omega)^\top F(\omega) c_i, c_i \rangle}{n-1} \leq \frac{5}{3} \sigma_1(V_f),$$

Next, the first inequality of $\langle 3.22 \rangle$ ensures that all ω in the above intersection satisfy

$$\hat{\sigma}_h(\omega) \geq \frac{1}{2} \sigma_h(V_f) - \frac{10}{3} \sigma_1(V_f) \varepsilon' = \frac{1}{3} \sigma_h(V_f).$$

Similar arguments verify that elements ω of this intersection also satisfy

$$\hat{\sigma}_h(\omega) \leq \min_{i \leq q} \frac{\langle F(\omega)^\top F(\omega) c_i, c_i \rangle}{n-1} \leq \frac{3}{2} \min_{i \leq q} \langle V_f c_i, c_i \rangle \leq \frac{3}{2} \left[\sigma_h(V_f) + 2\sigma_1(V_f)\varepsilon' \right] \leq 2\sigma_h(V_f).$$

These inequalities hold simultaneously with probability at least $1 - \delta$, $\delta \in (0, 1)$, if

$$n-1 \geq \frac{1}{\bar{C} \min_{i \leq q} \min\{\zeta_{c_i}^2, \zeta_{c_i}\}} \left(\log \left[41 \frac{\sigma_1(V_f)}{\sigma_h(V_f)} \right] h + \log \frac{2}{\delta} \right). \quad \langle 3.24 \rangle$$

Lemma 3.10 provides a lower bound on the denominator of the second factor on the right-hand side. Therein, the diagonal matrices A_0, \dots, A_l exhibit a *uniform decay rate* $\alpha > 0$ if $\sum_{i=k}^l \|A_i c\| \leq (\sum_{i=0}^l \|A_i c\|) \exp(1 - \alpha k)$ for all $c \in \mathbb{R}^{l+1}$ and $0 \leq k \leq l$. Every sequence A_0, \dots, A_l exhibits a uniform decay rate of $1/l$. However, larger uniform decay rates are possible. In particular, if $A_0 = \rho I$, $A_i = \rho D^i$, $\rho > 0$, wherein I and D symbolize the $h \times h$ identity matrix and a diagonal matrix with nonzero diagonal entries $d_{i,i} \in (-1, 1)$, respectively, then one has $\sum_{i=k}^l \|A_i c\| \leq \bar{d}^k \sum_{i=0}^l \|A_i c\| \leq (\sum_{i=0}^l \|A_i c\|) \exp(1 - k \log(1/\bar{d}))$, wherein \bar{d} represents the maximal absolute diagonal entry $\max_{i \leq h} |d_{i,i}| < 1$ of D .

uniform decay rate

Lemma 3.10. *If the sequence A_0, \dots, A_l exhibits a uniform decay rate α , then using the above notation one has $\min_{i \leq q} \min\{\zeta_{c_i}^2, \zeta_{c_i}\} \geq \bar{\zeta}^2$ with $\bar{\zeta} = \alpha / (3C^2(3 + \alpha))$.*

Lemma 3.10 reveals that the number of observations n has to exceed a constant times the number of factors h for the above inequalities to hold with “high” probability. Therein, the constant grows with the subgaussian norms $\|v_{t,j}\|_{\psi_2}$ and decreases as the uniform decay rate α increases. A proof starts on page 84 in appendix 3.b.

A comparable analysis—starting on page 85 in appendix 3.b—leads to lemma 3.11. The final paragraph of section 3.5.2 mentions that the present analysis targets the case $m \geq n - 1 \geq h$. The above discussion reveals the importance of the second inequality $n - 1 \geq h$. Lemma 3.11 requires the first inequality $m \geq n - 1$. The case $m < n - 1$ necessitates a modified argument and leads to a different result.

Lemma 3.11. *If $0 < h < m$, the sequence A_0, \dots, A_l exhibits a uniform decay rate α , the distribution of $v_{t,j}$ is subgaussian with $\|v_{t,j}\|_{\psi_2} \leq C$ for some $C > 0$ and all $(t, j) \in I_v$, and $m \geq n - 1$, then using the above notation one has*

$$\left\| \frac{V_2^\top F}{n-1} \right\|_{\text{op}} \leq \bar{\bar{C}} C^2 \left(\frac{1 + \alpha}{\alpha} \right)^{1/2} \sigma_1^{1/2}(V_f) \frac{m}{n-1}$$

with probability at least $1 - 1/2^{m-1}$, wherein $\bar{\bar{C}} > 1$ denotes a constant which does not depend on the model quantities and $V_f = \sum_{i=0}^l A_i^2$,

Lemma 3.12 focuses on the operator norm $\|G\|_{\text{op}}$ of G . The analysis leading to its assertion possesses the same structure as the two previous investigations but is complicated by the structure of the rows $(\bar{e}_{t,1}, \dots, \bar{e}_{t,m})$ of \bar{E} shown in $\langle 3.12 \rangle$. To simplify its statement, the sequence of diagonal matrices A_0, \dots, A_l is said to exhibit a *uniform*

autoregressive approximation factor $\beta \geq 0$ if $\|(A_i - \Gamma_* A_{i-1})c\| \leq \beta \max\{\|A_i c\|, \|A_{i-1} c\|\}$ for all $1 \leq i \leq l$ and unit length $c \in \mathbb{R}^h$, wherein $\Gamma_* = (\sum_{i=0}^l A_i^2)^{-1} \sum_{i=0}^{l-1} A_i A_{i+1}$. The Cauchy-Schwarz inequality implies that all diagonal entries of Γ_* lie in $[-1, 1]$. Consequently, every sequence A_0, \dots, A_l has a uniform autoregressive approximation factor of 2. At the other extreme, the above special case, namely, $A_0 = \rho I$, $A_i = \rho D^i$, $i \leq l$, $\rho > 0$, I being the $h \times h$ identity matrix, and D symbolizing a diagonal matrix with nonzero diagonal entries $d_{i,i} \in (-1, 1)$, has an approximation factor of $\bar{d}^{2l+1} / (\sum_{i=0}^l \bar{d}^{2i})$, wherein $\bar{d} = \max_{i \leq h} |d_{i,i}| < 1$. A proof of lemma 3.12 starts on 86 in appendix 3.b.

Lemma 3.12. *If $h > 0$, the sequence A_0, \dots, A_l exhibits a uniform decay rate $\alpha > 0$ and a uniform autoregressive approximation factor $\beta \geq 0$, the distributions of $v_{t,j}$ is subgaussian with $\|v_{t,j}\|_{\psi_2} \leq C$ for some $C > 0$ and all $(t, j) \in I_v$, then for $m \geq n - 1$ and using the above notation one has*

$$\|G\|_{\text{op}} = \left\| \frac{X^\top \bar{E} + \bar{E}^\top X}{2(n-1)} \right\|_{\text{op}} \leq \bar{C} C^2 (1 + \beta) \left(\frac{1 + \alpha}{\alpha} \sigma_1(V_f) + \rho^2 \right) \frac{m}{n-1}$$

with probability at least $1 - 1/2^{m-2}$, wherein $\bar{C} > 1$ represents a constant which is unrelated to the model quantities, $V_f = \sum_{i=0}^l A_i^2$, and $\rho = 0$ if $h = m$.

If $h = 0$, then the same result applies with $(1 + \alpha)\sigma_1(V_f)/\alpha = \beta = 0$.

Finally, combining lemma 3.10, 3.11, and 3.12 with proposition 3.9 yields proposition 3.13, whose details are proved on page 88 in appendix 3.b.

Proposition 3.13. *Let Θ_* be perfectly aligned with Λ in the above sense, $\text{rk } \Theta_* = h$, and the distribution of $v_{t,j}$ be subgaussian with $\|v_{t,j}\|_{\psi_2} \leq C$ for some $C > 1$ and all $(t, j) \in I_v$. If $h \geq 1$, then let the sequence A_0, \dots, A_l exhibit a uniform decay rate $\alpha > 0$ and a uniform autoregressive approximation factor $\beta \geq 0$. Under these conditions, there exist $C_1, C_2 > 1$, $C_3, C_4, C_5 > 0$ not depending on the model quantities such that*

$$m \geq n - 1 \geq C_1 C^4 \left(\frac{1 + \alpha}{\alpha} \right)^2 \left(\log \left[C_2 \frac{\sigma_1(V_f)}{\sigma_h(V_f)} \right] h + \log \frac{2}{\delta} \right), \quad \delta \in (0, 1),$$

together with the lower bounds

$$\xi \geq \bar{\xi} = \frac{C_3}{\sigma_{k-1}^q(\Lambda)} \left(\sigma_h(V_f) + \rho^2 C^4 \frac{1 + \alpha}{\alpha} \frac{\sigma_1(V_f)}{\sigma_h(V_f)} \left[\frac{m}{n-1} \right]^2 \right), \quad \text{and}$$

$$\lambda \geq \bar{\lambda} = C_4 C^2 (1 + \beta) \left(\frac{1 + \alpha}{\alpha} \sigma_1(V_f) + \rho^2 \right) \frac{m}{n-1}$$

guarantees that the unique minimizer $\hat{\Theta}$ of <3.15> satisfies the inequality

$$\|\hat{\Theta} - \Theta_*\| \leq C_5 \max \left\{ \frac{\lambda}{\bar{\kappa}} \sqrt{h}, \sqrt{\frac{\xi}{\bar{\kappa}}} \|\Lambda^{q/2} \Theta_*\| \right\} \quad \text{<3.25>}$$

with probability at least $1 - \delta - 1/2^{m-3}$. Herein, $\bar{\kappa} = \frac{\sigma_h(V_f)}{3}$ if $h > 0$ and one otherwise. In the latter case, all factor related quantities above vanish and $\beta = 0$.

The requirement on h, n, m in proposition 3.13 implies the inequalities $m \geq C'h + C'' \log(2/\delta) > h$ with $C', C'' > 1$, which precludes the equality $h = m$. If one disregards the other extreme case $h = 0$, then the lowest possible upper bound in proposition 3.13 follows when using $\lambda = \bar{\lambda}$ and $\xi = \bar{\xi}$. This upper bound can be simplified to

$$\begin{aligned} C_5 \max \left\{ \frac{\lambda}{\bar{\kappa}} \sqrt{h}, \sqrt{\frac{\xi}{\bar{\kappa}}} \|\Lambda^{q/2} \Theta_*\| \right\} &\leq C_5 \left(\frac{\lambda}{\bar{\kappa}} \sqrt{h} + \sqrt{\frac{\xi}{\bar{\kappa}}} \|\Lambda^{q/2} \Theta_*\| \right) \\ &\leq C'_5 C^2 \left(1 + \beta + \left[\frac{\sigma_k(\Lambda)}{\sigma_{k-1}(\Lambda)} \right]^{q/2} \right) \left(\frac{1 + \alpha}{\alpha} + \frac{\rho^2}{\sigma_1(V_f)} \right) \frac{\sigma_1(V_f)}{\sigma_h(V_f)} \sqrt{\frac{mh}{n-1}} \sqrt{\frac{m}{n-1}} \quad <3.26> \end{aligned}$$

with $C'_5 > 0$ representing a constant which does not depend on the model quantities. Therein, the term $(\sigma_k(\Lambda)/\sigma_{k-1}(\Lambda))^{q/2}$ is a consequence of perfect alignment of the transition matrix Θ_* with Λ , that is, $\text{img } \Theta_* = L_k$, and the definition of $\Lambda^{q/2}$ in section 3.4.2. The requirement $\sigma_k(\Lambda) > \sigma_{k-1}(\Lambda)$ implies that this term vanishes as $q \rightarrow \infty$.

Several interpretable quantities occur in <3.26>. The upper bound C on the subgaussian norms $\|v_{t,j}\|_{\psi_2}$, $(t, j) \in I_v$, controls how $\mathbb{P}\{|v_{t,j}| > w\}$ shrinks to zero as $0 < w \rightarrow \infty$. The decay rate α quantifies the temporal dependence of the factors. In addition, the ratio $\rho^2/\sigma_1(V_f)$ compares the “size” of the non-factor part and factor part of x_t ; similarly, $\sigma_1(V_f)/\sigma_h(V_f)$ measures the “size differences” between factors. The approximation factor β expresses how well the transition in <3.12> fits the true dynamics of the factors, while $(\sigma_k(\Lambda)/\sigma_{k-1}(\Lambda))^{q/2}$ represents the “misspecification” of Λ . The ratio $mh/(n-1)$ roughly equals the “number of parameters” in Θ_* per observation. In fact, section 2.5.4 expresses Θ_* as the sum $\Theta_+ - \Theta_-$, of two symmetric and positive semidefinite matrices Θ_+, Θ_- , which are uniquely determined by Θ_* and satisfy $\text{img } \Theta_+ \subset (\text{img } \Theta_-)^\perp$ and $\text{rk } \Theta_+ + \text{rk } \Theta_- = h$. Careful reading of section 2.4.2 reveals that each of the two summands uniquely corresponds to a $\text{rk } \Theta_\pm \times m$ row echelon matrix R^\pm with $r_{i,\inf I_i}^\pm > 0$ for all $i \leq \text{rk } \Theta_\pm$, wherein $I_i = \{j \leq m \mid r_{i,j} \neq 0\}$ as in section 2.2.1. Furthermore, if the number of factors h is “small” relative to m , then the representation as a pair of row echelon matrices with mutually orthogonal rows contains approximately mh entries.

Finally, two closing comments are in order. As previously mentioned, the use of the uniform decay rate α and the uniform autoregressive factor β entails no additional restrictions. These notions summarize two aspects of the sequence A_0, \dots, A_l and thereby simplify but also weaken the resulting statements. In fact, their definition is tailored to the special case $A_0 = \rho I$, $A_i = \rho D^i$, $i \leq l$, $\rho > 0$, wherein I represents the $h \times h$ identity matrix, and D denotes a diagonal matrix with nonzero diagonal entries in $(-1, 1)$. The resulting upper bounds are overly generous for “very different” cases.

The lower bounds $\bar{\xi}$ and $\bar{\lambda}$ in proposition 3.13 provide no guidance in practical situations as their definition involves several unobserved quantities including $\sigma_{k-1}(\Lambda)$ since $k = m + 1 - h$ is unknown. Hence, proposition 3.13 merely states that minimizing an objective function of the form <3.15> is a “meaningful” approach in the sense that

there exist $\lambda, \xi > 0$ which imply that an inequality of the form $\|\hat{\Theta} - \Theta_*\| \leq C^* \sqrt{h}$ with $C^* > 1$ holds with “high” probability once the number of observations exceeds a multiple of the number of factors. However, this guarantee is practically useless as the alternative estimator $\tilde{\Theta} = 0$ satisfies $\|\tilde{\Theta} - \Theta_*\| = \|\Theta_*\| \leq \sqrt{h}$ with probability one.

Comments and references

Section 3.1 Rockafellar (1970, part I–IV) presents all topics of this section in full detail. He shows that a nonempty interior is merely a sufficient condition for the existence of supporting hyperplanes at boundary points; a refined existence criterion relies on the concept of a *relative interior* (Rockafellar, 1970, sec. 6, thm. 11.6). Boyd and Vandenberghe (2004, sec. 2.1–2.3, 2.5) discuss basic properties of convex sets. Their presentation contains numerous examples and illustrations. In particular, their figure 2.8 partially resembles panel (B) of figure 3.1. Figure B.5 of Bertsekas (1999, app. B) covers other aspects of that figure. Figure 3.2 can be found in Chandrasekaran et al. (2012, fig. 1, panel (b)). The basis choice $\bar{B}_{1,1}, \bar{B}_{1,2}, \bar{B}_{2,2}$ is commonplace in literature; the alternative $\binom{1}{1}/\sqrt{2}, \binom{1}{-1}/\sqrt{2}, \binom{1}{-1}/\sqrt{2}$ leads to a more symmetric appearance of $\{\|\cdot\|_{\text{nuc}} \leq \ell\}$, $\ell > 0$. Chandrasekaran et al. (2012, sec. 2.2) identify $\{\|\cdot\|_{\text{nuc}} \leq \ell\}$ as the convex hull of the rank one matrices. The assertion of all proper and closed faces of the latter set being exposed can be obtained by adapting the proof of So (1990, thm. 3), which applies to $\mathbb{R}^{m \times m}$, to \mathbb{S}^m . Ziętak (1993, sec. 4, p. 140) indicates this possibility but is concerned with $\mathbb{R}^{m \times m}$. His equation (2.2) (section 2) contains the nucleus of lemma 3.2. Drusvyatskiy et al. (2015, sec. 2.3, def. 2.8) treat sets of the form $\{d \in \partial f(\cdot)\}$ under the name *exposed face of f* and relate these faces to the exposed faces of $\text{epi } f$ in their theorem 2.9. Nesterov (2004, def. 2.1.2) caters the definition of strong convexity for differentiable functions. His theorem 2.1.9 contains an alternative definition of this concept, which does not presuppose differentiability. Hiriart-Urruty and Lemaréchal (1996, ch. X) discuss conjugates (*Legendre(-Fenchel) transforms*) in-depth.

Section 3.2 Negahban and Wainwright (2011a, sec. 2.2) consider the criterion l_λ in slightly different form. Fazel (2002, sec. 5.1.4) links $\|\cdot\|_{\text{nuc}}$ and the rank via conjugation; her analysis may also be found in Fazel et al. (2001). Yuan et al. (2007, sec. 2) motivate the use of $\|\cdot\|_{\text{nuc}}$ via a general framework for dimension reduction. Using $g = \lambda \|\cdot\|_{\text{nuc}}$ (as in l_λ) in place of λrk leads to a convex criterion function and thus facilitates the computation of the minimizer $\hat{\Theta}$. Herein, the superscript $\hat{\cdot}$ does not indicate a projection—as in section 2.3; the same comment applies to the use of $\tilde{\cdot}$ in this section. If Θ ranges over all of $\mathbb{R}^{n \times p}$ with $Y \in W^{n \times p}$, then there exists a simple solution to the minimization of $\|Y - X\Theta\|^2 / (2n) + \lambda \text{rk}$ (Bunea et al., 2011, sec. 2). Thus, computational gains alone do not justify the use of $\|\cdot\|_{\text{nuc}}$ as a replacement for rk in this case.

The derivation of the dual problem <3.3> and the associated optimality conditions is an application of Borwein and Lewis (2010, sec. 3.3, thm. 3.3.5, exercise 9, sec. 4.2, exercise 17). Bach et al. (2011, prop. 2.2 and below, eq. 2.9–2.11)[†] present a compa-

[†]I learned about this work only after the review process.

rable duality analysis based on the same tools but tailored to a minimization problem over \mathbb{R}^p . In particular, their equations (2.7) and (2.8) are akin to <3.2>. Mishra et al. (2013, prop. 2.3)[†] adapt these considerations to a minimization problem over $\mathbb{R}^{m \times p}$. An analogous perturbation of the second summand g together with Hiriart-Urruty and Lemaréchal (1996, thm. 2.1.1) provides an alternative route to the same dual problem. Pong et al. (2010) consider the minimization of (a multiple of) l_λ but over $\Theta \in \mathbb{R}^{m \times p}$ (and with $Y \in \mathbb{R}^{n \times p}$). This problem generally allows the reduction to the case of X with linearly independent columns. Then, the uniqueness assertion corresponding to proposition 3.6 is immediate. The equivalent of the final condition in that proposition is their equation (19). These authors also derive a dual problem based on the mentioned reduction, which is akin to the above mentioned alternative perturbation approach. As a consequence, the equivalent of <3.3> matches their expression (8) after some rephrasing. Figure 3.5 resembles Bertsekas (1999, fig. 5.4.3). It focuses on a positive slope D ; negative slopes lead to a negative value of the dual objective <3.3>. Boyd and Vandenberghe (2004, sec. 3.3, ex. 3.26, 3.27) provide the arguments used to derive g_\star and f_\star .

Section 3.3 Bertsekas (2014, sec. 6.10.1) discusses the improvement $(f \circ X)(\widehat{\Theta}^{(k-1)}) \geq (f \circ X)(\widetilde{\Theta}^{(k)})$ in the first summand of the criterion <3.1> via a gradient step. Ji and Ye (2009, sec. 3) provide most of the algorithm <3.7> and the arguments in section 3.3. Pong et al. (2010, sec. 4, eq. (22)) suggest termination based on the *duality gap* $p_k - d_k$, however, in the form $(p_k - d_k)/(1 + d_k) < \varepsilon$ to allow for a problem independent choice of $\varepsilon > 0$. Bach et al. (2011, sec. 2.1.2.3)[†] present an analogous stopping rule to that used in the recipe <3.7> (line 6–11) for a minimization problem over \mathbb{R}^p . Toh and Yun (2010, sec. 3.5) suggest an alternative (subgradient-based) stopping rule. The algorithm <3.7> is not an optimal gradient-based scheme. It ensures $l_\lambda(\widehat{\Theta}^{(k)}) - l_\lambda(\widehat{\Theta}) \leq C/k$, wherein $C > 0$ is a problem-specific constant (proposition 3.7). The accelerated algorithm of Ji and Ye (2009, sec. 4) replaces the latter bound with C/k^2 . Herein, the role of <3.7> is to show that iterative schemes with low complexity of each iteration suffice to approximately recover $\widehat{\Theta}$. The first summand of the criterion function l_λ exhibits strong convexity if and only if $\ker X = \{0\}$ and then exhibits a curvature constant equal to a multiple of the square of the m -th singular value $\sigma_m(X) > 0$ of X .

Section 3.4 Bickel and Doksum (2015, sec 1.1) describe the usual model for mathematical statistics. The setup sketched in first paragraph of section 3.4 resembles these ideas. The present setting is *semiparametric* in the sense that only a few properties of the *joint distribution* of $v_{t,j}$, that is, the image measure of $(v_{t,j})$ under \mathbb{P} , enter the analysis.

Factor models have received much attention in literature; Stock and Watson (2011) provide an introductory overview. The *poor man's factor model* <3.8>, a stripped-down variant of (the simplest variation of) the original concept, combines a factor interpretation with the availability of an approximating *vector autoregressive model* $X_t = X_{t-1}\Theta_* + G_t$ as discussed below <3.12> with symmetric transition matrix Θ_* . Lütkepohl (2007, sec. 2.1) covers the latter model class with transition matrix $\Theta_* \in \mathbb{R}^{m \times m}$. Basu (2014, sec. V) considers vector autoregressive approximations to (more) general factor

models. The asymptotic setup alluded to in the approximation is akin to the usual *triangular array* setting of stochastic limit theory (Pollard, 2002, sec. 7.2). Koltchinskii and Rangel (2013, sec. 1) provide the notion of smoothness of u_1, \dots, u_h , the overall measure $\|\Lambda^{1/2}\Theta_*\|^2$, and the approximation $P_{V_k}\Theta_*$ in terms of $\bar{O}_{i,j}$ discussed in section 3.4.2. The pair $(\{1, \dots, m\}, (i, j) \mapsto w_{i,j})$ forms a *weighted graph* with *Laplacian* (matrix) Λ . Bapat (2014, lem. 4.2 (ii)) provides the rank of Λ but for unweighted graphs.

Section 3.5 Koltchinskii and Rangel (2013, eq. (4.1)) consider the objective function <3.15> but in a somewhat different setting. The proofs of their theorems 4 and 6 provide most of the techniques used in the proof of lemma 3.8. However, their analysis of the consequences of including a smoothness term is more refined than assuming perfect alignment. Incorporating these refinements should be the first step towards a practically useful upper bound. Negahban et al. (2012, sec. 2,3) advocate the two step approach of first deriving a data-based bound (section 3.5.2) followed by a probabilistic analysis (section 3.5.3). Their section 2.2 provides an in-depth discussion of the decomposability visualized in figure 3.9 as well as associated subspace pairs. However, their strategy for deriving data-based bounds requires the structure-based term of the objective function to satisfy the triangle inequality, hence, cannot be applied to <3.15>. The particular subspace pairs (V_k, \bar{V}_k^\perp) , $k \leq m$, are from Koltchinskii and Rangel (2013, sec. 1).

Negahban and Wainwright (2011a) analyze $\|\cdot\|_{\text{nuc}}$ -penalized estimation for a first-order vector autoregressive model using the just mentioned two step approach. The strategy of the proof of proposition 3.13 is borrowed from the proofs of their lemma 4 and 5 (Negahban and Wainwright, 2011b, appendix G). As a consequence of including merely a $\|\cdot\|_{\text{nuc}}$ -based part—no smoothness term, their result requires the number of observations n to be larger than m . Moreover, they assume Gaussian distributions. These two limitations also apply to the more general results of Basu (2014, prop. V.1, V.2) for the vector autoregressive setting. In this text, the generalization to subgaussian distributions becomes possible by restricting considerations to *l-dependence* (Shumway and Stoffer, 2006, appendix A.2). Basu and Michailidis (2015, sec. 2.3) suggest the use of the Hanson-Wright inequality, but focus on spectral techniques in *high dimensional* problems, that is, with m “large” compared to n . Vershynin (2012, lem. 5.4) states the upper bound on the Lipschitz constant of $c \mapsto \langle \frac{F^T F}{n-1} c, c \rangle$ in section 3.5.3.

The present analysis considers the transition matrix Θ_* to be a low rank element of $\text{span}\{\bar{O}_{i,j} \mid j \geq i \geq k\}$ with k smaller but close to $\text{rk } \Lambda + 1$. A different formulation considers $\Theta_* = \sum_{i \leq j \leq m} \beta_{h_{i,j}} \bar{O}_{i,j}$ with only few nonzero β_h , that is, $\beta \in \mathbb{R}^{m(m+1)/2}$ being *sparse*. Again, $\text{rk } \Lambda < m - 1$ results in some ambiguity, which needs to be resolved. In this framework, the induced basis $\bar{O}_{i,j}$ is often called a *dictionary* (Koltchinskii, 2011, sec. 1.6, 8.1). In comparison, mere $\|\cdot\|_{\text{nuc}}$ -penalization assumes sparsity but with respect to an unknown dictionary (Negahban and Wainwright, 2011a, sec. 2.1).

Finally, turning the insights of this section into a practically useful procedure requires finding data-based terms $\tilde{\lambda}$ and $\tilde{\xi}$ such that $\tilde{\lambda} \geq \bar{\lambda}$ and $\tilde{\xi} \geq \bar{\xi}$ hold simultaneously with “high probability”, wherein $\bar{\lambda}$ and $\bar{\xi}$ are defined in proposition 3.13. Then minimizing <3.15> with $\lambda = \tilde{\lambda}$ and $\xi = \tilde{\xi}$ is feasible and satisfies <3.25> with $\lambda = \tilde{\lambda}$ and $\xi = \tilde{\xi}$ and—by virtue of a union bound argument—with “high probability”. However, the

resulting data-based bound still contains unknown quantities.

Appendixes Vershynin (2012, sec. 5.2.3) motivates the subgaussian property via the well-known Gaussian distribution(s). His lemma 5.5 shows the equivalence of several alternative characterizations of this property; however, his definitions $\psi_2 = \exp(\cdot^2 - 1)$ and $\|x\|_{\psi_2} = \sup_{n \geq 1} (\mathbb{E}|x|^n)^{1/n} / \sqrt{n}$ are slightly different. Pollard (2002, sec. 4.4, ex. 26) explains the argument based on Tonelli’s theorem in detail.

The subgaussian random variables together with $\|\cdot\|_{\psi_2}$ form a so-called *Orlicz space*. Such spaces result when replacing ψ_2 by the composition $\psi \circ |\cdot|$ of the absolute value $|\cdot|$ and a monotone increasing, convex function $\psi : [0, \infty) \rightarrow [0, \infty)$ such that $\psi(0) = 0$ and $\psi(x) \rightarrow \infty$ when $x \rightarrow \infty$ (Pollard, 2002, exercise 22). Vaart and Wellner (2000, sec. 2.2) list several examples of such spaces including the \mathcal{L}^p -spaces and the subgaussian random variables. Their definition of ψ_2 coincides with the one used here.

Literature abounds with probability inequalities for subgaussian random variables; Vershynin (2012, sec. 5.2.3) contains an introductory treatment. Rudelson and Vershynin (2013, sec. 1) supply the Hanson-Wright inequality alongside an accessible proof. These authors use the above mentioned alternative definition of $\|\cdot\|_{\psi_2}$, which is subsequently denoted by $\|\cdot\|_{\psi_2}^*$. The distinction between $\|\cdot\|_{\psi_2}$ and $\|\cdot\|_{\psi_2}^*$ is immaterial to lemma 3.14. More specifically, there exist (compatibility) constants $0 < d < D$ such that the inequalities $d\|\cdot\|_{\psi_2} \leq \|\cdot\|_{\psi_2}^* \leq D\|\cdot\|_{\psi_2}$ hold for all subgaussian random variables. Therefore, the Hanson-Wright inequality is valid with either definition of the subgaussian norm, but the two cases require different constants $\bar{C} > 0$.

- Bach, F., R. Jenatton, J. Mairal, and G. Obozinski (2011). Convex optimization with sparsity-inducing norms. In S. Sra, S. Nowozin, and S. J. Wright (Eds.), *Optimization for Machine Learning* (1 ed.), Neural Information Processing, pp. 19–53. MIT Press.
- Bapat, R. B. (2014). *Graphs and Matrices* (2 ed.). Universitext. London: Springer.
- Basu, S. (2014). *Modeling and Estimation of High-dimensional Vector Autoregressions*. Ph. D. thesis, The University of Michigan.
- Basu, S. and G. Michailidis (2015). Regularized estimation in sparse high-dimensional time series models. *The Annals of Statistics* 43(4), 1535–1567.
- Bertsekas, D. P. (1999). *Nonlinear programming* (2 ed.). Athena Scientific.
- Bertsekas, D. P. (2014). *Supplementary Chapter 6 on Convex Optimization Algorithms*, Chapter 6, pp. 249–460. Athena Scientific.
- Bickel, P. J. and K. A. Doksum (2015). *Mathematical Statistics: Basic Ideas and Selected Topics, volume I*, Volume 117. CRC Press.
- Borwein, J. M. and A. S. Lewis (2010). *Convex analysis and nonlinear optimization: theory and examples* (2 ed.), Volume 3. Springer.
- Boyd, S. P. and L. Vandenberghe (2004). *Convex optimization* (Reprint. with corrections ed.). Cambridge [u.a.]: Cambridge University Press.
- Bunea, F., Y. She, and M. H. Wegkamp (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *The Annals of Statistics* 39(2), 1282–1309.
- Chandrasekaran, V., B. Recht, P. A. Parrilo, and A. S. Willsky (2012). The convex geometry of linear inverse problems. *Foundations of Computational Mathematics* 12(6), 805–849.
- Drusvyatskiy, D., S. A. Vavasis, and H. Wolkowicz (2015). Extreme point inequalities and geometry of the rank sparsity ball. *Mathematical Programming* 152(1-2), 521–544.

- Fazel, M. (2002). *Matrix rank minimization with applications*. Ph. D. thesis, Department of Electrical Engineering, Stanford University.
- Fazel, M., H. Hindi, and S. P. Boyd (2001). A rank minimization heuristic with application to minimum order system approximation. In *American Control Conference 2001, Proceedings of the*, Volume 6, pp. 4734–4739. IEEE.
- Hiriart-Urruty, J.-B. and C. Lemaréchal (1996). *Convex analysis and minimization algorithms II: Advanced theory and bundle methods*, Volume 306 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin.
- Ji, S. and J. Ye (2009). An accelerated gradient method for trace norm minimization. In *Proc. of the 26th annual international conference on machine learning*, pp. 457–464. ACM.
- Koltchinskii, V. (2011). *Oracle inequalities in empirical risk minimization and sparse recovery problems*, Volume 2033 of *Lecture notes in mathematics; 2033: Ecole d’Eté de Probabilités de Saint-Flour*. Berlin: Springer.
- Koltchinskii, V. and P. Rangel (2013). Low rank estimation of smooth kernels on graphs. *The Annals of Statistics* 41(2), 604–640.
- Lütkepohl, H. (2007). *New introduction to multiple time series analysis* (1, corr. 2. print. ed.). Berlin: Springer.
- Mishra, B., G. Meyer, F. Bach, and R. Sepulchre (2013). Low-rank optimization with trace norm penalty. *SIAM Journal on Optimization* 23(4), 2124–2149.
- Negahban, S., P. Ravikumar, M. J. Wainwright, and B. Yu (2012). A unified framework for the analysis of regularized M -estimators. *Statistical Science* 27(4), 538–557.
- Negahban, S. and M. J. Wainwright (2011a). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics* 39(2), 1069–1097.
- Negahban, S. and M. J. Wainwright (2011b). Supplementary Material for: Estimation of (near) low-rank matrices with noise and high-dimensional scaling. available online.
- Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*, Volume 87 of *Applied Optimization*. Springer Science & Business Media, New York.
- Pollard, D. (2002). *A user’s guide to measure theoretic probability*, Volume 8 of *Cambridge series in statistical and probabilistic mathematics*. Cambridge: Cambridge University Press.
- Pong, T. K., P. Tseng, S. Ji, and J. Ye (2010). Trace norm regularization: Reformulations, algorithms, and multi-task learning. *SIAM Journal on Optimization* 20(6), 3465–3489.
- Rockafellar, R. T. (1970). *Convex analysis*, Volume 28 of *Princeton mathematical series*. Princeton, NJ: Princeton University Press.
- Rudelson, M. and R. Vershynin (2013). Hanson-Wright inequality and sub-gaussian concentration. *Electron. Commun. Probab* 18(82), 1–9.
- Shumway, R. H. and D. S. Stoffer (2006). *Time series analysis and its applications with R-examples* (3 ed.). Springer texts in statistics. New York: Springer.
- So, W. (1990). Facial structures of Schatten p -norms. *Linear and Multilinear Algebra* 27(3), 207–212.
- Stock, J. H. and M. W. Watson (2011). Dynamic factor models. *Oxford Handbook of Economic Forecasting* 1, 35–59.
- Toh, K.-C. and S. Yun (2010). An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems. *Pacific Journal of Optimization* 6(615-640), 15.
- Vaart, A. W. v. d. and J. A. Wellner (2000). *Weak convergence and empirical processes* (Corr. 2. print. ed.). Springer series in statistics. New York: Springer.
- Vershynin, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In Y. C. Eldar and G. Kutyniok (Eds.), *Compressed Sensing*, pp. 210–268. Cambridge University Press.
- Yuan, M., A. Ekici, Z. Lu, and R. Monteiro (2007). Dimension reduction and coefficient estimation in multivariate linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69(3), 329–346.
- Ziętak, K. (1993). Subdifferentials, faces, and dual matrices. *Linear algebra and its applications* 185, 125–141.

Appendix

3.a. Subgaussian random variables

This section considers properties of the distribution $\mu_x = \mathbb{P} \circ x^{-1}$, given by $\mathcal{R}^1 \ni A \mapsto \mathbb{P}\{x \in A\}$, of a zero mean random variable x defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The measure μ_x is called *Gaussian* if either $\mu_x A = 1$ if and only if $0 \in A$ or it is absolutely continuous with respect to the Lebesgue measure ν on $(\mathbb{R}, \mathcal{R}^1)$ and with *density*, that is, Radon-Nikodym derivative,

*Gaussian
density*

$$\frac{d\mu_x}{d\nu}(y) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\left(\frac{y}{\sigma}\right)^2\right) \quad \text{for some } \sigma > 0.$$

In the latter case, the implied symmetry $\mu_x(-\infty, -t) = \mu_x(t, \infty)$ for all $t > 0$ leads to

$$\begin{aligned} \mathbb{P}\{|x| > t\} &= \frac{2}{\sqrt{2\pi\sigma^2}} \int_{(t, \infty)} \exp\left(-\frac{(y-t+t)^2}{2\sigma^2}\right) \nu(dy) \\ &\leq \frac{2}{\sqrt{2\pi\sigma^2}} \int_{(t, \infty)} \exp\left(-\frac{(y-t)^2}{2\sigma^2} - \frac{t^2}{2\sigma^2}\right) \nu(dy) \leq \exp\left(1 - \frac{t^2}{2\sigma^2}\right). \end{aligned}$$

The comparison between the leftmost and rightmost term is also valid in the former case with arbitrary $\sigma > 0$. More generally, μ_x is called *subgaussian* if there exists a positive number s such that the inequality $\mu_x((-\infty, t) \cup (t, \infty)) = \mathbb{P}\{|x| > t\} \leq \exp(1 - t^2/s^2)$ holds for all $t > 0$. Alternatively, these inequalities hold if and only if there exists $\bar{s} > 0$ with $\mathbb{E}\psi_2(x/\bar{s}) \leq 1$, wherein $\psi_2(z) = \exp(z^2) - 1$. In fact, if μ_x is subgaussian, then Tonelli's theorem implies

subgaussian

$$\begin{aligned} \mathbb{E}\psi_2\left(\frac{x}{\bar{s}}\right) &= \int \nu[0, \psi_2(\frac{y}{\bar{s}})) \mu_x(dy) = \int_{[0, \infty)} \int 1_{\{(y', t') \mid \psi_2(\frac{y'}{\bar{s}}) > t'\}}(y, t) \mu_x(dy) \nu(dt) \\ &= \int_{[0, \infty)} \mathbb{P}\{\psi_2(x/\bar{s}) > t\} \nu(dt) = \int_{[0, \infty)} \mathbb{P}\{|x| > \bar{s} \log^{1/2}(1+t)\} \nu(dt) \leq e \int_0^\infty e^{-\frac{\bar{s}^2}{s^2} \log(1+t)} dt, \end{aligned}$$

wherein $e = \exp(1) < 3$ represents Euler's number and $\int_a^b \dots dt$ indicates Riemann integration. The final term is bounded above by $e \int_1^\infty t^{-(\bar{s}/s)^2} dt \leq e/3 < 1$ whenever $\bar{s} \geq 2s$. Conversely, if $\mathbb{E}\psi_2(x/\bar{s}) \leq 1$ for some positive \bar{s} , then one has the (in)equalities

$$\mathbb{P}\{|x| > t\} = \mathbb{P}\{\psi_2(\frac{x}{\bar{s}}) > \psi_2(\frac{t}{\bar{s}})\} \leq \frac{1}{\psi_2(t/\bar{s})} \mathbb{E}\psi_2(\frac{t}{\bar{s}}) 1_{\{\psi_2(\frac{\bullet}{\bar{s}}) > \psi_2(\frac{t}{\bar{s}})\}}(x) \leq \frac{1}{\psi_2(t/\bar{s})}.$$

Therein, the comparison between the leftmost and rightmost term is uninformative unless $t \geq \bar{s} \log^{1/2} 2$. If the latter inequality holds, then

$$\mathbb{P}\{|x| > t\} = \left[\frac{1}{\exp(\frac{t^2}{2\bar{s}^2}) - 1} \right] \left[\frac{1}{\exp(\frac{t^2}{2\bar{s}^2}) + 1} \right] \leq \frac{\exp(-t^2/(2\bar{s}^2))}{\sqrt{2} - 1} \leq \exp\left(1 - \frac{t^2}{(\sqrt{2}\bar{s})^2}\right),$$

wherein the final term exceeds one unless $t \geq \bar{s}\sqrt{2} > \bar{s} \log^{1/2} 2$ and therefore holds in general.

If the numbers $s, \bar{s} > 0$ satisfy the respective of the above requirements, then the same applies to all elements of (s, ∞) and (\bar{s}, ∞) , respectively. An unambiguous surrogate for \bar{s} comes in form of $\|x\|_{\psi_2} = \inf\{s > 0 \mid \mathbb{E}\psi_2(x/s) \leq 1\}$. The latter—defined for the random

variables on $(\Omega, \mathcal{F}, \mathbb{P})$ with subgaussian distribution—satisfies $\|ax\|_{\psi_2} = |a|\|x\|_{\psi_2}$ as well as $\|x + x'\|_{\psi_2} \leq \|x\|_{\psi_2} + \|x'\|_{\psi_2}$. The latter follows from convexity of ψ_2 —implied by the convexity of $x \mapsto x^2$ together with convexity of the monotone increasing function \exp —and

$$\mathbb{E}\psi_2\left(\frac{x + x'}{s + s'}\right) = \mathbb{E}\psi_2\left(\frac{s}{s + s'}\frac{x}{s} + \frac{s'}{s' + s}\frac{x'}{s'}\right).$$

In particular, the random variables with subgaussian distribution form a vector space (with addition and multiplication with numbers as defined in section 2.1.1). Furthermore, the inequality $\mathbb{E}\psi_2(x/\bar{s}) \geq \mathbb{E}(x/\bar{s})^2$ implies that $\|x\|_{\psi_2} = 0$ or equivalently $\mathbb{E}\psi_2(x/\bar{s}) \leq 1$ for all $\bar{s} > 0$ is possible only if $\mathbb{E}x^2 = 0$. Hence, the *subgaussian norm* $\|\cdot\|_{\psi_2}$ provides a norm on finite dimensional subspaces after adjusting the representation as in appendix 2.a if needed. subgaussian norm

Moreover, the above arguments show that there exists a number $D > 0$ such that $\mathbb{P}\{|x| > t\} \leq \exp(1 - (t/[D\|x\|_{\psi_2}])^2)$ for every random variable x with subgaussian distribution and $\|x\|_{\psi_2} > 0$. Consequently, the subgaussian norm controls the decay of $\mathbb{P}\{|x| > t\}$ as $t \rightarrow \infty$. The inequality $\mathbb{E}\psi_2(x) \geq \mathbb{E}x^2 + \mathbb{E}x^4/2 > \mathbb{E}x^2$ implies that $\|x\|_{\psi_2} \geq 1$ whenever $\mathbb{E}x^2 = 1$.

Lemma 3.14 contains the *Hanson-Wright inequality*, which plays an important role in verifying the assertions of section 3.5.3. Hanson-Wright inequality

Lemma 3.14. *If x_1, \dots, x_k , $k \in \mathbb{N}$, are independent random variables such that each x_i has zero mean as well as a subgaussian distribution with $\|x_i\|_{\psi_2} \leq C$ for some common upper bound $C > 0$, then there exists $\bar{C} > 0$ such that the random vector $x = (x_1, \dots, x_k)$ satisfies*

$$\mathbb{P}\{|x^\top Ax - \mathbb{E}x^\top Ax| > w\} \leq 2 \exp\left(-\bar{C} \min\left\{\frac{w^2}{C^4\|A\|^2}, \frac{w}{C^2\|A\|_{\text{op}}}\right\}\right)$$

for every $A \in \mathbb{R}^{k \times k}$ and $w > 0$. Here, $\bar{C} > 0$ does not depend on the distribution of (x_1, \dots, x_k) .

3.b. Proofs

Proof of lemma 3.1. Firstly, convexity of $\text{int } C$ guarantees the second inclusion in $\{\|x - \cdot\| \leq \varepsilon\} \subset \text{conv}\{x \pm \varepsilon\sqrt{k}y_i \mid i \leq k\} \subset \text{int } C$, wherein $\varepsilon \in (0, 1)$ is small enough to ensure $\{\|\cdot - x\| \leq \varepsilon\sqrt{k}\} \subset \text{int } C$; y_1, \dots, y_k symbolizes an orthonormal basis of W ; and $k = \dim W$. The first inclusion exploits the form of the second set: if the difference $z - x = \sum_{i \leq k} c_i y_i$ has squared length $\sum_{i \leq k} c_i^2 \leq \varepsilon^2$, then $z = \sum_{i \leq k} (|c_i|/s)(x + \text{sign}(c_i)sy_i)$. Herein, $s = \sum_{i \leq k} |c_i| = \langle (|c_1|, \dots, |c_k|), (1, \dots, 1) \rangle \leq \varepsilon\sqrt{k}$ by the Cauchy-Schwarz inequality, and $x \pm sy_i \in \text{conv}\{x \pm \varepsilon\sqrt{k}y_i\}$. Secondly, convexity of f implies boundedness of f on $\text{conv}\{x \pm \varepsilon\sqrt{k}y_i \mid i \leq k\}$ as well as $\sup_{\|y-x\| \leq c\varepsilon} f(y) \leq (1-c)f(x) + c \sup_{\|y-x\| \leq \varepsilon} f(y)$ for all $c \in [0, 1]$. Hence, if $\varepsilon \in (0, 1)$ is as above and $\varepsilon' = \varepsilon \min\{1, \delta\}/4K$ with $K > 1$ exceeding $\sup_{\|y-x\| \leq \varepsilon} f(y)$, then $\{\|(x, f(x) + \delta) - \cdot\| < \varepsilon'\} \subset \text{epi } f$, and $(x, f(x) + \delta)$ is therefore contained in an open subset of $\text{epi } f$. ■

Proof of lemma 3.2. If $d \in \partial\lambda\|0\|'$, then $\lambda\|x'\|' \geq \lambda\|0\|' + \langle d, x' - 0 \rangle = \langle d, x' \rangle$; thus, $\|d\|'_d \leq \lambda$. Conversely, if $\|d\|'_d \leq \lambda$, then $\lambda\|0\|' + \langle d, x' - 0 \rangle = \langle d, x' \rangle \leq \lambda\|x'\|'$. Subsequently, $\|x\| > 0$ is assumed. If $d \in \partial\lambda\|x\|'$ and $x' \in \{\|\cdot\|' \leq \|x\|'\}$, then x is a boundary point of the closed and convex set $\{\|\cdot\|' \leq \|x\|'\}$ with nonempty interior $\{\|\cdot\|' < \|x\|'\}$ and $0 \geq \lambda\|x'\|' - \lambda\|x\|' \geq \langle d, x' - x \rangle$. Furthermore, one has $\lambda = \lambda\|y\|' \geq \lambda\|x + y\|' - \lambda\|x\|' \geq \langle d, y \rangle$ for all $y \in \{\|\cdot\|' = 1\}$ and $0 = \lambda\|0\|' \geq \lambda\|x\|' - \langle d, x \rangle$. Thus, $\|d\|'_d = \lambda$. If $d \in \text{ncone}(\{\|\cdot\|' \leq \|x\|'\}, x)$, then

$\langle d, x \rangle \geq \langle d, x' \rangle$ for all $x' \in \{\|\cdot\|' = \|x\|'\}$ and therefore $\langle d, x \rangle = \|d\|'_d \|x\|'$. If, in addition to this equality, one has $\|d\|'_d = \lambda$, then the inequality $\lambda\|x\|' + \langle d, x' - x \rangle = \langle d, x' \rangle \leq \lambda\|x'\|'$ holds for all $x' \in W$, that is, $d \in \partial\lambda\|x\|'$. ■

Proof of lemma 3.3. If $S \in \{\|\cdot\|_{\text{op}} \leq 1\} \subset \mathbb{S}^m$ and $B \in \mathbb{S}^m$, then

$$\langle S, B \rangle - \lambda \text{rk } S \leq \sum_{i \leq \text{rk } S} (\sigma_i(S)\sigma_i(B) - \lambda) \leq \sum_{i \leq \text{rk } S} \sigma_i(S)(\sigma_i(B) - \lambda) \leq \sum_{i \leq \text{rk } B} \max\{\sigma_i(B) - \lambda, 0\}$$

by the von Neumann trace inequality, wherein $\sigma_i(\cdot)$ symbolizes the i -th singular value of its argument with the understanding that $\sigma_{\text{rk} \cdot + p}(\cdot) = 0$ for all $p \in \mathbb{N}$. Equality is attained in case $S = S' = \sum_{i \in \{\sigma_i(B) > \lambda\}} u_i \langle v_i, \cdot \rangle$, wherein u_i, v_i denote an i -th singular vector pair of B with $u_i = \pm v_i$. An analogous argument shows that

$$\langle S, B \rangle - \lambda \|S\|_{\text{nuc}} \leq \sum_{i \leq \text{rk } S} \sigma_i(S)(\sigma_i(B) - \lambda) \leq \sum_{i \leq \text{rk } B} \max\{\sigma_i(B) - \lambda, 0\} \quad \langle \text{A3.1} \rangle$$

with equality for $S = S'$. The convexity of g' guarantees $\partial g'(S) \neq \emptyset$ and therefore $g'(S) = g'_{\star\star}(S)$ for all $S \in \text{int } H = \{\|\cdot\|_{\text{op}} < 1\}$. If S is a boundary point of H , then

$$g'_{\star\star}(S) = 2 \left[\frac{1}{2} g'_{\star\star}(S) + \frac{1}{2} g'_{\star\star}(0) \right] \geq 2g'_{\star\star}(0/2 + S/2) = \lambda \|S\|_{\text{nuc}} \geq g'_{\star\star}(S),$$

wherein 0 denotes the $m \times m$ zero matrix. In summary, one has $h'_{\star\star} = g'_{\star\star} = g'$. ■

Proof of lemma 3.4. The quantity $v(Z)$ is attained at some $\Theta \in \mathbb{S}^m$ for all $Z \in W^{\times m}$ as the bracketed term shares its form with the least-squares criterion l_λ in $\langle \text{3.1} \rangle$. In particular, the function $v : W^{\times m} \rightarrow \mathbb{R}$ is well-defined. If $(Z, t), (Z', t') \in \text{epi } v$, and $v(Z) \leq t, v(Z') \leq t'$ are attained at $\Theta, \Theta' \in \mathbb{S}^m$, respectively, then for $c \in [0, 1]$ one has

$$\begin{aligned} v(cZ + (1-c)Z') &\leq f(X(c\Theta + (1-c)\Theta') + cZ + (1-c)Z') + g(c\Theta + (1-c)\Theta') \\ &\leq c[f(X\Theta + Z) + g(\Theta)] + (1-c)[f(X\Theta' + Z') + g(\Theta')] \leq ct + (1-c)t'. \end{aligned}$$

To derive the conjugate g_\star , let $D \in \mathbb{S}^m$ be nonzero and with spectral decomposition $D = \sum_{i \leq \text{rk } D} s_i \sigma_i v_i \langle v_i, \cdot \rangle$, $s_i \in \{\pm 1\}$. Then the $\|\cdot\|_{\text{nuc}}/\|\cdot\|_{\text{op}}$ -duality implies $\langle D, \Theta \rangle - \lambda \|\Theta\|_{\text{nuc}} \leq \|\Theta\|_{\text{nuc}}(\|D\|_{\text{op}} - \lambda)$. Equality is achieved by setting $\Theta = s_1 \eta v_1 \langle v_1, \cdot \rangle$, which leads to the equality $\langle D, \Theta \rangle - \lambda \|\Theta\|_{\text{nuc}} = \eta(\|D\|_{\text{op}} - \lambda)$ for every $\eta > 0$. The optimal choice is given by $\eta = 0$ if $\|D\|_{\text{op}} \leq \lambda$ and leads to $g_\star(D) = 0$; otherwise, no finite value of η is optimal, and therefore the quantity is unbounded. The conjugate f_\star follows from

$$\begin{aligned} \langle M, Z \rangle - \frac{1}{2\bar{\mu}} \|Y - Z\|^2 &= -\frac{\bar{\mu}}{2} \|Y/\bar{\mu}\|^2 + [\langle Y/\bar{\mu} + M, Z \rangle - \frac{1}{2\bar{\mu}} \|Z\|^2] \\ &\leq -\frac{\bar{\mu}}{2} \|Y/\bar{\mu}\|^2 + [\|Y/\bar{\mu} + M\| \|Z\| - \frac{1}{2\bar{\mu}} \|Z\|^2] \leq \frac{\bar{\mu}}{2} [\|Y/\bar{\mu} + M\|^2 - \|Y/\bar{\mu}\|^2], \end{aligned}$$

wherein the final inequality results from maximization with respect to $\|Z\|$. Equality is attained by choosing $Z = Y + \bar{\mu}M$. Finally, the conjugate of v derives from

$$\begin{aligned} \langle M, Z \rangle - v(Z) &= \sup_{\Theta \in \mathbb{S}^m} [\langle M, Z + X\Theta \rangle - f(Z + X\Theta)] + [\langle -M, X\Theta \rangle - g(\Theta)] \\ &\leq f_*(M) + \sup_{\Theta \in \mathbb{S}^m} [\langle -M, X\Theta \rangle - g(\Theta)] \leq f_*(M) + g_*\left(-\frac{\langle X, M \rangle + \langle M, X \rangle}{2}\right), \quad \langle \text{A3.2} \rangle \end{aligned}$$

wherein the first inequality follows by maximization over $Z \in W^{\times m}$. Moreover, $\langle -M, X\Theta \rangle = -\text{tr} \langle M, X\Theta \rangle = -\text{tr}(\langle X, M \rangle^\top \Theta) = -\langle \langle X, M \rangle, \Theta \rangle = \langle -(\langle X, M \rangle + \langle M, X \rangle)/2, \Theta \rangle$ due to $\Theta \in \mathbb{S}^m$ and $(\langle X, M \rangle + \langle M, X \rangle)/2$ being the orthogonal projection of $\langle X, M \rangle \in \mathbb{R}^{m \times m}$ onto \mathbb{S}^m . The implied inclusion $-(\langle X, M \rangle + \langle M, X \rangle)/2 \in \mathbb{S}^m$ leads to the second inequality. Finally, the rightmost term in $\langle \text{A3.2} \rangle$ is the least upper bound of $\{\langle M, Z \rangle - v(Z) \mid Z \in W^{\times m}\}$. ■

Proof of lemma 3.5. If the common $\|\cdot\|_{\text{nuc}}$ -length $\hat{\ell}$ of all minimizers equals zero, then uniqueness is uncontroversial. Otherwise, $\|\hat{\Theta}\|_{\text{nuc}} = \hat{\ell} > 0$. In this case, the lower part of the left-hand side of $\langle 3.4 \rangle$ together with lemma 3.2 implies that a minimizer $\hat{\Theta}$ satisfies $\langle \hat{\Theta}, \hat{G} \rangle = \|\hat{G}\|_{\text{op}} \|\hat{\Theta}\|_{\text{nuc}}$ and, in addition, $\|G\|_{\text{op}} = \lambda$. Then, corollary 2.6 guarantees the existence of two symmetric and positive semidefinite matrices S^+ , S^- with $\text{tr} S^+ + \text{tr} S^- = \hat{\ell}$ and $\hat{\Theta} = U_1^+ S^+ \langle U_1^+, \cdot \rangle - U_1^- S^- \langle U_1^-, \cdot \rangle$, wherein the columns of U_1^+ and U_1^- provide orthonormal bases of $\ker(\hat{G} \mp \lambda \text{id})$, respectively, which are independent of the choice of the minimizer $\hat{\Theta}$. Moreover, every minimizer $\hat{\Theta}$ satisfies the equality $X\hat{\Theta} = \bar{\mu}\hat{D} + Y$. The latter implies $XU_1^+ S^+ = (\bar{\mu}\hat{D} + Y)U_1^+$ and $XU_1^- S^- = -(\bar{\mu}\hat{D} + Y)U_1^-$ as all entries of $\langle U_1^-, U_1^+ \rangle$ equal zero. If $\ker(\hat{G} \mp \lambda \text{id}) \cap \ker X = \{0\}$, then $\ker XU_1^\pm = \{0\}$ and S^\pm are uniquely determined. ■

Proof of proposition 3.6. Minimizers $\hat{\Theta}$ of $\langle 3.1 \rangle$ and maximizers \hat{D} of $\langle 3.3 \rangle$ generally exist. If the common $\|\cdot\|_{\text{nuc}}$ -length $\hat{\ell}$ of all minimizers equals zero, then the lefthand side of $\langle 3.4 \rangle$ implies $\hat{D} = -Y/\bar{\mu}$ and $\|(\langle X, Y \rangle + \langle Y, X \rangle)/(2\bar{\mu})\|_{\text{op}} = \|\hat{G}\|_{\text{op}} \leq \lambda$. Conversely, if $\|(\langle X, Y \rangle + \langle Y, X \rangle)/(2\bar{\mu})\|_{\text{op}} \leq \lambda$, then $(\hat{\Theta}, \hat{D}) = (0, -Y/\bar{\mu})$ satisfies the lefthand side of $\langle 3.4 \rangle$. Hence, the latter choice is an optimal pair and therefore $\hat{\ell} = 0$. If $\hat{\ell} > 0$, then the uniqueness assertion follows from lemma 3.5 and the subsequent discussion. ■

Proof of proposition 3.7. Using $\langle 3.6 \rangle$ with $\Theta = \hat{\Theta}$ and $\alpha = \bar{\mu}/\|X\|_{\text{op}}^2$ yields

$$\begin{aligned} l_\lambda(\hat{\Theta}) - l_\lambda(\hat{\Theta}^{(j)}) &\geq \frac{1}{\alpha} \|\hat{\Theta}^{(j)} - \hat{\Theta}\|^2 - \frac{1}{2\alpha} [\|\hat{\Theta}^{(j)} - \hat{\Theta}^{(j-1)}\|^2 - 2\langle \hat{\Theta}^{(j-1)} - \hat{\Theta}, \hat{\Theta} - \hat{\Theta}^{(j)} \rangle] \\ &= \frac{1}{2\alpha} [\|\hat{\Theta}^{(j)} - \hat{\Theta}\|^2 - \|\hat{\Theta}^{(j-1)} - \hat{\Theta}\|^2]. \end{aligned}$$

Furthermore, the inequality $l_\lambda(\hat{\Theta}^{(i)}) \leq l_\lambda(\hat{\Theta}^{(i-1)})$ for all $i \geq 1$ leads to

$$l_\lambda(\hat{\Theta}^{(k)}) - l_\lambda(\hat{\Theta}) \leq \frac{1}{k} \sum_{j \leq k} l_\lambda(\hat{\Theta}^{(j)}) - l_\lambda(\hat{\Theta}) \leq \frac{1}{2\alpha k} [\|\hat{\Theta}^{(0)} - \hat{\Theta}\|^2 - \|\hat{\Theta}^{(k)} - \hat{\Theta}\|^2].$$

Consequently, $p_k = l_\lambda(\hat{\Theta}^{(k)}) \downarrow l_\lambda(\hat{\Theta})$. If $\delta > 0$ and $n \in \mathbb{N}$ is sufficiently large, then $\{l_\lambda \leq n\} \cap \{\|\hat{\Theta} - \cdot\| \geq \delta\}$ is nonempty and compact. The continuous function l_λ attains its minimum l_λ^δ on

this compact set. Thus, $l_\lambda(\widehat{\Theta}^{(k)}) - l_\lambda(\widehat{\Theta}) < l_\lambda^\delta - l_\lambda(\widehat{\Theta})$ implies $l_\lambda(\widehat{\Theta}^{(k)}) \leq n$, $\|\widehat{\Theta}^{(k)} - \widehat{\Theta}\| < \delta$, and therefore $\widehat{\Theta}^{(k)} \rightarrow \widehat{\Theta}$. Thus, $\widehat{D}^{(k)} \rightarrow \widehat{D} = (X\widehat{\Theta} - Y)/\bar{\mu}$ and $\widehat{G}^{(k)} \rightarrow \widehat{G} = -(\langle X, \widehat{D} \rangle + \langle \widehat{D}, X \rangle)/2$. If $\|\widehat{G}\|_{\text{op}} > 0$, then $\lambda\widehat{D}^{(k)}/\|\widehat{G}^{(k)}\|_{\text{op}} \rightarrow \widehat{D}$. If $\|\widehat{G}\|_{\text{op}} = 0$, then $0 \in \partial\lambda\|\widehat{\Theta}\|_{\text{nuc}}$, that is, $\widehat{\Theta} = 0$. Consequently, continuity of $-f_*$ implies $d_k \uparrow -f_*(\widehat{D}) = l_\lambda(\widehat{\Theta})$. \blacksquare

Proof of lemma 3.8. Rearranging terms in the inequality $\bar{l}_{\lambda,\xi}(\widehat{\Delta}) \leq \bar{l}_{\lambda,\xi}(0)$, that is,

$$\frac{\|\bar{E} - X\widehat{\Delta}\|^2}{2(n-1)} + \lambda\|\Theta_* + \widehat{\Delta}\|_{\text{nuc}} + \xi\|\Lambda^{q/2}(\Theta_* + \widehat{\Delta})\|^2 \leq \frac{\|\bar{E}\|^2}{2(n-1)} + \lambda\|\Theta_*\|_{\text{nuc}} + \xi\|\Lambda^{q/2}\Theta_*\|^2,$$

leads to an upper bound on $\|X\widehat{\Delta}\|^2/(2(n-1)) + \xi\|\Lambda^{q/2}\widehat{\Delta}\|^2$ in the form of

$$\frac{\langle \bar{E}, X\widehat{\Delta} \rangle}{n-1} + 2(\sqrt{2\xi}\|\Lambda^{q/2}\Theta_*\|) \left(\sqrt{\frac{\xi}{2}}\|\Lambda^{q/2}\widehat{\Delta}\| \right) + \lambda(\|\Theta_*\|_{\text{nuc}} - \|\Theta_* + \widehat{\Delta}\|_{\text{nuc}}).$$

The elementary inequality $2ab \leq a^2 + b^2$ for $a, b \in \mathbb{R}$, the (in)equalities $\langle \bar{E}, X\widehat{\Delta} \rangle/(n-1) = \langle X^T \bar{E}, \widehat{\Delta} \rangle/(n-1) = \langle G, \widehat{\Delta} \rangle \leq \|G\|_{\text{op}}\|\widehat{\Delta}\|_{\text{nuc}}$ implied by the symmetry of $\widehat{\Delta}$ and $\|\cdot\|_{\text{op}}/\|\cdot\|_{\text{nuc}}$ -duality, and $\|\Theta_* + \widehat{\Delta}\|_{\text{nuc}} - \|\Theta_*\|_{\text{nuc}} \geq \langle M_*, \widehat{\Delta} \rangle$ for any subgradient $M_* \in \partial\|\Theta_*\|_{\text{nuc}}$ lead to

$$\frac{1}{2}\left\|\frac{X\widehat{\Delta}}{\sqrt{n-1}}\right\|^2 + \frac{\xi}{2}\|\Lambda^{q/2}\widehat{\Delta}\|^2 \leq \|G\|_{\text{op}}(\|P_{\widehat{V}_k}\widehat{\Delta}\|_{\text{nuc}} + \|P_{\widehat{V}_k^\perp}\widehat{\Delta}\|_{\text{nuc}}) + \lambda\langle -M_*, \widehat{\Delta} \rangle + 2\xi\|\Lambda^{q/2}\Theta_*\|^2.$$

A more specific result follows by setting $M_* = \text{sign } \Theta_* + P_{\widehat{V}_k^\perp} \text{sign}(P_{\widehat{V}_k^\perp}\widehat{\Delta})$, wherein $\text{sign } S = \sum_{i \leq \text{rk } S} s_i v_i \langle v_i, \cdot \rangle$ is (uniquely) defined for every symmetric matrix S via a spectral decomposition $\sum_{i \leq \text{rk } S} s_i \sigma_i(S) v_i \langle v_i, \cdot \rangle$. In particular, $\langle \text{sign } S, S \rangle = \sum_{i,j \leq \text{rk } S} s_i s_j \sigma_i(S) \langle v_i, v_j \rangle^2 = \sum_{i \leq \text{rk } S} \sigma_i(S) = \|S\|_{\text{nuc}}$. Herein, the inclusion $M_* \in \partial\|\Theta_*\|_{\text{nuc}}$ follows from lemma 3.2. In fact, the perfect alignment of Θ_* with respect to Λ implies $\langle M_*, \Theta_* \rangle = \|\Theta_*\|_{\text{nuc}}$ as well as $\|M_* x\|^2 \leq \|\text{sign } \Theta_* x\|^2 + \|P_{L_k^\perp} x\|^2 \leq 1$ for all unit length $x \in \mathbb{R}^m$. If $h > 0$, then the latter inequality holds with equality for $x = u_1$; thus, $\|M_*\|_{\text{op}} = 1$ and $\langle M_*, \Theta_* \rangle = \|M_*\|_{\text{op}}\|\Theta_*\|_{\text{nuc}}$. If $h = 0$, then $\|M_*\|_{\text{op}} \leq 1$. Using $\lambda \geq \|G\|_{\text{op}}$ and $\langle P_{\widehat{V}_k^\perp} \text{sign}(P_{\widehat{V}_k^\perp}\widehat{\Delta}), \widehat{\Delta} \rangle = \|P_{\widehat{V}_k^\perp}\widehat{\Delta}\|_{\text{nuc}}$ leads to

$$\left\|\frac{X\widehat{\Delta}}{\sqrt{n-1}}\right\|^2 + \xi\|\Lambda^{q/2}\widehat{\Delta}\|^2 \leq 2\lambda\|P_{\widehat{V}_k}\widehat{\Delta}\|_{\text{nuc}} + 2\lambda\langle -\text{sign } \Theta_*, \widehat{\Delta} \rangle + 4\xi\|\Lambda^{q/2}\Theta_*\|^2. \quad \langle \text{A3.3} \rangle$$

Next, $\langle -\text{sign } \Theta_*, \widehat{\Delta} \rangle = \langle -\text{sign } \Theta_*, P_{L_k}\widehat{\Delta} \rangle \leq \|\text{sign } \Theta_*\|_{\text{op}}\|P_{L_k}\widehat{\Delta}\|_{\text{nuc}} \leq \sqrt{h}\|\widehat{\Delta}\|$, which relies on the final display of section 2.5.2, $\dim L_k = h$ by perfect alignment, and $\|P_{L_k}\|_{\text{op}} = 1$. The inequality $\text{rk } P_{\widehat{V}_k}\widehat{\Delta} \leq 2h$ justified at the end of section 3.5.1 implies $2\lambda\|P_{\widehat{V}_k}\widehat{\Delta}\|_{\text{nuc}} \leq 2\lambda\sqrt{2h}\|P_{\widehat{V}_k}\widehat{\Delta}\| \leq 3\lambda\sqrt{h}\|\widehat{\Delta}\|$. Consequently, the righthand side in $\langle \text{A3.3} \rangle$ may be replaced by $5\lambda\sqrt{h}\|\widehat{\Delta}\| + 4\xi\|\Lambda^{q/2}\Theta_*\|^2$. Finally, an orthonormal basis o_1, \dots, o_m of \mathbb{R}^m obtained as an extension of a sequence of singular vectors $o_1, \dots, o_{\text{rk } \Lambda}$ of Λ induces an orthonormal basis $O_{i,j} = o_i o_j^T$, $i, j \leq m$, of $\mathbb{R}^{m \times m}$ —comparable to e_i and $B_{i,j}$ in examples (a) and (b) of section 2.1.1. This orthonormal basis reveals the inequality

$$\|\Lambda^{q/2}\widehat{\Delta}\|^2 = \left\| \sum_{i,j \leq m} \langle O_{i,j}, \widehat{\Delta} \rangle \Lambda^{q/2} O_{i,j} \right\|^2 = \sum_{i,j \leq m} \langle O_{i,j}, \widehat{\Delta} \rangle^2 \sigma_i^q(\Lambda) \geq \sigma_{k-1}^q(\Lambda) \sum_{i < k, i \leq j \leq m} \langle O_{i,j}, \widehat{\Delta} \rangle^2.$$

In addition, symmetry of $\widehat{\Delta}$ implies $\langle O_{i,j}, \widehat{\Delta} \rangle = \langle O_{j,i}, \widehat{\Delta} \rangle$. Consequently,

$$\begin{aligned} \sum_{i < k, i \leq j \leq m} \langle O_{i,j}, \widehat{\Delta} \rangle^2 &\geq \frac{1}{2} \left(\sum_{i < k} \langle O_{i,i}, \widehat{\Delta} \rangle^2 + \sum_{\substack{i < k, \\ i < j \leq m}} \langle O_{i,j}, \widehat{\Delta} \rangle^2 + \sum_{\substack{j < k, \\ i < j \leq m}} \langle O_{i,j}, \widehat{\Delta} \rangle^2 \right) \\ &= \frac{1}{2} \left\| \sum_{i < k} \langle \bar{O}_{i,i}, \widehat{\Delta} \rangle \bar{O}_{i,i} + \sum_{\substack{i < k, \\ i < j \leq m}} \langle O_{i,j}, \widehat{\Delta} \rangle (O_{i,j} + O_{j,i}) \right\|^2 = \frac{1}{2} \left\| \sum_{\substack{i < k, \\ i < j \leq m}} \langle \bar{O}_{i,j}, \widehat{\Delta} \rangle \bar{O}_{i,j} \right\|^2, \end{aligned}$$

wherein the final term equals $\frac{1}{2} \|P_{V_k^\perp} \widehat{\Delta}\|^2$. Replacing $\xi \|\Lambda^{q/2} \widehat{\Delta}\|^2$ on the lefthand side of [<A3.3>](#) accordingly verifies [<3.18>](#). \blacksquare

Proof of proposition 3.9. If $h = 0$, then the requirement $\lambda \geq \|G\|_{\text{op}}$ ensures that $\|\widehat{\Delta}\| = 0$; thus, the inequality [<3.20>](#) holds. If $h > 0$ and $k = 1$, then $X = FU_1^\top$ and lemma 3.8 implies

$$\left\langle \frac{F^\top F}{n-1} U_1^\top \widehat{\Delta}, U_1^\top \widehat{\Delta} \right\rangle = \left\| \frac{X \widehat{\Delta}}{\sqrt{n-1}} \right\|^2 \leq 5\sqrt{h}\lambda \|\widehat{\Delta}\|.$$

Using a spectral decomposition $\sum_{i \leq h} \sigma_i (F^\top F / (n-1)) a_i \langle a_i, \cdot \rangle = F^\top F / (n-1)$, wherein a_1, \dots, a_h forms an orthonormal basis of \mathbb{R}^h with $h = m$, yields the inequalities

$$\left\langle \frac{F^\top F}{n-1} U_1^\top \widehat{\Delta}, U_1^\top \widehat{\Delta} \right\rangle = \sum_{i \leq h} \sigma_i \left(\frac{F^\top F}{n-1} \right) \|a_i a_i^\top U_1^\top \widehat{\Delta}\|^2 \geq \sigma_h \left(\frac{F^\top F}{n-1} \right) \|U_1^\top \widehat{\Delta}\|^2 \geq \kappa \|\widehat{\Delta}\|^2,$$

wherein $\langle a_i a_i^\top U_1^\top \widehat{\Delta}, U_1^\top \widehat{\Delta} \rangle = \|a_i^\top U_1^\top \widehat{\Delta}\|^2 = \|a_i a_i^\top U_1^\top \widehat{\Delta}\|^2$ as $\|a_i\| = 1$, and the inequalities rely on the orthogonality of $a_i a_i^\top U_1^\top \widehat{\Delta}$ and $a_j a_j^\top U_1^\top \widehat{\Delta}$ for $i \neq j$ as well as $\|U_1^\top \widehat{\Delta}\| = \|U_1 U_1^\top \widehat{\Delta}\| = \|\widehat{\Delta}\|$. Combining these two inequalities leads to the conclusion [<3.20>](#).

If $h > 0$ and $k > 1$, then using the representation in [<3.19>](#) leads to

$$\begin{aligned} \|X \widehat{\Delta}\|^2 &= \|FU_1^\top \widehat{\Delta} + \rho V_2 U_2^\top \widehat{\Delta}\|^2 = \|FU_1^\top \widehat{\Delta}\|^2 + \rho^2 \|V_2 U_2^\top \widehat{\Delta}\|^2 + 2\rho \langle V_2^\top F U_1^\top \widehat{\Delta}, U_2^\top \widehat{\Delta} \rangle \\ &\geq \|FU_1^\top \widehat{\Delta}\|^2 + \rho^2 \|V_2 U_2^\top \widehat{\Delta}\|^2 - 2 \left(\sqrt{2} \rho \frac{\|V_2^\top F\|_{\text{op}} \|U_2^\top \widehat{\Delta}\|}{\sigma_h^{1/2}(F^\top F)} \right) \left[\sigma_h^{1/2}(F^\top F) \frac{\|U_1^\top \widehat{\Delta}\|}{\sqrt{2}} \right] \\ &\geq \frac{\sigma_h(F^\top F)}{2} \|U_1^\top \widehat{\Delta}\|^2 - 2\rho^2 \frac{\|V_2^\top F\|_{\text{op}}^2}{\sigma_h(F^\top F)} \|U_2^\top \widehat{\Delta}\|^2, \end{aligned} \tag{A3.4}$$

wherein the second and third line rely on $\sigma_h(F^\top F) \geq \kappa > 0$. In addition, the second line uses

$$\langle V_2^\top F U_1^\top \widehat{\Delta}, U_2^\top \widehat{\Delta} \rangle \leq \sqrt{\|V_2^\top F U_1^\top \widehat{\Delta}\|^2 \|U_2^\top \widehat{\Delta}\|^2} = \sqrt{\sum_{i \leq \text{rk } V_2^\top F} \sigma_i^2(V_2^\top F) \|d_i f_i^\top U_1^\top \widehat{\Delta}\|^2} \|U_2^\top \widehat{\Delta}\|,$$

wherein the final term is upper bounded by $\sigma_1(V_2^\top F) \|U_1^\top \widehat{\Delta}\| \|U_2^\top \widehat{\Delta}\|$, and furthermore $V_2^\top F = \sum_{i \leq \text{rk } V_2^\top F} \sigma_i(V_2^\top F) d_i \langle f_i, \cdot \rangle$ represents a singular value decomposition of $V_2^\top F \in \mathbb{R}^{(m-h) \times h}$. The inequality [<A3.4>](#) follows by dropping $\rho^2 \|V_2 U_2^\top \widehat{\Delta}\|^2$, applying the elementary inequality $a^2 + b^2 \geq 2ab$ for $a, b \in \mathbb{R}$ to the final term of the previous line, and $\|FU_1^\top \widehat{\Delta}\|^2 \geq$

$\sigma_h(F^\top F)\|U_1^\top \widehat{\Delta}\|^2$, whose derivation is similar to the preceding display. The inequality

$$\|X\widehat{\Delta}\|^2 \geq \frac{\sigma_h(F^\top F)}{2}\|\widehat{\Delta}\|^2 - \left(\frac{\sigma_h(F^\top F)}{2} + 2\rho^2 \frac{\|V_2^\top F\|_{\text{op}}^2}{\sigma_h(F^\top F)} \right) \|P_{V_k^\perp} \widehat{\Delta}\|^2$$

results from $\|\widehat{\Delta}\|^2 = \|U_1 U_1^\top \widehat{\Delta} + U_2 U_2^\top \widehat{\Delta}\|^2 = \|U_1^\top \widehat{\Delta}\|^2 + \|U_2^\top \widehat{\Delta}\|^2$ and $\|U_2^\top \widehat{\Delta}\|^2 = \|U_2 U_2^\top \widehat{\Delta}\|^2 = \|P_{L_k^\perp} \widehat{\Delta}\|^2 = \|P_{L_k^\perp} \widehat{\Delta} P_{L_k^\perp}\|^2 + \|P_{L_k^\perp} \widehat{\Delta} P_{L_k}\|^2 \leq \|P_{V_k^\perp} \widehat{\Delta}\|^2$, which uses the representation in <3.17> and the perfect alignment of Θ_* with Λ , that is, $L_k = \text{img } U_1 = (\text{img } U_2)^\perp$. Plugging these results into <3.18> yields the inequality

$$\begin{aligned} \sigma_h \left(\frac{F^\top F}{n-1} \right) \|\widehat{\Delta}\|^2 + \left[\xi \sigma_{k-1}^q(\Lambda) - \left(\sigma_h \left(\frac{F^\top F}{n-1} \right) + 4\rho^2 \frac{\|V_2^\top F/(n-1)\|_{\text{op}}^2}{\sigma_h(F^\top F/(n-1))} \right) \right] \|P_{V_k^\perp} \widehat{\Delta}\|^2 \\ \leq 10\sqrt{h}\lambda \|\widehat{\Delta}\| + 8\xi \|\Lambda^{q/2} \Theta_*\|^2. \end{aligned}$$

Next, the lower bound on ξ , $\sigma_h(F^\top F/(n-1)) \geq \kappa$, and $a+b \leq 2\max\{a,b\}$ for $a, b \in \mathbb{R}$ imply the inequality $\kappa \|\widehat{\Delta}\|^2 \leq \max\{20\sqrt{h}\lambda \|\widehat{\Delta}\|, 16\xi \|\Lambda^{q/2} \Theta_*\|^2\}$. Consequently, at least one of $\|\widehat{\Delta}\| \leq 20\sqrt{h}\frac{\lambda}{\kappa}$ and $\|\widehat{\Delta}\| \leq 4\sqrt{\frac{\xi}{\kappa}} \|\Lambda^{q/2} \Theta_*\|$ holds. \blacksquare

Proof of lemma 3.10. Partitioning a unit length $a \in \mathbb{R}^{m(n+l)}$ in $n+l$ vectors $a_i \in \mathbb{R}^m$ leads to $\langle a, A_c a \rangle = \sum_{t \leq n-1} \langle a, B_t^\top c c^\top B_t a \rangle = \sum_{t \leq n-1} \langle c, \sum_{i=n-t+1}^{n-t+1+l} \bar{A}_{i-(n-t+1)} a_i \rangle^2$ for an arbitrary unit length $c \in \mathbb{R}^h$. This representation yields

$$\langle a, A_c a \rangle = \sum_{i=0}^l \sum_{j=0}^l \langle \bar{A}_i^\top c, \left(\sum_{t \leq n-1} a_{n-t+1+i} a_{n-t+1+j}^\top \right) \bar{A}_j^\top c \rangle \leq \left(\sum_{i=0}^l \|A_i c\| \right)^2,$$

wherein the final inequality follows from the Cauchy-Schwarz inequality together with

$$\begin{aligned} \left\| \sum_{t \leq n-1} a_{n-t+1+i} a_{n-t+1+j}^\top \right\|_{\text{op}}^2 &\leq \left\| \sum_{t \leq n-1} a_{n-t+1+i} a_{n-t+1+j}^\top \right\|^2 \\ &= \sum_{t \leq n-1} \sum_{s \leq n-1} \langle a_{n-t+1+i}, a_{n-s+1+i} \rangle \langle a_{n-t+1+j}, a_{n-s+1+j} \rangle \\ &\leq \left[\sum_{t \leq n-1} \|a_{n-t+1+i}\| \|a_{n-t+1+j}\| \right]^2 = \left\langle \begin{pmatrix} \|a_{n+i}\| \\ \vdots \\ \|a_{2+i}\| \end{pmatrix}, \begin{pmatrix} \|a_{n+j}\| \\ \vdots \\ \|a_{2+j}\| \end{pmatrix} \right\rangle^2 \leq \left(\sum_{i \leq n+l} \|a_i\|^2 \right)^2, \end{aligned}$$

wherein the final term does not exceed one, and $\|\bar{A}_i^\top c\| = \|A_i c\|$ —as A_i is a diagonal matrix and therefore symmetric—for $0 \leq i \leq l$. In summary, the symmetric and positive semidefinite matrix A_c satisfies $\|A_c\|_{\text{op}} \leq \left(\sum_{i=0}^l \|A_i c\| \right)^2$. Thus, $\langle V_f c, c \rangle = \sum_{i=0}^l \|A_i c\|^2$ implies

$$\zeta_c = \frac{1}{2C^2} \frac{\sum_{i=0}^l \|A_i c\|^2}{\left(\sum_{i=0}^l \|A_i c\| \right)^2} \geq \frac{1}{2C^2} \frac{\sum_{i=0}^{k'} \|A_i c\|^2}{\left(\frac{6}{5} \sum_{i=0}^{k'} \|A_i c\| \right)^2} \geq \frac{1}{3C^2(k'+1)}, \quad \text{<A3.5>}$$

wherein k' equals the largest nonnegative integer not exceeding $\min\{l, (1 + \log 6)/\alpha\}$, which is

therefore less than $k' + 1$. Thus, if $k' < l$, then $k' + 1 > (1 + \log 6)/\alpha$ and

$$\begin{aligned} \sum_{i=0}^l \|A_i c\| &= \sum_{i=0}^{k'} \|A_i c\| + \sum_{i=k'+1}^l \|A_i c\| \leq \sum_{i=0}^{k'} \|A_i c\| + \exp(1 - \alpha(k' + 1)) \sum_{i=0}^l \|A_i c\| \\ &\leq \sum_{i=0}^{k'} \|A_i c\| + \exp(-\log 6) \sum_{i=0}^l \|A_i c\| = \sum_{i=0}^{k'} \|A_i c\| + \frac{1}{6} \sum_{i=0}^l \|A_i c\|. \end{aligned}$$

Consequently, $\sum_{i=0}^l \|A_i c\| \leq \frac{6}{5} \sum_{i=0}^{k'} \|A_i c\| \leq \frac{6}{5} \sqrt{k'+1} (\sum_{i=0}^{k'} \|A_i c\|^2)^{1/2}$, wherein the final inequality follows as in the final display of section 2.5.2. The latter inequalities also hold for the case $k' = l$ and explain the first and second inequality in <A3.5>. The inequality $(3 + \alpha)/\alpha \geq (1 + \log 6 + \alpha)/\alpha \geq k' + 1$ holds in general. Finally, $C \geq 1$ implies that the resulting $\bar{\zeta} = \alpha/(3C^2(3 + \alpha))$ lies in $[0, 1]$. \blacksquare

Proof of lemma 3.11. Let c_1, \dots, c_q and d_1, \dots, d_p be \subset -minimal $\frac{1}{4}$ -nets of the unit sphere $\{\|\cdot\| = 1\}$ of \mathbb{R}^h and \mathbb{R}^{m-h} , respectively. Then $\bar{H} = \{(c_i, d_j) \in \mathbb{R}^m \mid i \leq q, j \leq p\}$ contains no more than $9^h 9^{m-h} = 9^m$ elements. In addition, if $(c, d) \in H = \{(c', d') \in \mathbb{R}^{h+(m-h)} \mid \|c'\| = \|d'\| = 1\}$, then there exist $i \leq q, j \leq p$ such that $\|(c, d) - (c_i, d_j)\| = (\|c - c_i\|^2 + \|d - d_j\|^2)^{1/2} \leq 1/\sqrt{8}$. Consequently, \bar{H} provides an $\frac{1}{2\sqrt{2}}$ -net of H . In addition, the map $H \ni (c, d) \mapsto \langle \frac{V_2^\top F}{n-1} c, d \rangle$, wherein $F = F(\omega)$, $V_2 = V_2(\omega)$ represent images of some ω , satisfies

$$\left| \left\langle \frac{V_2^\top F}{n-1} c, d \right\rangle - \left\langle \frac{V_2^\top F}{n-1} c', d' \right\rangle \right| \leq \left\| \frac{V_2^\top F}{n-1} \right\|_{\text{op}} [\|c - c'\| + \|d - d'\|] \leq \sqrt{2} \left\| \frac{V_2^\top F}{n-1} \right\|_{\text{op}} \left\| \begin{pmatrix} c \\ d \end{pmatrix} - \begin{pmatrix} c' \\ d' \end{pmatrix} \right\|,$$

which utilizes $\|V_2^\top F/(n-1)\|_{\text{op}} = \|F^\top V_2/(n-1)\|_{\text{op}}$ as well as $a + b = \langle \begin{pmatrix} a \\ b \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \rangle \leq \sqrt{2} \sqrt{a^2 + b^2}$ for $a, b \in \mathbb{R}$. Consequently, the second inequality in <2.1> implies the inequality

$$\left\| \frac{V_2^\top F}{n-1} \right\|_{\text{op}} = \sup_{(c,d) \in H} \left\langle \frac{V_2^\top F}{n-1} c, d \right\rangle \leq \max_{i \leq q, j \leq p} \left\langle \frac{V_2^\top F}{n-1} c_i, d_j \right\rangle + \sqrt{2} \left\| \frac{V_2^\top F}{n-1} \right\|_{\text{op}} \frac{1}{2\sqrt{2}},$$

which in turn implies $\|V_2^\top F/(n-1)\|_{\text{op}} \leq 2 \max_{i,j} \langle \frac{V_2^\top F}{n-1} c_i, d_j \rangle$. Subsequently, F represents $\omega \mapsto F(\omega)$ and likewise for V_2 . For given $(c, d) \in H$, it holds that $\langle V_2^\top F c, d \rangle = \sum_{i \leq n-1} \langle d, v_i^{(2)} \rangle \langle c, f_i \rangle$, wherein the factors of the summands are independent. In particular, $\mathbb{E} \langle V_2^\top F c, d \rangle = 0$. Using

$$v_t^{(2)} = \bar{B}_t v = \begin{bmatrix} \underbrace{0 \dots 0}_{\substack{n-t \text{ zero matrices} \\ \text{in } \mathbb{R}^{m-h \times m}}} & \underbrace{0}_{\substack{\text{zero matrix} \\ \text{in } \mathbb{R}^{m-h \times h}}} & \underbrace{I}_{\substack{\text{identity matrix} \\ \text{in } \mathbb{R}^{m-h \times m-h}}} & \underbrace{0 \dots 0}_{\substack{t+l-1 \text{ zero matrices} \\ \text{in } \mathbb{R}^{m-h \times m}}} \end{bmatrix} v \quad \text{<A3.6>}$$

and the notation of the main text leads to $\langle V_2^\top F c, d \rangle = \sum_{t \leq n-1} \langle \bar{B}_t v, d \rangle \langle c, B_t v \rangle = \langle v, A_{c,d} v \rangle$, wherein $\sum_{t \leq n-1} \bar{B}_t^\top d c^\top B_t$. The squared Frobenius norm of the latter matrix equals

$$\|A_{c,d}\|^2 = \text{tr} \left(\sum_{t \leq n-1} \sum_{s \leq n-1} B_t^\top c d^\top \bar{B}_t \bar{B}_s^\top d c^\top B_s \right) = \sum_{t \leq n-1} \langle c, B_t B_t^\top c \rangle = (n-1) \langle V_f c, c \rangle,$$

wherein the second equality uses $d^\top \bar{B}_t \bar{B}_s^\top d = 0$ if $s \neq t$ and one otherwise. If $a, a' \in \mathbb{R}^{m(n+l)}$

are of unit length and each consisting of $n + l$ vectors $a_i, a'_i \in \mathbb{R}^m$, respectively, then

$$\begin{aligned} \langle a', \sum_{t \leq n-1} \bar{B}_t^\top d c^\top B_t a \rangle &= \sum_{t \leq n-1} \langle d, [0 \quad I] a'_{n-t+1} \rangle \langle c, \sum_{i=0}^l \bar{A}_i a_{i+n-t+1} \rangle \\ &= \sum_{i=0}^l d^\top [0 \quad I] \left(\sum_{t \leq n-1} a'_{n-t+1} a_{i+n-t+1}^\top \right) \bar{A}_i^\top c \leq \sum_{i=0}^l \|\bar{A}_i^\top c\| = \sum_{i=0}^l \|A_i c\|, \end{aligned}$$

wherein the inequality results from using the Cauchy-Schwarz inequality as exemplified in the proof of lemma 3.10, and $\|\bar{A}_i^\top c\| = \|A_i c\|$ holds by definition. In particular, one has $\sum_{i=0}^l \|A_i c\| \geq \|A_{c,d}\|_{\text{op}}$. Using the Hanson-Wright inequality with $w_{c_i} = \bar{C} C^2 \sum_{k=0}^l \|A_k c_i\| m$, wherein the constant \bar{C} is given by $\max\{\bar{C}^{-1} \log 18, 1\} \geq 1$, together with similar arguments as in the proof of lemma 3.10 and $m \geq n - 1$ reveals that for given $i \leq q$ and $j \leq p$ the probability

$$\mathbb{P} \left\{ \left| \langle d_j, \frac{V_2^\top F}{n-1} c_i \rangle \right| > \bar{C} C^2 \frac{m}{n-1} \sum_{k=0}^l \|A_k c_i\| \right\} = \mathbb{P} \left\{ |\langle v, A_{c_i, d_j} v \rangle| > \bar{C} C^2 m \sum_{k=0}^l \|A_k c_i\| \right\}$$

is less or equal to $2 \exp(-(\log 9 + \log 2)m)$. As a consequence, the inequality $\sum_{k=0}^l \|A_k c_i\| \leq \frac{6}{5} \left(\frac{3+\alpha}{\alpha}\right)^{1/2} \langle c_i, V_f c_i \rangle^{1/2} \leq \frac{6\sqrt{3}}{5} \left(\frac{1+\alpha}{\alpha}\right)^{1/2} \sigma_1^{1/2}(V_f)$, which uses $\sum_{k=0}^l \|A_k c_i\|^2 = \langle c_i, V_f c_i \rangle$ and is derived in the proof of lemma 3.10, together with an application of the union bound shows that

$$\left\| \frac{V_2(\omega)^\top F(\omega)}{n-1} \right\|_{\text{op}} \leq 2 \max_{i \leq q, j \leq p} \left\langle \frac{V_2(\omega)^\top F(\omega)}{n-1} c_i, d_j \right\rangle \leq \bar{C} C^2 \left(\frac{1+\alpha}{\alpha}\right)^{1/2} \sigma_1^{1/2}(V_f) \frac{m}{n-1},$$

wherein $\bar{C} = \frac{12\sqrt{3}}{5} \bar{C}$, holds for all ω in an element of the σ -field \mathcal{F} , whose probability is bounded below by $1 - 2 \exp(m \log 9 - (\log 9 + \log 2)m) = 1 - 1/2^{m-1}$. \blacksquare

Proof of lemma 3.12. The proof considers $0 < h < m$; the case $h \in \{0, m\}$ follows by similar arguments. Let c_1, \dots, c_q be a \subset -minimal $\frac{1}{4}$ -net of the unit sphere $\{\|\cdot\| = 1\}$ in \mathbb{R}^m . Then $\bar{H} = \{(c_i, \pm c_i) \mid i \leq q\}$ contains no more than 2×9^m elements. Moreover, if $(c, d) \in H = \{(c', d') \in \mathbb{R}^{m+m} \mid c' = \pm d', \|c'\| = \|d'\| = 1\}$, then there exists $i \leq q$ and a sign choice $s \in \{-1, 1\}$ such that $\|(c, d) - (c_i, s c_i)\| \leq 1/\sqrt{8}$. Consequently, \bar{H} provides an $\frac{1}{2\sqrt{2}}$ -net of H . An upper bound on the $\|\cdot\|$ -Lipschitz constant for the map $(c, d) \mapsto \langle Gc, d \rangle$, wherein $G = (X^\top \bar{E} + \bar{E}^\top X)/(2(n-1))$ and $X = X(\omega) \in \mathbb{R}^{n-1 \times m}$ as well as $\bar{E} = \bar{E}(\omega) \in \mathbb{R}^{n-1 \times m}$ represent images of some ω , follows from $|\langle Gc, d \rangle - \langle Gc', d' \rangle| \leq \|G\|_{\text{op}} (\|c - c'\| + \|d - d'\|) \leq \sqrt{2} \|G\|_{\text{op}} \|(c, d) - (c', d')\|$, which relies on $\left\langle \begin{pmatrix} a \\ b \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\rangle \leq \sqrt{2} \sqrt{a^2 + b^2}$ for $a, b \in \mathbb{R}$. Consequently, lemma 2.4 and the second inequality in <2.1> imply $\|G\|_{\text{op}} = \sup_{(c,d) \in H} \langle Gc, d \rangle \leq 2 \max_{i \leq q, s \in \{-1, 1\}} \langle Gc_i, s c_i \rangle$.

Below, X and \bar{E} symbolize the random matrices $\omega \mapsto X(\omega)$ and $\omega \mapsto \bar{E}(\omega)$, respectively. The (transposed) rows of the former equal $x_t = U_1 f_t + \rho U_2 v_t^{(2)} = (U_1 B_t + \rho U_2 \bar{B}_t) v$, $t \leq n - 1$, wherein v, B_t and \bar{B}_t are defined in the main text and <A3.6>, respectively. The (transposed) rows of the latter are given by $\bar{e}_{t+1} = U_1 A_0 v_{t+1}^{(1)} + \rho U_2 v_{t+1}^{(2)} + \sum_{i \leq t} U_1 (A_i - \Gamma_* A_{i-1}) v_{t+1-i}^{(1)} - U_1 \Gamma_* A_i v_{t-i}^{(1)}$, wherein the representation follows from <3.12>, and Γ_* is defined in <3.11>.

The latter amounts to $\bar{e}_{t+1} = (U_1 \bar{B}_{t+1} + \rho U_2 \bar{B}_{t+1})v$, wherein \bar{B}_{t+1} is defined in <A3.6> and

$$\mathbb{R}^{h \times m(n+l)} \ni \bar{B}_{t+1} = \left[\underbrace{\begin{matrix} 0 & \dots & 0 \\ n-t-1 \text{ zero} \\ \text{matrices in } \mathbb{R}^{h \times m} \end{matrix}} & \bar{A}_0 & \dots & \bar{A}_{l+1} & \underbrace{\begin{matrix} 0 & \dots & 0 \\ t-1 \text{ zero} \\ \text{matrices in } \mathbb{R}^{h \times m} \end{matrix}} \right], \quad t \leq n-1,$$

with $\bar{A}_0 = [A_0 \ 0] \in \mathbb{R}^{h \times h+(m-h)}$, $\bar{A}_i = [A_i - \Gamma_* A_{i-1} \ 0]$, $i \leq l$, and $\bar{A}_{l+1} = [-\Gamma_* A_l \ 0]$. For given $(c, d) \in H$, one has $\mathbb{E}\langle Gc, d \rangle = 0$ as $\mathbb{E}\langle X^\top \bar{E}c, d \rangle = \sum_{t \leq n-1} d^\top (\mathbb{E}x_t \bar{e}_{t+1}^\top) c$ with $\mathbb{E}x_t \bar{e}_{t+1}^\top = \langle X_t, \bar{E}_{t+1} \rangle = 0$, $X_t = [x_{t,1} \ \dots \ x_{t,m}]$, $\bar{E}_{t+1} = [\bar{e}_{t+1,1} \ \dots \ \bar{e}_{t+1,m}]$ by the definition above <3.11> and likewise for $\mathbb{E}\langle X^\top \bar{E}d, c \rangle$. The above representation leads to

$$\begin{aligned} \langle Gc, d \rangle &= \left\langle \frac{X^\top \bar{E} + \bar{E}^\top X}{2(n-1)} c, d \right\rangle = \frac{1}{2(n-1)} \sum_{t \leq n-1} \left[\langle \bar{e}_{t+1}, c \rangle \langle d, x_t \rangle + \langle \bar{e}_{t+1}, d \rangle \langle c, x_t \rangle \right] = \\ &= \frac{1}{n-1} v^\top \left[\underbrace{\frac{1}{2} \sum_{t \leq n-1} (\bar{B}_{t+1}^\top U_1^\top + \rho \bar{B}_{t+1}^\top U_2^\top) c d^\top (U_1 B_t + \rho U_2 \bar{B}_t) + \frac{A_{d,c}}{2}}_{A_{c,d}} \right] v = v^\top \left[\frac{A_{c,d} + A_{d,c}}{2(n-1)} \right] v. \end{aligned}$$

The equality $c = \pm d$ implies $A_{c,d} = A_{d,c} = \pm A_{c,c}$ and thereby $\|A_{c,d} + A_{d,c}\|_{\text{op}} = \|\pm 2A_{c,c}\|_{\text{op}} = 2\|A_{c,c}\|_{\text{op}}$. If $a, a' \in \mathbb{R}^{m(n+l)}$ consist of $n+l$ vectors $a_i, a'_i \in \mathbb{R}^m$, respectively, then

$$\begin{aligned} \langle a', A_{c,c} a \rangle &= \underbrace{\langle a', \sum_{t \leq n-1} \bar{B}_{t+1}^\top U_1^\top c c^\top U_1 B_t a \rangle}_{T_1} + \rho \underbrace{\langle a', \sum_{t \leq n-1} \bar{B}_{t+1}^\top U_2^\top c c^\top U_1 B_t a \rangle}_{T_2} \\ &\quad + \rho \underbrace{\langle a', \sum_{t \leq n-1} \bar{B}_{t+1}^\top U_1^\top c c^\top U_2 \bar{B}_t a \rangle}_{T_3} + \rho^2 \underbrace{\langle a', \sum_{t \leq n-1} \bar{B}_{t+1}^\top U_2^\top c c^\top U_2 \bar{B}_t a \rangle}_{T_4}. \end{aligned}$$

In case $\|a\| = \|a'\| = 1$, arguments similar to those used in the proof of lemma 3.10 yield

$$T_1 \leq \left(\sum_{i=0}^{l+1} \|\bar{A}_i^\top c_1\| \right) \left(\sum_{i=0}^l \|A_i c_1\| \right), \quad T_2 \leq \|c_2\| \left(\sum_{i=0}^l \|A_i c_1\| \right), \quad T_3 \leq \left(\sum_{i=0}^{l+1} \|\bar{A}_i^\top c_1\| \right) \|c_2\|,$$

and $T_4 \leq \|c_2\|^2$, wherein $c_i = U_i^\top c$ for $i \in \{1, 2\}$. The approximation factor β leads to

$$\begin{aligned} \sum_{i=0}^{l+1} \|\bar{A}_i^\top c_1\| &= \|A_0 c_1\| + \sum_{i=1}^l \|(A_i - \Gamma_* A_{i-1}) c_1\| + \|\Gamma_* A_l c_1\| \\ &\leq \|A_0 c_1\| + \sum_{i=1}^l \beta \|A_i c_1\| + \sum_{i=0}^{l-1} \beta \|A_i c_1\| + \|A_l c_1\| \leq \underbrace{\max\{1 + \beta, 2\beta\}}_{M_\beta} \sum_{i=0}^l \|A_i c_1\|, \end{aligned}$$

wherein the first inequality uses $\max\{a, b\} \leq a+b$, $a, b \in \mathbb{R}$ and the fact that A_i, Γ_* are diagonal matrices with $\|\Gamma_*\|_{\text{op}} \leq 1$. Setting $S_c = \sum_{i=0}^l \|A_i c_1\| / \|c_1\|$ if $c_1 \neq 0$ and $S_c = 0$ if $c_1 = 0$ yields

$$\|A_{c,c}\|_{\text{op}} \leq M_\beta [S_c^2 \|c_1\|^2 + 2(S_c \|c_1\|)(\rho \|c_2\|) + \rho^2 \|c_2\|^2] = M_\beta (S_c \|c_1\| + \rho \|c_2\|)^2 = M_\beta (S_c^2 + \rho^2).$$

Therein, the inequality uses $M_\beta = \max\{1 + \beta, 2\beta\} \geq 1$, in particular, $(M_\beta + 1) \leq 2M_\beta$, and the final equality is by virtue of the Cauchy-Schwarz inequality together with $\|c_1\|^2 + \|c_2\|^2 = \|U_1 U_1^\top c\|^2 + \|U_2 U_2^\top c\|^2 = \|c\|^2 = 1$. The Hanson-Wright inequality together with the inequali-

ties $\|(A_{c_i, \pm c_i} + A_{\pm c_i, c_i})/2\|^2 = \|A_{c_i, c_i}\|^2 \leq (n-1)\|A_{c_i, c_i}\|_{\text{op}}^2 = (n-1)\|(A_{c_i, \pm c_i} + A_{\pm c_i, c_i})/2\|_{\text{op}}^2$ and $m \geq n-1$ as well as the equalities $w_{c_i} = \bar{C}C^2M_\beta(S_{c_i}^2 + \rho^2)m$, $\bar{C} = \max\{\bar{C}^{-1} \log 18, 1\} \geq 1$, and \bar{C} being the unspecified constant in the Hanson-Wright inequality implies that

$$\mathbb{P}\left\{|\langle Gc_i, \pm c_i \rangle| > \bar{C}C^2M_\beta(S_{c_i}^2 + \rho^2)\frac{m}{n-1}\right\} = \mathbb{P}\left\{\left|v^\top \frac{A_{c_i, \pm c_i} + A_{\pm c_i, c_i}}{2}v\right| > \bar{C}C^2M_\beta(S_{c_i}^2 + \rho^2)m\right\}$$

is bounded above by $2 \exp(-(\log 9 + \log 2)m)$. The final step uses the inequality $\sum_{i=0}^l \|A_i c\| \leq \frac{6}{5} \left[\frac{3+\alpha}{\alpha}\right]^{1/2} \langle c, V_f c \rangle^{1/2} \leq \frac{6\sqrt{3}}{5} \left[\frac{1+\alpha}{\alpha}\right]^{1/2} \sigma_1^{1/2}(V_f)$, which relies on $\sum_{i=0}^l \|A_i c\|^2 = \langle c, V_f c \rangle$, applies to all $c \in \{\|\cdot\| = 1\}$ as well as $c = 0$, and is derived in the proof of lemma 3.10. Using the union bound on the events indexed by $(c_i, \pm c_i) \in \bar{H}$ reveals that

$$\|G(\omega)\|_{\text{op}} = \left\| \frac{X(\omega)^\top \bar{E}(\omega) + \bar{E}(\omega)^\top X(\omega)}{2(n-1)} \right\|_{\text{op}} \leq \bar{\bar{C}}C^2(1+\beta) \left(\frac{1+\alpha}{\alpha} \sigma_1(V_f) + \rho^2 \right) \frac{m}{n-1},$$

which relies on $\max\{1+\beta, 2\beta\} \leq 2(1+\beta)$ and $\bar{\bar{C}} = 2\frac{6^3}{5^2}\bar{C} \leq 18\bar{C}$, holds for all ω in an element of \mathcal{F} with probability larger or equal to $1 - 1/2^{m-2}$. ■

Proof of proposition 3.13. If $h = 0$, then lemma 3.12 together with $C_4 = \bar{\bar{C}}$ —the latter being the constant in lemma 3.12—guarantees that $\lambda \geq \|G\|_{\text{op}}$ and thereby $\|\hat{\Theta} - \Theta_*\| = 0$ with probability at least $1 - 1/2^{m-2} \geq 1 - \delta - 1/2^{m-3}$. If $h \geq 1$, then the requirements of lemma 3.10 are met. Thus, choosing the constant $C_1 = \max\{1, 81/\bar{C}\}$ and $C_2 = 41$, wherein \bar{C} denotes the (unspecified) constant in the Hanson-Wright inequality, implies that <3.24> holds and hence $m \geq C'h + C'' \log(2/\delta) \geq h + \log(2/\delta)$ with $C', C'' > 1$, in particular, $h < m$. Therefore,

$$\bar{\kappa} = \frac{1}{3} \sigma_h(V_f) \leq \sigma_h \left(\frac{F(\omega)^\top F(\omega)}{n-1} \right) \leq 2\sigma_h(V_f)$$

applies to all elements ω of some $S_1 \in \mathcal{F}$ with probability greater or equal to $1 - \delta$. Furthermore, the above inequalities together with lemma 3.11 implies that

$$\hat{\xi} = \hat{\xi}(\omega) \leq \frac{C_3}{\sigma_{k-1}^q(\Lambda)} \left(\sigma_h(V_f) + \rho^2 C_4 \frac{1+\alpha}{\alpha} \frac{\sigma_1(V_f)}{\sigma_h(V_f)} \left[\frac{m}{n-1} \right]^2 \right) = \bar{\xi} \leq \xi$$

holds for all $w \in S_1 \cap S_2$ for some $S_2 \in \mathcal{F}$ with $\mathbb{P}S_2 \geq 1 - 1/2^{m-1}$, wherein $C_3 = \max\{2, 12\bar{\bar{C}}^2\}$, $\bar{\bar{C}}$ symbolizes the constant of lemma 3.11, and $\hat{\xi}$ is defined in proposition 3.9. Finally, lemma 3.12 guarantees that the inequality

$$\hat{\lambda} = \hat{\lambda}(\omega) \leq C_4 C^2 (1+\beta) \left(\frac{1+\alpha}{\alpha} \sigma_1(V_f) + \rho^2 \right) \frac{m}{n-1} = \bar{\lambda} \leq \lambda$$

holds for all $\omega \in S_3$ with $S_3 \in \mathcal{F}$ and $\mathbb{P}S_3 \geq 1 - 1/2^{m-2}$, wherein $C_4 = \bar{\bar{C}}$ denotes the constant of lemma 3.12—as in case $h = 0$, and $\hat{\lambda}$ is defined in proposition 3.9. Consequently, the latter proposition ensures that the inequality <3.25> holds for all $\omega \in S_1 \cap S_2 \cap S_3$ and with $C_5 = 20 = \max\{20, 4\}$. An application of the union bound shows that the previous intersection has probability greater or equal to $1 - \delta - 3/2^{m-1} \geq 1 - \delta - 1/2^{m-3}$. ■

4. Prediction techniques

4.1. Oblique approximation

4.1.1. Distance between projectors

This section bounds the final summand in <2.5>, that is, the (squared) length of the difference $(P_V - P_{V/U})x$ between the orthogonal projection $P_V x$ and the oblique projection $P_{V/U} x$. Herein, x symbolizes an element of a Euclidean space W , and $V, U \subset W$ denote complementary subspaces of W ; the particular form $x = \sum_{i \leq q} a_i x_i$ considered in <2.5> bears no significance for the subsequent discussion. Bounding this distance has practical relevance as the first summand $\|x - P_V x\|^2$ of the rightmost term in <2.5> equals the minimal—and therefore unavoidable—(squared) approximation error when substituting x by an element of the set V based solely on the Euclidean space structure.

Orthogonal projections amount to oblique projections along the orthogonal complement $V^\perp = \cap_{x \in V} \{\langle x, \cdot \rangle = 0\}$ of V (in W). Hence, the equality $U = V^\perp$, in particular $U = \{0\} = V^\perp$, renders the endeavor of bounding the images of $P_{V/U} - P_V$ meaningless. More precisely, the kernel of $P_{V/U} - P_V$ takes the form $V + (V^\perp \cap U)$. Thus, its orthogonal complement $V^\perp \cap (V^\perp \cap U)^\perp$ is nontrivial if and only if $V^\perp \neq U$. In that case, the linearity of projectors allows restricting considerations to elements x of that complement. Then, proposition 4.1 provides alternative expressions for the multipliers (of $\|x\|$) in

$$\left[\inf_{v'} \|(P_{V/U} - P_V)v'\| \right] \|x\| \leq \|(P_{V/U} - P_V)x\| \leq \left[\sup_{v'} \|(P_{V/U} - P_V)v'\| \right] \|x\|, \quad <4.1>$$

wherein the infimum and supremum extend over all unit length elements of the orthogonal complement $V^\perp \cap (U \cap V^\perp)^\perp = V^\perp \cap (U^\perp + V)$ of the kernel of $P_{V/U} - P_V$.

More specifically, these two multipliers equal the tangent of the least nonzero *principal angle* $\theta_{\min, \neq 0} = \theta_{\min, \neq 0}(V^\perp, U)$ and that of the largest principal angle $\theta_{\max} = \theta_{\max}(V^\perp, U)$ between the two equal dimensional subspaces V^\perp and U . As shown below, the specification of the angles $\theta_{\min, \neq 0} \in [0, \pi/2]$ and $\theta_{\max} \in [0, \pi/2]$ in form of

$$\cos \theta_{\min, \neq 0}(V^\perp, U) = \sup_{v' \in V^\perp \cap (V^\perp \cap U)^\perp \cap \{\|\bullet\|=1\}} \sup_{u \in U \cap \{\|\bullet\|=1\}} \langle u, v' \rangle \quad \text{and} \quad <4.2a>$$

$$\cos \theta_{\max}(V^\perp, U) = \inf_{v' \in V^\perp \cap \{\|\bullet\|=1\}} \sup_{u \in U \cap \{\|\bullet\|=1\}} \langle u, v' \rangle \quad <4.2b>$$

is unequivocal due to the strict monotonicity of the cosine function on $[0, \pi/2]$.

More generally, the definition <4.2b> extends to an arbitrary $\ell = \text{codim } V$ dimensional subspace U of W , which needs not be complementary to V . Using this more general understanding, lemma 4.2 shows that $\theta_{\max}(V^\perp, U) = 0$ is equivalent to the equality $U = V^\perp$, that is, $V^\perp \cap (V^\perp \cap U)^\perp = \{0\}$. In particular, an analogous generalized understanding of definition <4.2a> is possible whenever $U \neq V^\perp$. Moreover, the inequality $\theta_{\max}(U, V^\perp) < \pi/2$ in connection with $\dim U = \text{codim } V$ is tantamount to the subspace U being complementary to V . Consequently, the expressions in proposition 4.1 are meaningful—in light of the definition of \tan —and are proved in section 4.1.2.

principal angle




Figure 4.1

The figure illustrates the two bounds in <4.4>. To this end, it considers the projections of $x \in \mathbb{R}^3$ onto a one dimensional space $V = \text{span}\{v_2\}$ along its orthogonal complement $V^\perp = \text{span}\{v'_1, v'_2\}$ as well as onto V along another complement $U = \text{span}\{u_1, u_2\}$. Herein $\theta_{\min, \neq 0} = \theta_{\max} = \theta_2$, $x \in V + (V^\perp \cap (V^\perp \cap U)^\perp)$, and thus the inequalities in <4.4> hold with equality.

Proposition 4.1. *Let $U \neq V^\perp$ be a $\ell \geq 1$ dimensional complement of a subspace V of a Euclidean space W , then the subspace $V'' = V^\perp \cap (V^\perp \cap U)^\perp$ is nontrivial and*

$$\inf_{v'' \in V'' \cap \{\|\bullet\|=1\}} \|(P_{V/U} - P_V)v''\| = \tan \theta_{\min, \neq 0} \quad \text{and} \quad <4.3a>$$

$$\sup_{v'' \in V'' \cap \{\|\bullet\|=1\}} \|(P_{V/U} - P_V)v''\| = \tan \theta_{\max}. \quad <4.3b>$$

Panel (B) of figure 2.2 shows how the multipliers in <4.3a> and <4.3b> depend on the corresponding principal angles $\theta_{\min, \neq 0}$ and θ_{\max} . Combining this result with <4.1> yields

$$\tan \theta_{\min, \neq 0} \|P_{V''}x\| \leq \|(P_{V/U} - P_V)x\| \leq \tan \theta_{\max} \|P_{V''}x\| \leq \tan \theta_{\max} \|\tilde{x}_V\|, \quad <4.4>$$

for all $x \in W$ and subspaces $V^\perp \neq U$, wherein $V'' = V^\perp \cap (V^\perp \cap U)^\perp$, and the principle angles refer to the pair of complements (V^\perp, U) of V . Figure 4.1 illustrates the possibility of—even simultaneous—nontrivial attainment of the bounds in <4.4>. This figure considers a one dimensional subspace $V = \text{span}\{v_2\}$ of \mathbb{R}^3 together with a complement $U = \text{span}\{u_1, u_2\}$. The latter shares the direction $u_1 = v'_1$ with the orthogonal complement V^\perp . Consequently, $V^\perp \cap (V^\perp \cap U)^\perp = \text{span}\{v'_1, v'_2\} \cap (\text{span}\{v'_1\})^\perp = \text{span}\{v'_2\}$, and furthermore the least nonzero and the largest principal angle between V and U both equal θ_2 . The upper righthand corner magnifies $x \in V + V''$ alongside its two projections $\hat{x} = P_V x$ and $\hat{x}_/ = P_{V/U} x$ to visualize <4.3a> and <4.3b>. Panel (B) of figure 4.2 reconsiders this setting, and its notation is adapted to the discussion in section 4.1.2.

Finally, a comment on the role of W in <4.4> is in order. An individual projection of a given $x \in W$ onto V bears no reference to directions outside $\text{span}\{x\} + V$. Hence, only (at most) one dimensional complements are relevant. In this minimal setting, the first upper bound in <4.4> coincides with the lower one. However, the inequalities <4.4> also apply to W with higher dimension. In that case, these bounds hold uniformly over all $x \in W$, but at the price of less favorable multipliers. In any case, the influence of the dimension of W exhibits an overall limit in that $(P_{V/U} - P_V)x$ differs from zero if and only if x extends into $V^\perp \cap (V^\perp \cap U)^\perp$. More specifically, the dimension of $V^\perp \cap U$ is immaterial, and its codimension does not exceed $2 \dim V$. Therein, the latter dimension bound follows from $(V^\perp \cap U)^\perp = V + U^\perp$ together with $\dim U^\perp = \dim V$.

4.1.2. Angles between subspaces

This section considers two complementary subspaces V and U of a Euclidean space W and investigates the concept of principal angles between the equal dimensional subspaces V^\perp and U . If these two subspaces differ, then the least nonzero principal angle $\theta_{\min, \neq 0}(V^\perp, U)$ is defined by <4.2a>. In any case, the largest of these angles $\theta_{\max}(V^\perp, U)$ is determined by <4.2b>. At first sight, the order of subspaces in these definitions matters. Lemma 4.2 negates this impression. Until then, $\theta_{\min, \neq 0}$ and θ_{\max} symbolize $\theta_{\min, \neq 0}(V^\perp, U)$ and $\theta_{\max}(V^\perp, U)$ as defined in <4.2a> and <4.2b>, respectively.

The final term $\phi_U(v') = \sup_{u \in U \cap \{\|\bullet\|=1\}} \langle u, v' \rangle$ of either of these two definitions is key to understanding their construction. Compactness of $\{\|\bullet\|=1\}$ and continuity of $\langle \bullet, v' \rangle$ ensure that this supremum is attained. Consequently, $\phi_U(v')$ equals the cosine of the smallest (small) angle between a unit length element $u \in U$ and the given unit length v' . It is conventionally interpreted as the *angle* between v' and the subspace U . *angle*

If $v' \notin U^\perp$, then the unit length element $u_* = \hat{v}'_U / \|\hat{v}'_U\|$ exists, and $\langle u, v' \rangle = \langle u, \hat{v}'_U \rangle$ ensures that the supremum $\phi_U(v') = \|\hat{v}'_U\|$ is solely attained at u_* . If $v' \in U^\perp$, then $\phi_U(v') = 0 = \|\hat{v}'_U\|$, which is clearly attained for all $u \in U \cap \{\|\bullet\|=1\}$. The subspace U of \mathbb{R}^3 and $v' \in \mathbb{R}^3$ shown in panel (A) of figure 4.2 exemplify these arguments for the case that $v' \notin U$ and $v' \notin U^\perp$. Then, $\phi_U(v')$ lies in the interval $(0, 1)$. If (and only if) $v' \in U^\perp$, then $\phi_U(v') = 0$; if (and only if) $v' \in U$, then $\phi_U(v') = 1$. Thus, there generally exists a unique angle $\theta_{\max} \in [0, \pi/2]$ satisfying <4.2b>. Moreover, if $V^\perp \neq U$, then $V^\perp \cap (U \cap V^\perp)^\perp \neq \{0\}$ and the angle $\theta_{\min, \neq 0} \in (0, \pi/2]$ in <4.2a> is well defined.

Furthermore, continuity of $\|P_U \bullet\|$ guarantees that $\inf_{v' \in V^\perp \cap \{\|\bullet\|=1\}} \phi_U(v')$ is attained, and therefore $V^\perp \cap U^\perp \neq \{0\}$ is tantamount to $\theta_{\max}(V^\perp, U) = \pi/2$. In particular, if U and V are complementary, then so are U^\perp and V^\perp , that is, $U^\perp \cap V^\perp = \{0\}$ or equivalently $\theta_{\max}(V^\perp, U) < \pi/2$. Lemma 4.2 asserts the converse and justifies calling θ_{\max} an angle between V^\perp and U , which after all suggest some sort of symmetry.

Lemma 4.2. *If U and V are two ℓ dimensional subspaces of a Euclidean space W , then*

- (a) $\theta_{\max} = \theta_{\max}(U, V^\perp) = \theta_{\max}(V, U^\perp)$,
- (b) U and V are complementary if and only if $\theta_{\max} < \pi/2$, and
- (c) U and V^\perp are equal if and only if $\theta_{\max} = 0$.

Furthermore, if $\theta_{\max} > 0$, then the notion of $\theta_{\min, \neq 0}$ is meaningful and satisfies

$$\theta_{\min, \neq 0} = \theta_{\min, \neq 0}(U, V^\perp) = \theta_{\min, \neq 0}(V, U^\perp).$$

To address the claim in lemma 4.2, let V be a subspace of a Euclidean space W with positive codimension $\ell = \text{codim } V$ and U be another ℓ dimensional subspace of W . Initially, the inequality $U \neq V^\perp$ is dispensable as only θ_{\max} is of concern.

Compactness of $U \cap \{\|\cdot\| = 1\} \times V^\perp \cap \{\|\cdot\| = 1\} \neq \emptyset$ and continuity of $\langle \cdot, \cdot \rangle$ —both with respect to the product (of the norm) topology on $W \times W$ —guarantee the equalities

$$\sup_{v' \in V^\perp \cap \{\|\cdot\|=1\}} \phi_U(v') = \sup_{(v', u) \in V^\perp \cap \{\|\cdot\|=1\} \times U \cap \{\|\cdot\|=1\}} \langle v', u \rangle = \sup_{u \in U \cap \{\|\cdot\|=1\}} \phi_{V^\perp}(u), \quad <4.5>$$

and thereby supply a symmetric definition of a least principal angle θ_1 . The same argument ensures the existence of a unit length $v'_1 \in V^\perp$ and a unit length $u_1 \in U$ such that $\cos \theta_1 = \langle v'_1, u_1 \rangle$. Any valid choice of v'_1, u_1 satisfies $\phi_U(v'_1) = \langle u_1, v'_1 \rangle = \phi_{V^\perp}(u_1)$. Hence, the above discussion reveals $P_{V^\perp} u_1 \in \text{span}\{v'_1\}$ and $P_U v'_1 \in \text{span}\{u_1\}$ and consequently $P_{V^\perp} u_1 = \langle u_1, v'_1 \rangle v'_1 = \cos \theta_1 v'_1$ and $P_U v'_1 = \cos \theta_1 u_1$.

If $\dim U = \ell \geq 2$, then a further principal angle θ_2 may be defined based on a specific choice of v'_1 and u_1 . More specifically, using $U \cap (\text{span}\{u_1\})^\perp$ in place of U and $V^\perp \cap (\text{span}\{v'_1\})^\perp$ instead of V^\perp in <4.5> yields a second smallest principal angle $\theta_2 \geq \theta_1$. As before, any two points v'_2, u_2 with $\cos \theta_2 = \langle v'_2, u_2 \rangle$ satisfy

$$P_{V^\perp \cap (\text{span}\{v'_1\})^\perp} u_2 = (\cos \theta_2) v'_2 \quad \text{and} \quad P_{U \cap (\text{span}\{u_1\})^\perp} v'_2 = (\cos \theta_2) u_2.$$

Moreover, $\langle v'_2 - (\cos \theta_2) u_2, u_1 \rangle = \langle v'_2, u_1 \rangle = \cos \theta_1 \langle v'_2, v'_1 \rangle = 0$ ensures that $(\cos \theta_2) u_2 = P_U v'_2$. As a consequence, one has $\text{span}\{u_1, v'_1\} \subset (\text{span}\{u_2, v'_2\})^\perp$ and $(\cos \theta_2) v'_2 = P_{V^\perp} u_2$.

If $\ell \geq 3$, then the definition of θ_3 based on previously chosen v'_1, v'_2 and u_1, u_2 proceeds by analogous steps and so forth. Orthonormality of u_1, u_2, \dots as well as v'_1, v'_2, \dots ensures that the iteration terminates after ℓ steps. Its output comprises ℓ principal angles $\theta_1 \leq \theta_2 \leq \dots \leq \theta_\ell \leq \pi/2$ as well as (orthonormal) bases v'_1, \dots, v'_ℓ of V^\perp and u_1, \dots, u_ℓ of U . These quantities satisfy $\text{span}\{u_i, v'_i\} \subset (\text{span}\{u_j, v'_j\})^\perp$, $P_U v'_j = (\cos \theta_j) u_j$, and $P_{V^\perp} u_j = (\cos \theta_j) v'_j$ for $i < j \leq \ell$. Thus, the orthogonal projections of $u \in U$ and $v' \in V^\perp$ onto V^\perp and U , respectively, may be represented as

$$\sum_{i \leq \ell} \langle u_i, u \rangle (\cos \theta_i) v'_i = \sum_{i \leq \ell} \langle v'_i, u \rangle v'_i = \hat{u}_{V^\perp} \quad \text{and} \quad \sum_{i \leq \ell} \langle v'_i, v' \rangle (\cos \theta_i) u_i = \hat{v}'_U. \quad <4.6>$$

The second of these equalities shows that the set of orthogonal projections $\{\hat{v}'_U \mid v' \in V^\perp, \|v'\| = 1\}$ is given by the ellipsoid $\{\sum_{i \leq \ell} \langle \cdot, u_i \rangle^2 / \cos^2 \theta_i = 1\} \cap U$ with principal semi-axes lengths $\cos \theta_1 \geq \cos \theta_2 \geq \dots \geq \cos \theta_\ell$. Panel (B) of figure 4.2 illustrates the elliptical shape of $\{\hat{v}'_U \mid v' \in V^\perp, \|v'\| = 1\}$ for a complement U of $V \subset \mathbb{R}^3$ and $\ell = 2$.

The representation in <4.6> implies that every element $v' \in V^\perp$ satisfies

$$\phi_U(v') = \sup_{u \in U \cap \{\|\cdot\|=1\}} \langle P_{V^\perp} u, v' \rangle = \sup_{u \in U \cap \{\|\cdot\|=1\}} \sum_{i \leq \ell} \cos \theta_i \langle u_i, u \rangle \langle v_i, v' \rangle \geq \cos \theta_\ell \langle v_\ell, v' \rangle.$$




Figure 4.2

The figure illustrates the concept of the angle between a unit length element v' and a subspace U (panel (A)) and the notion of principal angles between two equal dimensional subspaces U and V^\perp (panel (B)). Panel (B) also shows $\{\hat{v}'_U \mid v' \in V^\perp, \|v'\| = 1\}$ (solid black line).

Hence, the largest principal angle $\theta_{\max} = \theta_{\max}(V^\perp, U)$ —as defined in <4.2b>—coincides with θ_ℓ , which verifies $\theta_{\max} = \theta_{\max}(U, V^\perp)$. This equality together with the remark above lemma 4.2 implies that $\theta_{\max} < \pi/2$ is equivalent to $U \cap (V^\perp)^\perp = U \cap V = \{0\}$, which completes the verification of (b) of lemma 4.2. In fact, if $\theta_{\max} < \pi/2$, then $\dim U = \text{codim } V$ implies that combining a basis of U and one of V yields a basis of W . The equality $\theta_{\max}(V^\perp, U) = \theta_{\max}(U, V^\perp)$ also handles the assertion (c). More specifically, the two interpretations of $\theta_{\max} = 0$ are equivalent to $V^\perp \subset U$ and $V^\perp \supset U$, respectively.

If $\theta_{\max} = 0$, then $V^\perp = U$ and consequently $\theta_{\max}(V, U^\perp) = 0$. Otherwise $\theta_{\max} > 0$, and there exists a least $j \leq \ell$ such that $\theta_j \neq 0$. Then, <4.6> reveals that $\theta_j = \theta_{\min, \neq 0}$ and therefore $\theta_{\min, \neq 0} = \theta_{\min, \neq 0}(U, V^\perp)$. In fact, if $j > 1$, then v'_1, \dots, v'_{j-1} or equivalently u_1, \dots, u_{j-1} form an orthonormal basis of $U \cap V^\perp$. The same applies to v'_j, \dots, v'_ℓ with respect to $V^\perp \cap (V^\perp \cap U)^\perp$ and u_j, \dots, u_ℓ with respect to $U \cap (V^\perp \cap U)^\perp$. Panel (B) of figure 4.2 exemplifies these relations with $j = 2$, thus, $U \cap V^\perp = \text{span}\{v'_1\} = \text{span}\{u_1\}$.

Panel (B) of figure 4.2 also indicates that the residuals $P_V u_i = u_i - P_{V^\perp} u_i$ and $P_{U^\perp} v'_i = v'_i - P_U v'_i$ for $i \in \{j, \dots, \ell\}$ exhibit length $\sin \theta_i > 0$. Thus,

$$v_i = \frac{P_V u_i}{\sin \theta_i} \in \text{span}\{u_i, v'_i\} \subset (U^\perp \cap V)^\perp \quad \text{and} \quad u'_i = \frac{P_{U^\perp} v'_i}{\sin \theta_i} \in \text{span}\{u_i, v'_i\}$$

are well-defined unit length elements of $V \cap (U^\perp \cap V)^\perp$ and $U^\perp \cap (U^\perp \cap V)^\perp$, respectively. These elements serve to verify $\theta_{\max} = \theta_{\max}(V, U^\perp)$ and $\theta_{\min, \neq 0} = \theta_{\min, \neq 0}(V, U^\perp)$ in case $V^\perp \neq U$ or equivalently $V \neq U^\perp$. The case $U^\perp = V$ is covered by (c). More specifically,

$$\langle -u'_i, v_i \rangle = \frac{1}{\sin^2 \theta_i} \langle P_{U^\perp} v'_i, P_{V^\perp} u_i - u_i \rangle = \frac{1}{\sin^2 \theta_i} \langle v'_i - \cos \theta_i u_i, \cos \theta_i v'_i \rangle = \cos \theta_i$$

imply $v_i + \cos \theta_i u'_i \in (\text{span}\{u'_i\})^\perp$, that is, $P_{U^\perp} v_i = \cos \theta_i (-u'_i)$, and therefore $\phi_{U^\perp}(v_i) = \cos \theta_i$ for all $j \leq i \leq \ell$. Choosing $i = \ell, j$ implies $\cos \theta_{\max}(V, U^\perp) \leq \cos \theta_{\max}$ and $\cos \theta_{\min, \neq 0}(V, U^\perp) \geq \cos \theta_{\min, \neq 0}$, respectively, as $v_j \in V \cap (U^\perp \cap V)^\perp \subset V$. Reversing the roles of the respective orthogonal complements concludes the proof of lemma 4.2.

Proposition 4.1 presupposes $\theta_{\max} \neq 0$, then $P_{V/U} - P_V = P_{V^\perp} - P_{U/V}$ equals

$$\sum_{i \leq \ell} v'_i \langle v'_i, \cdot \rangle - \sum_{i \leq \ell} \frac{1}{\cos \theta_i} u_i \langle v'_i, \cdot \rangle = \sum_{j \leq i \leq \ell} \frac{\cos \theta_i v'_i - u_i}{\cos \theta_i} \langle v'_i, \cdot \rangle = \sum_{j \leq i \leq \ell} \tan \theta_i (-v_i) \langle v'_i, \cdot \rangle,$$

wherein $P_{U/V} = \sum_{i \leq \ell} (\cos \theta_i)^{-1} u_i \langle v'_i, \cdot \rangle$ may be verified from the definition of the projector. The equalities <4.3a> and <4.3b> follow from the final representation of $P_{V/U} - P_V$.

4.2. Subordinate directions

4.2.1. Down-weighting directions

An orthonormal basis of a Euclidean space W represents an exhaustive set of perpendicular directions. The coordinates of $x \in W$ with respect to this basis quantify its extent into the individual directions. Prioritizing some directions when projecting x onto the span V of a nontrivial sequence y_1, \dots, y_k in W implements an oblique projection.

More specifically, let U be a subspace of W and denote by \hat{Y} and \tilde{Y} the composition of $Y = [y_1 \cdots y_k]$ with the orthogonal projector onto U and U^\perp , respectively. Proposition 4.3 asserts that the notion of a projector onto $V = \text{img } Y$ along $V_*^\perp = (\text{img } Y_*)^\perp$ with $Y_* = \hat{Y} + \tilde{Y} P_{\ker \hat{Y}}$, thus $\tilde{Y}_* = P_{U^\perp} Y_* = \tilde{Y} P_{\ker \hat{Y}}$, is well defined. The latter assigns a subordinate role to the directions in U^\perp . In fact, P_{V/V_*^\perp} may be thought of as an “intermediate” between P_V and P_{V_*} . The decomposition $P_{V_*} = P_{\text{img } \hat{Y}} + P_{\text{img } \tilde{Y}_*}$, which relies on $\text{img } \hat{Y} \subset (\text{img } \tilde{Y}_*)^\perp$ as well as $\ker \hat{Y} \supset (\ker \tilde{Y}_*)^\perp$, points to the prime focus of P_{V_*} on $\text{img } \hat{Y} \subset U$; the modification of \tilde{Y} highlights the subordinate role of the part of Y in U^\perp . In contrast, $P_V = P_{\text{img } Y}$ does not generally decompose in an analogous way.

Proposition 4.3. *Let U be a subspace of a Euclidean space W , $Y \in W^{\times k}$ with image $V \neq \{0\}$ and $\hat{Y} = \hat{Y}_U$, $\tilde{Y} = \tilde{Y}_U$. In addition, consider $Y_* = \hat{Y} + \tilde{Y} P_{\ker \hat{Y}}$ with image V_* .*

Then, one has $\ker Y = \ker Y_$ and thereby $\dim V = \dim V_*$ as well as $V \cap V_* = \text{img } \hat{Y}_\times + \text{img } \tilde{Y}_*$, wherein $\tilde{Y}_* = P_{U^\perp} Y_* = \tilde{Y} P_{\ker \hat{Y}}$ and $\hat{Y}_\times = \hat{Y} P_{\{\tilde{Y} \in \text{img } \tilde{Y}_*\}}$. Consequently, the equality $V = V_*$ is tantamount to $\text{img } \tilde{Y}_* = \text{img } \tilde{Y}$, and if neither of these equalities holds, then $K = (\ker P_{(\text{img } \hat{Y}_\times)^\perp} \hat{Y})^\perp = (\ker P_{(\text{img } \tilde{Y}_*)^\perp} \tilde{Y})^\perp$ is nontrivial,*

$$\tan \theta_{\min, \neq 0}(V_*^\perp, V^\perp) = \inf_{c \in K} \frac{\|P_{(\text{img } \tilde{Y}_*)^\perp} \tilde{Y} c\|}{\|P_{(\text{img } \hat{Y}_\times)^\perp} \hat{Y} c\|}, \quad \text{and}$$

$$\tan \theta_{\max}(V_*^\perp, V^\perp) = \sup_{c \in K} \frac{\|P_{(\text{img } \tilde{Y}_*)^\perp} \tilde{Y} c\|}{\|P_{(\text{img } \hat{Y}_\times)^\perp} \hat{Y} c\|} \leq \frac{\|P_{(\text{img } \tilde{Y}_*)^\perp} \tilde{Y}\|_{\text{op}}}{\sigma_{\min, \neq 0}(P_{(\text{img } \hat{Y}_\times)^\perp} \hat{Y})} < \infty.$$

In particular, V_^\perp generally provides a complement of V .*




Figure 4.3

The figure compares the orthogonal projection \hat{x}_V of x onto $V = \text{span}\{y_1, y_2\}$ with the oblique projection $\hat{x}_/$ of x onto V along V_*^\perp . Panel (A) shows the entire setting: an orthonormal basis q_1, q_2, q_3 of \mathbb{R}^3 , the spanning set y_1, y_2 , as well as x and its two projections. Panel (B) amplifies the surrounding of x to highlight the (limited) $\|\cdot\|$ -optimality associated with \hat{x}_V and $\hat{x}_/$.

A proof of proposition 4.3 starts on page 121 in appendix 4.a. The remainder of this section identifies P_{V/V_*^\perp} as an “intermediate” between the projectors P_V and P_{V_*} .

More specifically, the definition of $P_{V/V_*^\perp} x = Yb_*$ implies that its coordinate vectors b_* with respect to Y are characterized by $P_{V_*^\perp/V} x = x - Yb_* \in V_*^\perp = (\text{img } Y_*)^\perp$, that is,

$$\langle Y_*, x - Yb_* \rangle = \langle \hat{Y} + \tilde{Y}_*, x - \hat{Y}b_* - \tilde{Y}b_* \rangle = \langle \hat{Y}, x - \hat{Y}b_* \rangle + \langle \tilde{Y}_*, x - \tilde{Y}b_* \rangle.$$

Therein, $\langle \hat{Y}, x - \hat{Y}b_* \rangle \in (\ker \hat{Y})^\perp$ and $\langle \tilde{Y}_*, x - \tilde{Y}b_* \rangle = P_{\ker \hat{Y}} \langle \tilde{Y}, x - \tilde{Y}b_* \rangle \in \ker \hat{Y}$ imply that both of these terms equal zero. In particular, one has $\hat{Y}b_* = P_{\text{img } \hat{Y}} x$. More specifically, every choice of b_* may be thought of as the result of the following two step procedure. Firstly, choose the part of b_* in $(\ker \hat{Y})^\perp$ to minimize $\|\hat{Y}b_* - x\|$. Secondly, use the part in $\ker \hat{Y}$ to minimize $\|\tilde{Y}b_* - x\| = \|\tilde{Y}_*b_* - (x - \tilde{Y}P_{(\ker \hat{Y})^\perp} b_*)\|$. In summary,

$$\inf_{c \in H} \|\tilde{Y}c - x\|, \quad H = \{c' \in \mathbb{R}^k \mid \|\hat{Y}c' - x\| = \inf_{c'' \in \mathbb{R}^k} \|\hat{Y}c'' - x\|\}. \quad <4.7>$$

Hence, if Yb'_* equals the orthogonal projection $P_{V_*} x$, then the difference $b_* - b'_*$ lies in $\ker \hat{Y}$, but is usually nonzero as $P_{\ker \hat{Y}} b'_*$ minimizes the criterion $\|\tilde{Y}_*b'_* - x\|$.

Figure 4.3 illustrates the workings of P_{V/V_*^\perp} in \mathbb{R}^3 for $k = 2$, $U = \text{span}\{q_1\}$, and $\hat{x}_U \in \text{img } \hat{Y}$. The latter inclusion implies $\hat{Y}b_* = P_{\text{img } \hat{Y}} x = \hat{x}_U = q_1 \langle q_1, x \rangle$. Panel (A) shows the entire setting including the orthonormal basis q_1, q_2, q_3 , the orthogonal projection \hat{x}_V , and the oblique projection $\hat{x}_/ = P_{V/V_*^\perp} x$. The latter shares its extent into the sole

direction q_1 of U with $P_{\text{img } \hat{Y}} x = \hat{x}_U = q_1 \langle q_1, x \rangle$ as required by <4.7>. Panel (B) zooms in on the two projections to illustrate the $\|\cdot\|$ -inferiority of the oblique projection P_{V/V_*^\perp} relative to P_V as well as the $\|\cdot\|$ -optimality of the former under the constraint in <4.7>.

Finally, the case $\text{img } Y \cap U^\perp = \{0\}$ is noteworthy. Then $\hat{Y}c = 0$ implies $\tilde{Y}c = Yc \in U^\perp \cap \text{img } Y = \{0\}$, which effectively removes the second stage in <4.7>.

4.2.2. Inner products and linear space structure

This section considers an additional inner product $\langle \cdot, \cdot \rangle_*$ on a linear space W' , thus, handles the geometry of $(W', \langle \cdot, \cdot \rangle_*)$ alongside that of $(W', \langle \cdot, \cdot \rangle)$. As mentioned in section 2.4.1, the orthogonal complement $V^{\perp*}$ of a subspace V in $(W', \langle \cdot, \cdot \rangle_*)$ usually differs from the orthogonal complement V^\perp of V in $(W', \langle \cdot, \cdot \rangle)$. Consequently, one has $P_{V/V^{\perp*}} \neq P_{V/V^\perp}$, and $P_{V/V^{\perp*}}$ is therefore $\|\cdot\|$ -suboptimal. Proposition 4.1 bounds its lack of $\|\cdot\|$ -optimality in terms of $\langle \cdot, \cdot \rangle$ -angle $\theta_{\max}(V^\perp, V^{\perp*})$.

Figure 4.4 illustrates this setting in $W' = \mathbb{R}^2$. Therein, $V = \text{span}\{e_1\}$ equals the first coordinate axis, and the alternative inner product of $a, b \in \mathbb{R}^2$ has the form

$$\langle a, b \rangle_* = a^\top \begin{pmatrix} 1 & \rho\psi \\ \psi\rho & \psi^2 \end{pmatrix} b, \quad 0 < \rho < 1 < \psi. \quad <4.8>$$

In particular, elements $a = (a_1, a_2)$ of the $\langle \cdot, \cdot \rangle_*$ -orthogonal complement $V^{\perp*}$ of V satisfy $a_1 = -\psi\rho a_2$. Thus, $V^{\perp*}$ differs from the $\langle \cdot, \cdot \rangle$ -orthogonal complement V^\perp . The largest—in fact, the sole—principal $\langle \cdot, \cdot \rangle$ -angle θ_{\max} between $V^{\perp*}$ and $V^\perp = \text{span}\{e_2\}$ satisfies $\cos \theta_{\max} = (1 + \psi^2 \rho^2)^{-1/2}$, $\sin \theta_{\max} = \psi\rho / (1 + \psi^2 \rho^2)^{1/2}$, and $\tan \theta_{\max} = \psi\rho$. The configuration of ρ, ψ in figure 4.4 is such that $\psi\rho < 1$, thus, $\theta_{\max} < \pi/4$. This principal angle bounds the $\|\cdot\|$ -loss as shown in <4.4> when using $P_{V/V^{\perp*}}$ in place of P_{V/V^\perp} .

The inner product <4.8> is defined by specification of a $\langle \cdot, \cdot \rangle_*$ -Gramian of the spanning set e_1, e_2 of \mathbb{R}^2 . Lemma 2.3 grants full generality to this approach. That is, inner products $\langle \cdot, \cdot \rangle_*$ on a linear space W' spanned by the columns of $Z = [z_1 \cdots z_m]$ are in one-to-one correspondence with the possible $\langle \cdot, \cdot \rangle_*$ -Gramians G of Z given by the symmetric and positive semidefinite $m \times m$ matrices with kernel $\ker Z$. Hence, such a matrix identifies an oblique projector $P_{V/V^{\perp*}}$ onto a subspace V of W' as the $\langle \cdot, \cdot \rangle_*$ -orthogonal one. This representation has practical implications regarding the computation of projections $P_{V/V^{\perp*}} x$, $x \in W'$. In particular, computational schemes designed for orthogonal projections—such as the Gram-Schmidt orthogonalization—become applicable.

Below, the space W' often forms a subspace of a larger Euclidean space $(W, \langle \cdot, \cdot \rangle)$. An alternative (to $\langle \cdot, \cdot \rangle$) inner product $\langle \cdot, \cdot \rangle_*$ on W' defines an orthogonal projector $P_{V/V^{\perp*}}$ projecting elements $x \in W'$ onto a subspace $V \subset W'$. If this projector coincides with the restriction to W' of a projector P_V defined on W and projecting onto V , then the above mentioned computational gains also extend to $P_V x$ for $x \in W'$. The oblique projector P_{V/V_*^\perp} of section 4.2.1 provides an example, wherein $W' = \text{img } [Y \ x]$, $V = \text{img } Y$ with $y_1, \dots, y_k, x \in W$ and $Y = [y_1 \cdots y_k]$. In this case, \hat{Y} and \tilde{Y} denote the composition of Y with the orthogonal projectors P_U and P_{U^\perp} , respectively, wherein U denotes a subspace of W . If $b \in (\ker \hat{Y})^\perp$ represents a coordinate vector with respect to \hat{Y}




Figure 4.4

The figure compares the orthogonal projections \hat{x}_V and \hat{x}_{V/V^\perp} of an element $x \in \mathbb{R}^2$ onto $V = \text{span}\{e_1\}$ derived under the standard inner product $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle_*$, respectively. The distance (in terms of $\langle \cdot, \cdot \rangle$) between the two projections is governed by the largest principal $\langle \cdot, \cdot \rangle$ -angle θ_{\max} between the two orthogonal complements V^\perp and $V^{\perp*}$.

of $P_{\text{img } \hat{Y}} x = \hat{Y} b$, $x_* = x - \tilde{Y} b$, and $Y_* = \hat{Y} + \tilde{Y} P_{\ker \hat{Y}}$, then the linear map $[Y_* \ x_*]$ shares its kernel with $[Y \ x]$. Thus, taking the Gramian of the former linear map as the $\langle \cdot, \cdot \rangle_*$ -Gramian of $[Y \ x]$ introduces a further inner product $\langle \cdot, \cdot \rangle_*$ on W' . In this setting, the projection $P_{V/V^{\perp*}} x$ uniquely determines the $\langle \cdot, \cdot \rangle_*$ -orthogonal projector defined on W' . Coordinate vectors b_* of this projection with respect to Y are characterized by

$$0 = \langle Y, x - Y b_* \rangle_* = \langle Y_*, x_* - Y_* b_* \rangle = \langle \hat{Y}, x - \hat{Y} b_* \rangle + \langle \tilde{Y}_*, x - \tilde{Y}(b + P_{\ker \hat{Y}} b_*) \rangle.$$

Therein, the first of the rightmost summands lies in $(\ker \hat{Y})^\perp$. The second summand is an element of $\ker \hat{Y}$. Hence, both terms equal zero. The first of the implied equalities, that is, $\langle \hat{Y}, x - \hat{Y} b_* \rangle = 0$, guarantees $\hat{Y} b_* = P_{\text{img } \hat{Y}} x$ and thereby $P_{(\ker \hat{Y})^\perp} b_* = b$. Consequently, the second implied equality becomes $\langle \tilde{Y}_*, (x - \tilde{Y} P_{(\ker \hat{Y})^\perp} b_*) - \tilde{Y}_* b_* \rangle = 0$. The latter coincides with the optimality condition for the second stage in $\langle 4.7 \rangle$. In total, the two implied equalities guarantee $P_{V/V^{\perp*}} x = Y b_* = P_{V/V^\perp} x$. That is, the restriction of the projector $P_{V/V^{\perp*}}$ to the subspace W' of W equals the $\langle \cdot, \cdot \rangle_*$ -orthogonal projector P_{V/V^\perp} .

An extension of the special case $\text{img } Y \cap U^\perp = \{0\}$ —mentioned at the end of section 4.2.1—allows a simpler construct. If $\text{img } [Y \ x] \cap U^\perp = \{0\}$, then the equality $\ker [Y \ x] = \ker [\hat{Y} \ \hat{x}]$, wherein $\hat{Y} = \hat{Y}_U$ and $\hat{x} = \hat{x}_U$, follows in analogy with the consideration in section 4.2.1. Consequently, the Gramian of $[\hat{Y} \ \hat{x}]$ induces an inner product $\langle \cdot, \cdot \rangle_\circ$ on $W' = \text{img } [Y \ x]$. The orthogonality conditions characterizing coordinate vectors $b_\circ \in \mathbb{R}^k$ with respect to Y of the $\langle \cdot, \cdot \rangle_\circ$ -orthogonal projections \hat{x}_{V/V^\perp} amount to

$$0 = \langle Y, x - Y b_\circ \rangle_\circ = \langle \hat{Y}, \hat{x} - \hat{Y} b_\circ \rangle = \langle Y_*, (x - \tilde{Y} b) - Y_* b_\circ \rangle = \langle Y, x - Y b_\circ \rangle_*$$

due to the equality $Y_* = \hat{Y}$, which follows from $\text{img } Y \cap U^\perp \subset \text{img } [Y \ x] \cap U^\perp = \{0\}$.

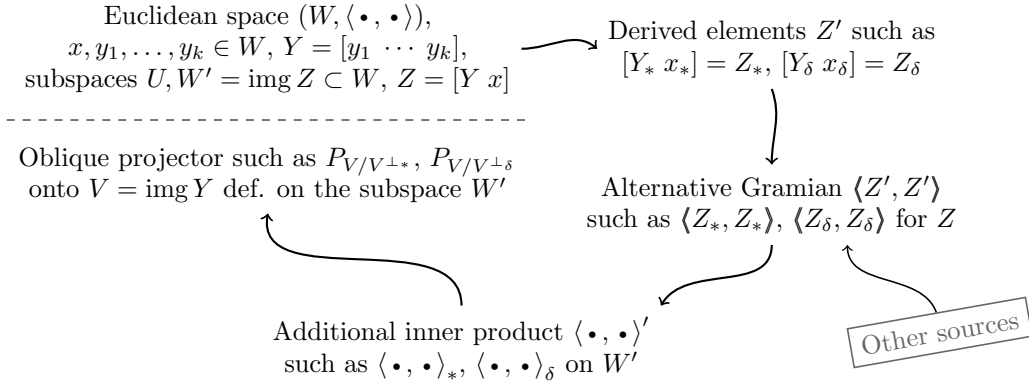


Figure 4.5

The figure visualizes the derivation of oblique projectors (defined on the subspace $W' = \text{img } [Y \ x]$ of W) considered in the (present) section 4.2.2.

Thus, $P_{V/V^{\perp_*}}x = Yb_* = Yb_* = P_{V/V^{\perp_*}}x$ coincides with the oblique projection $P_{V/V^{\perp_*}}x$.

Another example of an alternative inner product comes in the form of the inner products $\langle \cdot, \cdot \rangle_\delta$ induced by the Gramians of $[Y_\delta \ x_\delta]$ with $Y_\delta = \hat{Y}/\delta + \tilde{Y}$, $x_\delta = \hat{x}/\delta + \tilde{x}$, and $\delta \in (0, 1]$. The $\langle \cdot, \cdot \rangle_\delta$ -orthogonal projectors $P_{V/V^{\perp_\delta}}$ increasingly focus on the parts \hat{Y} and \hat{x} in U as $\delta \downarrow 0$. In fact, $P_{V/V^{\perp_*}}$ results as a limit case of $(P_{V/V^{\perp_\delta}})_{\delta \in (0, 1]}$ as $\delta \rightarrow 0$. More precisely, the argument on page 122 in appendix 4.a shows that

$$\lim_{\delta \rightarrow 0} \sup_{z' \in W' \cap \{\|\cdot\|=1\}} \|(P_{V/V^{\perp_*}} - P_{V/V^{\perp_\delta}})z'\| = 0. \quad <4.9>$$

Two observations are key to the verification of <4.9>. Firstly, the two projectors $P_{V/V^{\perp_*}}$ and $P_{V/V^{\perp_\delta}}$ —defined on $W' = \text{img } [Y \ x]$ and a given $\delta \in (0, 1]$ —differ solely on the (at most) one dimensional complements of V . Therefore, the statement essentially reduces to the convergence of a sequence $((P_{V/V^{\perp_*}} - P_{V/V^{\perp_{\delta_n}}})z)_{n \in \mathbb{N}}$ in W' with $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. Secondly, norm equivalence—discussed in section 2.1.2—allows to verify the convergence using the norm $\|\cdot\|_*$ induced by $\langle \cdot, \cdot \rangle_*$, that is, perform the analysis inside $(W', \langle \cdot, \cdot \rangle_*)$.

Using the Gramian of $[Y'_\delta \ x'_\delta]$ with $Y'_\delta = \hat{Y} + \delta\tilde{Y}$, $x'_\delta = \hat{x} + \delta\tilde{x}$, and $\delta \in (0, 1]$ leads to the same projector $P_{V/V^{\perp_\delta}}$ on W' . Hence, the use of $P_{V/V^{\perp_*}}$ may be motivated by either increasing amplification of the directions in U or gradual neglect of the directions in U^\perp .

Figure 4.5 summarizes the construction of oblique projectors on a subspace W' of a Euclidean space W as discussed in this section. Therein, elements y_1, \dots, y_k, x of W with $Z = [y_1 \dots y_k \ x]$ are modified in such a way that the kernel of the map $Z' = [y'_1 \dots y'_k \ x']$ built from the modifications y'_1, \dots, y'_k, x' equals $\ker Z$. By lemma 2.3, the latter provide an alternative Gramian $\langle Z', Z' \rangle$ for Z , which in turn induces an alternative inner product $\langle \cdot, \cdot \rangle'$ on the subspace $W' = \text{img } Z$. Other sources of Gramian substitutes are conceivable. Finally, the induced inner product $\langle \cdot, \cdot \rangle'$ comes with a new understanding of orthogonality and an orthogonal projector onto $V = \text{span}\{y_1, \dots, y_k\}$ defined on W' . The latter enjoys $\|\cdot\|'$ -optimality, but is usually suboptimal with respect to $\|\cdot\|$. Finally, the linearity of projectors implies that the case of multiple x_1, \dots, x_q can be handled

by consideration of the linear maps $[Y x_j]$, $j \leq q$, one after the other. Alternatively, a minor extension of the arguments in this section leads to the appropriate generalization.

4.2.3. (Sub-)Optimality analysis on superspaces

This section reconsiders the case of an oblique projector P_{V/V_*^\perp} as defined in section 4.2.1. More specifically, V and U symbolize subspaces of a Euclidean space W . The former is spanned by a nontrivial sequence y_1, \dots, y_k with $Y = [y_1 \ \dots \ y_k]$. The composition \hat{Y} and \tilde{Y} of Y with the orthogonal projectors P_U and P_{U^\perp} , respectively, allow the construction of $Y_* = \hat{Y} + \tilde{Y}_*$ with $\tilde{Y}_* = \tilde{Y}P_{\ker \hat{Y}}$. Proposition 4.3 ensures that the oblique projector P_{V/V_*^\perp} onto V and along the orthogonal complement V_*^\perp of $V_* = \text{img } Y_*$ is well defined and quantifies its $\|\cdot\|$ -performance relative to the orthogonal projector P_V .

Section 4.2.2 identifies the restriction of P_{V/V_*^\perp} to $W' = [Y x]$, $x \in W$, with the orthogonal projector onto V defined on W' with respect to an alternative inner product $\langle \cdot, \cdot \rangle_*$. If only projections $P_{V/V_*^\perp} z$, $z \in W'$, are of concern, then this identification facilitates the computations. Moreover, the inequalities <4.4> apply with the principal $\langle \cdot, \cdot \rangle_*$ -angles between the complements V^\perp and V_*^\perp in W' . The bounds resulting from <4.4> together with the principal $\langle \cdot, \cdot \rangle_*$ -angles between the complements V_*^\perp and V^\perp in W remain valid, but may offer only a rough characterization of the $\|\cdot\|$ -performance of P_{V/V_*^\perp} on W' .

The following example justifies this claim. It considers a Euclidean space W spanned by $y'_1, \dots, y'_k, x', y''_1, \dots, y''_k$, and x'' , wherein $\text{img } Y'$, $Y' = [y'_1 \ \dots \ y'_k]$, and $\text{img } Y''$, $Y'' = [y''_1 \ \dots \ y''_k]$, are nontrivial. The geometry of W is such that $[Y' x']$ and $[Y'' x'']$ share their Gramian G and $U = \text{img } [Y' x'] = (\text{img } [Y'' x''])^\perp$. Herein, the goal is to project $x = x' + x''$ onto the image V of $Y = Y' + Y''$. These elements satisfy $\hat{x}_U = \hat{x} = x'$, $\hat{Y}_U = \hat{Y} = Y'$, and the Gramian of $[Y x]$ equals $2G$. Hence, the equality $\ker \hat{Y} = \ker Y = \ker \tilde{Y}$ holds and implies $\text{img } \tilde{Y}_* = \{0\}$ as well as $\text{img } \tilde{Y}_x = \{0\}$. In addition, $V_* = \text{img } Y_* = \text{img } Y' \neq V$, and the key ratio in proposition 4.3 equals $\|Y''c\|/\|Y'c\| = 1$ for all $c \in (\ker Y)^\perp$, that is, $\theta_{\min, \neq 0}(V_*^\perp, V^\perp) = \pi/4 = \theta_{\max}(V_*^\perp, V^\perp)$.

Moreover, if $z \in \text{img } [Y x] \cap U^\perp$, then $z = [Y x]c$ for some $c \in \mathbb{R}^{k+1}$ and $0 = \langle [Y' x'], z \rangle = Gc$, which implies $z = 0$. Thus, the present setting amounts to an instance of the special case $\text{img } [Y x] \cap U^\perp = \{0\}$ considered in section 4.2.2. As a consequence, the equality $P_{V/V_*^\perp} z = P_{V/V_*^\perp} z$ holds for all $z \in W'$, wherein $\langle \cdot, \cdot \rangle_*$ denotes the inner product induced by G . The two inner products $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle_*$ on W' satisfy $\langle \cdot, \cdot \rangle = 2 \langle \cdot, \cdot \rangle_*$; in particular, the equality $P_{V/V^\perp} z = P_{V/V_*^\perp} z$ holds for all $z \in W'$. In summary, the space $\text{img } [Y x]$ containing all elements of interest amounts to a subspace of $\ker(P_{V/V_*^\perp} - P_V)$. However, the elements of $V^\perp \cap (V^\perp \cap V_*^\perp)^\perp = \text{img}(Y' - Y'')$ are needed to apply proposition 4.3, which requires U to be a subspace of W .

4.3. A prediction framework

4.3.1. The prediction task

This section considers the prediction of a numerical characteristic x of m locations at n points in time. These quantities are represented by random variables $x_{t,j}$, wherein

the index (t, j) ranges over $I_x = \{1, \dots, n\} \times \{1, \dots, m\}$, $m, n \in \mathbb{N}$. Herein, *prediction* refers to the pointwise (with respect to ω) evaluation of random variables of the type $r_{t,j} + P_{V/V_*^\perp}(x_{t,j} - r_{t,j})$, wherein $r_{t,j}$, $(t, j) \in I_x$, represent an additional characteristic r of the m locations, and V, V_*^\perp denote complementary subspaces. In this section, all coordinates are considered known; their estimation—as in section 3.5—is not addressed. These predictions inherit two properties from the underlying projector P_{V/V_*^\perp} . Firstly, linearity guarantees that a prediction of $\sum_{(t,j) \in I_x} c_{t,j} x_{t,j}$, $C \in \mathbb{R}^{n \times m}$, in form of this linear combination of the predictions of $x_{t,j}$ exhibits the same structure as its ingredients but with $r_{t,j}$ replaced by $\sum_{(t,j) \in I_x} c_{t,j} r_{t,j}$. Secondly, idempotence ensures that such predictions of linear combinations equal their known value if observed without error.

The overall setting amounts to the span W of \mathbb{P} -square integrable random variables

$$z_{i,t,j}, \quad (i, t, j) \in I_z = \{1, \dots, s\} \times I_x, \quad v_{t,j}, \quad (t, j) \in I_x, \quad \bar{v}_{t,i}, \quad (t, i) \in I_{\text{obs}},$$

and the constant function $\omega \mapsto 1$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Therein, the index set $I_{\text{obs}} \subset \mathbb{N} \times \mathbb{N}$ is finite and nonempty, and $s \in \mathbb{N}$. The sequences $v_{t,j}$, $(t, j) \in I_x$, and $\bar{v}_{t,i}$, $(t, i) \in I_{\text{obs}}$, form bases of their spans $U_v = \text{span}\{v_{t,j} \mid (t, j) \in I_x\}$ and $U_{\bar{v}} = \text{span}\{\bar{v}_{t,i} \mid (t, i) \in I_{\text{obs}}\}$, respectively. The intersection of the span U_z of $z_{i,t,j}$, $(i, t, j) \in I_z$, with $\text{span}\{1\}$ equals $\{0\}$; the symbol $U_{1,z} = \text{span}\{1\} + U_z$ denotes the joint span of these variables. The random variables $z_{i,t,j}$, $(i, t, j) \in I_z$, are such that the kernel $\ker [1 \ Z_{t,j}]$, wherein $Z_{t,j} = [z_{1,t,j} \ \dots \ z_{s,t,j}]$, amounts to $\{0\}$ for all $(t, j) \in I_x$. The same applies to the intersections of pairs of subspaces $U' \neq U''$ with $U', U'' \in \{U_{1,z}, U_v, U_{\bar{v}}\}$.

The definition $\langle x, y \rangle = \int x(\omega)y(\omega)\mathbb{P}(d\omega) = \mathbb{E}xy$ for every pair $x, y \in W$ endows this linear space with a Euclidean space structure. Herein, the inner product is such that the sequence $\bar{v}_{t,i}$, $(t, i) \in I_{\text{obs}}$, provides an orthonormal basis of its span $U_{\bar{v}}$. In addition, the three subspaces $U_{1,z}$, U_v , and $U_{\bar{v}}$ satisfy $U_{1,z}^\perp = U_v + U_{\bar{v}}$, $U_v^\perp = U_{1,z} + U_{\bar{v}}$, and $U_{\bar{v}}^\perp = U_{1,z} + U_v$. In particular, this specification implies the equalities $\langle 1, v_{t,j} \rangle = \mathbb{E}v_{t,j} = 0$ as well as $\mathbb{E}\bar{v}_{t,i} = 0$ and $\mathbb{E}\bar{v}_{t,i}^2 = 1$ for all $(t, j) \in I_x$ and $(t, i) \in I_{\text{obs}}$, respectively. Finally, the formal model is chosen relative to $\langle \cdot, \cdot \rangle$ as explained towards the end of appendix 2.a.

The random variables $z_{i,t,j}$ represent s additional numerical characteristic of the m locations at n points in time. A given $1 \leq s' \leq s$ splits these variables into two disjoint subsets $z_{i,t,j}$, $i < s'$, and $z_{i,t,j}$, $i \geq s'$, wherein (t, j) ranges over I_x . The case $s' = 1$ is possible. Then, the first set is empty and all summands consisting only of elements of that set vanish. These two groups play two different roles in the following development.

The random variables $x_{t,j}$ representing the numerical characteristic x are given by

$$x_{t,j} = \alpha_{t,j} + [Z_{a,t,j} \ Z_{b,t,j}] \begin{pmatrix} \beta_a \\ \beta_b \end{pmatrix} + v_{t,j} = \alpha_{t,j} + Z_{t,j}\beta + v_{t,j}, \quad (t, j) \in I_x, \quad \langle 4.10 \rangle$$

$$Z_{a,t,j} = [z_{1,t,j} \ \dots \ z_{s'-1,t,j}], \quad Z_{b,t,j} = [z_{s',t,j} \ \dots \ z_{s,t,j}],$$

for some $\beta \in \mathbb{R}^s$ and $\alpha_{t,j} \in \mathbb{R}$. The equalities $\ker [1 \ Z_{t,j}] = \{0\}$, $(t, j) \in I_x$, imply that the coordinates $\alpha_{t,j}, \beta$ are uniquely characterized by $\langle 4.10 \rangle$. The subsequent discussion mostly focuses on the modification $\bar{x}_{t,j} = x_{t,j} - r_{t,j}$, $(t, j) \in I_x$, with $r_{t,j} = Z_{a,t,j}\beta_a$.

Linear independence of $v_{t,j}$, $(t, j) \in I_x$, and $U_{1,z} \cap U_v = \{0\}$ imply $\ker [1 \ \bar{X}] = \{0\}$,

wherein $\bar{X} = [\bar{X}_n \dots \bar{X}_1]$ with $\bar{X}_t = [\bar{x}_{t,1} \dots \bar{x}_{t,m}]$. The following argument requires that this kernel equality continues to hold if the columns of \bar{X} are replaced by their orthogonal projections $P_U \bar{x}_{t,j} = \hat{x}_{t,j}$, $(t, j) \in I_x$, onto $U = U_{1,z} + U_{\bar{v}} = U_v^\perp$. These projections equal $\hat{x}_{t,j} = \alpha_{t,j} + Z_{b,t,j} \beta_b$. Thus, the equality $\ker[1 \hat{X}] = \{0\}$, wherein $\hat{X} = [\hat{X}_n \dots \hat{X}_1]$ with $\hat{X}_t = [\hat{x}_{t,1} \dots \hat{x}_{t,m}]$, requires some additional restrictions. Lemma 4.4 contains an appropriate condition. A proof starts on page 123 in appendix 4.a.

Lemma 4.4. *The equality $\ker[1 \hat{X}] = \{0\}$ holds if and only if*

$$\beta_b \notin B_b = \cup_{\substack{C \in \mathbb{R}^{n \times m} \\ \|C\|=1}} \ker \left(\sum_{(t,j) \in I_x} c_{t,j} Z_{b,t,j} \right) \subset \mathbb{R}^{s-s'+1} .$$

The kernels corresponding to indexes $C = B_{t,j}$ equal $\ker Z_{b,t,j} = \{0\}$ for all $(t, j) \in I_x$, wherein $B_{t,j}$ represents the t, j -th standard basis element of $\mathbb{R}^{n \times m}$ as in example (b) in section 2.1.1. However, the analogous equality in case of a general unit length matrix $C \in \mathbb{R}^{n \times m}$ necessitates some additional linear independence conditions. For example, if the sequence $z_{i,t,j}$, $s' \leq i$, $(t, j) \in I_x$ exhibits linear independence, then $B_b = \{0\}$ and the requirement reduces to $\beta_b \neq 0$. The latter scenario places a strong requirement on the random variables $z_{i,t,j}$, $s' \leq i$, $(t, j) \in I_x$, and therefore requires only the minimal condition $\beta_b \neq 0$ on β_b . The remainder of this section assumes that $\beta_b \notin B_b$ holds.

The givens of the initially mentioned prediction task include the images $z_{i,t,j}(\omega)$ under the random variables $z_{i,t,j}$, $i < s'$, $(t, j) \in I_x$, as well as the images $y_{t,i}(\omega)$ of $\omega \in \Omega$ under

$$y_{t,i} = X a_{t,i} + \rho_{t,i} \bar{v}_{t,i} , \quad (t, i) \in I_{\text{obs}} . \quad \langle 4.11 \rangle$$

The adjustment of the formal representation to $\langle \cdot, \cdot \rangle$ as explained in appendix 2.a ensures that the just mentioned images reflect the geometry of the space $(W, \langle \cdot, \cdot \rangle)$. Subsequently, these images and the corresponding random variables are referred to as observations (or data) and observables, respectively. The random variables in $\langle 4.11 \rangle$ amount to linear combinations of $x_{t,j}$, $(t, j) \in I_x$, and $\bar{v}_{t,i}$, $(t, i) \in I_{\text{obs}}$, with coordinates $0 \neq a_{t,i} \in \mathbb{R}^{nm}$ and $\rho_{t,i} \geq 0$. The index set I_{obs} has the form $\cup_{t \leq n} (\{t\} \times \{1, \dots, k_t\})$ with $k_t \in \mathbb{N} \cup \{0\}$ being the number of observations at t . If $k_t = 0$, then $\{t\} \times \{1, \dots, k_t\}$ equals the empty set; however, $I_{\text{obs}} \neq \emptyset$ is assumed below. The final summand in $\langle 4.11 \rangle$ embodies an observation error with variance $\rho_{t,i}^2$. The individual observation errors $\rho_{t,i} \bar{v}_{t,i}$ are pairwise orthogonal. Consequently, superfluous observables of the type $\langle 4.11 \rangle$ in form of linear combinations of other observables of the same type are possible only if all corresponding error variances $\rho_{t,i}^2$ equal zero. The case $a_{t,i} = a_{t,j}$ for some $i < j \leq k_t$ together with $\max\{\rho_{t,i}, \rho_{t,j}\} > 0$ represents the availability of different observations of the same element $X a_{t,i}$. Subsequently, the focus is on the modified observables $\bar{y}_{t,i} = \bar{X} a_{t,i} + \rho_{t,i} \bar{v}_{t,i}$, $(t, i) \in I_{\text{obs}}$, which amount to linear combinations of the observables $y_{t,i}$, $(t, i) \in I_{\text{obs}}$, and $z_{i,t,j}$, $i < s'$, $(t, j) \in I_x$; thus, their images of ω are available, too.

The modified observables $\bar{y}_{t,i}$, $(t, i) \in I_{\text{obs}}$, together with the constant function 1 span the image V of the projector underlying the predictions. Initially, this projector is only defined on the superspace $W' = \text{img}[1 \bar{Y} \bar{X}]$ of V , wherein $\bar{Y} = [\bar{Y}_n \dots \bar{Y}_1]$ with

$\bar{Y}_t = [\bar{y}_{t,1} \cdots \bar{y}_{t,k_t}]$, $1 \leq t \leq n$. If $k_t = 0$ for some $t \leq n$, then the respective \bar{Y}_t is not defined and missing from \bar{Y} . The inequality $I_{\text{obs}} \neq \emptyset$ ensures that at least one \bar{Y}_t is present. The linear maps \hat{Y} and \bar{V} are defined in analogy with respect to $\hat{y}_{t,i} = P_U y_{t,i} = \hat{X} a_{t,i} + \rho_{t,i} \bar{v}_{t,i}$ with $U = U_{1,z} + U_{\bar{v}}$ and $\bar{v}_{t,i}$, $(t, i) \in I_{\text{obs}}$. The equalities $\ker [1 \ \bar{X} \ \bar{V}] = \{0\} = \ker [1 \ \hat{X} \ \bar{V}]$, which are due to the assumption of the condition $\beta_b \notin B_b$ in lemma 4.4 and the equality $U_{1,z} \cap U_{\bar{v}} = \{0\}$, imply that the kernels of the two linear maps

$$[1 \ \bar{Y} \ \bar{X}] = [1 \ \bar{X} \ \bar{V}] \begin{pmatrix} 1 & & \\ & A^\top & \\ & S^\top & I \end{pmatrix} \quad \text{and} \quad [1 \ \hat{Y} \ \hat{X}] = [1 \ \hat{X} \ \bar{V}] \begin{pmatrix} 1 & & \\ & A^\top & \\ & S^\top & I \end{pmatrix} \quad \langle 4.12 \rangle$$

coincide with the kernel of their identical second factor. Therein, I symbolizes the $nm \times nm$ identity matrix, whose columns amount to the standard basis e_1, \dots, e_{nm} of \mathbb{R}^{nm} . The number of columns $k = \sum_{t \leq n} k_t > 0$ of Y equals the overall number of observations of the type $y_{t,i}$. Moreover, the rows of $A \in \mathbb{R}^{k \times mn}$ and the diagonal entries of the diagonal matrix $S \in \mathbb{R}^{k \times k}$ amount to $a_{t,i}$ and $\rho_{t,i}$, respectively, arranged in appropriate order.

Lemma 2.3 ensures that the Gramian \hat{G} of $[1 \ \hat{Y} \ \hat{X}]$ induces an additional inner product $\langle \cdot, \cdot \rangle_*$ on W' by identification of \hat{G} with the $\langle \cdot, \cdot \rangle_*$ -Gramian of $[1 \ \bar{Y} \ \bar{X}]$. The second equality in $\langle 4.12 \rangle$ shows that knowledge of the aggregation matrix A , the observation error matrix S , and the Gramian of $[1 \ \hat{X}]$ suffices to construct the Gramian \hat{G} and thereby $\langle \cdot, \cdot \rangle_*$. The latter leads to an additional orthogonal complement $V^{\perp*}$ of V in W' , which in turn defines the projector $P_{V/V^{\perp*}}$ (on W') underlying the predictions.

These predictions are of the form $r_{t,j}(\omega) + (P_{V/V^{\perp*}} \bar{x}_{t,j})(\omega)$, $(t, j) \in I_x$, wherein $\bar{x}_{t,j} = x_{t,j} - r_{t,j}$ and $r_{t,j} = Z_{a,t,j} \beta_a$. This construct enjoys the initially mentioned properties. In fact, a prediction of a linear combination $\sum_{(t,j) \in I_x} c_{t,j} x_{t,j}$ —obtained as the same linear combination of the individual predictions—amounts to the image of ω under

$$\left[\sum_{t,j} c_{t,j} Z_{a,t,j} \right] \beta_a + P_{V/V^{\perp*}} \left(\sum_{t,j} c_{t,j} x_{t,j} - \left[\sum_{t,j} c_{t,j} Z_{a,t,j} \right] \beta_a \right)$$

and thus equals $y_{t,i}$ if $c_{t,j}$ coincide with the entries of a row $a_{t,i}$ of A in $\langle 4.12 \rangle$ and $\rho_{t,i} = 0$.

Calculation of the predictions requires evaluation (at ω) of the $\langle \cdot, \cdot \rangle_*$ -orthogonal projection $P_{V/V^{\perp*}} \bar{x}_{t,j}$ of $\bar{x}_{t,j}$ onto $V = \text{img} [1 \ \bar{Y}]$ for all $(t, j) \in I_x$. The techniques of section 2.2 and 2.4 are crucial to this endeavor. More specifically, lemma 2.2 guarantees the existence of a $\langle \cdot, \cdot \rangle_*$ -representation of the columns of $[1 \ \bar{Y} \ \bar{X}]$ in form of the columns of a coordinate matrix R as in figure 4.6. Therein, the presence of 1 in the unitary map Q —in form of the function $\omega \mapsto 1$ —and R —as a number, respectively, follows from the equality $\|1\|_* = 1$, wherein $\|\cdot\|_*$ denotes the norm induced by $\langle \cdot, \cdot \rangle_*$ on W' . As a consequence, the vectors $r_{y,1} \in \mathbb{R}^k$ and $r_{x,1} \in \mathbb{R}^{nm}$ in the first row of R consists of entries $\langle 1, \bar{y}_{t,i} \rangle_* = \langle 1, \hat{y}_{t,i} \rangle = \mathbb{E} \hat{y}_{t,i}$ and $\langle 1, \bar{x}_{t,j} \rangle_* = \langle 1, \hat{x}_{t,j} \rangle = \mathbb{E} \hat{x}_{t,j}$, respectively. The latter equals $\langle 1, \hat{x}_{t,j} \rangle = \langle 1, \alpha_{t,j} + Z_{b,t,j} \beta_b \rangle = \alpha_{t,j} + \sum_{i=s'}^s \beta_i \mathbb{E} z_{i,t,j}$, wherein $\beta_b = (\beta_{s'}, \dots, \beta_s)$. In particular, the entries of $r_{x,1}$ usually differ from $\alpha_{t,j}$ as defined in $\langle 4.10 \rangle$.

Furthermore, the columns $q_1^*, \dots, q_{k'}^*$ of Q_y form an $\langle \cdot, \cdot \rangle_*$ -orthonormal basis of the

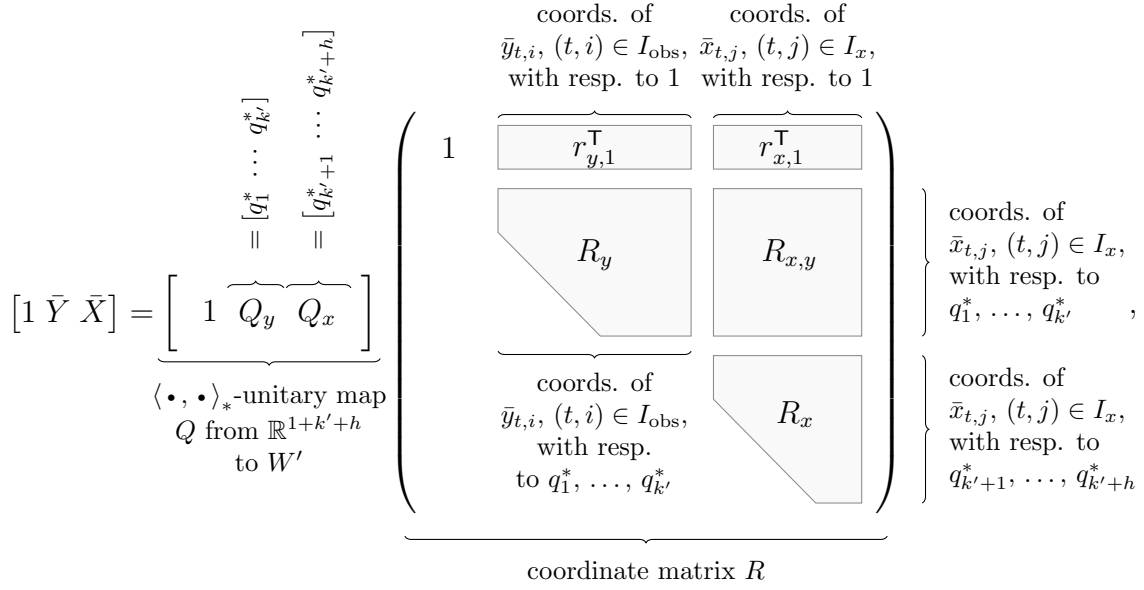


Figure 4.6

The figure visualizes the structure of a representation of $[1 \ \bar{Y} \ \bar{X}]$ as considered in lemma 2.2 and with respect to $\langle \cdot, \cdot \rangle_*$. Herein, the rank of $P_{(\text{span}\{1\})^{\perp*}/\text{span}\{1\}}\bar{Y} = \bar{Y} - r_{y,1}^\top$ and $P_{V^{\perp*}/V}\bar{X} = \bar{X} - r_{x,1}^\top - Q_y R_{x,y}$ with $V = \text{img}[1 \ \bar{Y}]$ is denoted by k' and h , respectively.

image of $P_{(\text{span}\{1\})^{\perp*}/\text{span}\{1\}}\bar{Y} = \bar{Y} - r_{y,1}^\top$, and $Q_y R_{x,y}$ contains the $\langle \cdot, \cdot \rangle_*$ -orthogonal projections of the columns of $P_{(\text{span}\{1\})^{\perp*}/\text{span}\{1\}}\bar{X} = \bar{X} - r_{x,1}^\top$ onto that space. The basis elements $q_1^*, \dots, q_{k'}^*$ can be evaluated (at ω) based on $\bar{Y} - r_{y,1}^\top = Q_y R_y$, which implies—due to the choice of the formal model—the pointwise (with respect to ω) equality

$$R_y^\top \begin{pmatrix} q_1^*(\omega) \\ \vdots \\ q_{k'}^*(\omega) \end{pmatrix} = \begin{pmatrix} \bar{y}_{n,1}(\omega) \\ \vdots \\ \bar{y}_{1,k_1}(\omega) \end{pmatrix} - r_{y,1}. \quad \langle 4.13 \rangle$$

Herein, the number of observations k_1 and k_n of the first and n -th time point are assumed to be nonzero to simplify the presentation. The row echelon matrix $R_y \in \mathbb{R}^{k' \times k}$, $k' = \text{rk}(\bar{Y} - r_{y,1}^\top) \leq k$, provides a Cholesky factor of the $\langle \cdot, \cdot \rangle_*$ -Gramian of $\bar{Y} - r_{y,1}^\top$; its rows are linearly independent—as elements of \mathbb{R}^k . Consequently, the equality $\langle 4.13 \rangle$ uniquely determines the row $q_y = Q_y(\omega) = (q_1^*(\omega), \dots, q_{k'}^*(\omega))$ of Q_y . Finally, the columns of $P_{V^{\perp*}/V}\bar{X} = r_{x,1}^\top + Q_y R_{x,y}$ equal the $\langle \cdot, \cdot \rangle_*$ -orthogonal projections $P_{V^{\perp*}/V}\bar{x}_{t,j}$, $(t, j) \in I_x$. Thus, the entries of $r_{x,1}^\top + R_{x,y}^\top q_y$ supply the required images $(P_{V^{\perp*}/V}\bar{x}_{t,j})(\omega)$.

By virtue of section 2.4.2, these computations require only the data and the $\langle \cdot, \cdot \rangle_*$ -Gramian of $[1 \ \bar{Y} \ \bar{X}]$, which may be recovered via $\langle 4.12 \rangle$ from A , S , the $\langle \cdot, \cdot \rangle_*$ -Gramian of $P_{(\text{span}\{1\})^{\perp*}/\text{span}\{1\}}\hat{\bar{X}}$, and $r_{x,1}$ whose entries equal $\langle 1, \hat{x}_{t,j} \rangle = \mathbb{E}\hat{x}_{t,j}$, $(t, j) \in I_x$. Section 4.4.3 shows how to exploit a particular structure of A and the $\langle \cdot, \cdot \rangle_*$ -Gramian of $P_{(\text{span}\{1\})^{\perp*}/\text{span}\{1\}}\bar{x}_{t,j}$, $(t, j) \in I_x$ when evaluating the $\langle \cdot, \cdot \rangle_*$ -orthogonal projections of $P_{(\text{span}\{1\})^{\perp*}/\text{span}\{1\}}\bar{x}_{t,j}$, $(t, j) \in I_x$, onto $\text{span}\{P_{(\text{span}\{1\})^{\perp*}/\text{span}\{1\}}\bar{y}_{t,i} \mid (t, i) \in I_{\text{obs}}\}$.

4.3.2. (Sub-)optimality of predictions

An alternative characterization of $P_{V/V^{\perp*}}$ facilitates its comparison in terms of $\|\cdot\|$ with the $\langle \cdot, \cdot \rangle$ -orthogonal projector $P_{V/V^{\perp}}$. More specifically, the current framework fits the scenario considered in section 4.2.1. In particular, the composition of $Y' = [1 \ \bar{Y}]$ with the orthogonal projector P_U onto the subspace $U = U_{1,z} + U_{\bar{v}}$ of W equals $\hat{Y}' = [1 \ \hat{Y}]$. As observed below <4.12>, the linear map \hat{Y}' shares its kernel with Y' . Consequently, $Y'_* = \hat{Y}' + P_{U^{\perp}} Y' P_{\ker \hat{Y}'} = \hat{Y}'$, and proposition 4.3 ensures that there exists an oblique projector $P_{V/V_*^{\perp}}$ (on W) onto $V = \text{img } Y'$ and along the orthogonal complement of $V_* = \text{img } Y'_* = \text{img } \hat{Y}'$. The restriction of the latter to W' coincides with $P_{V/V^{\perp*}}$ as defined in section 4.3.1. In fact, the linear map $[\hat{Y}' \ \hat{X}]$ amounts to the composition $P_U [Y' \ \bar{X}]$ and shares its kernel with $[Y' \ \bar{X}]$. Consequently, if $z' \in \text{img } [Y' \ \bar{X}] \cap U^{\perp}$, then $z' = [Y' \ \bar{X}] c$ for some $c \in \mathbb{R}^{1+k+nm}$, and its projection onto U equals $0 = P_U z' = [\hat{Y}' \ \hat{X}] c$. Thus, $c \in \ker [\hat{Y}' \ \hat{X}] = \ker [Y' \ \bar{X}]$, and therefore $z' = 0$. The resulting equality $\text{img } [Y' \ \bar{X}] \cap U^{\perp} = \{0\}$ identifies the present setting as an instance of the corresponding special case considered in section 4.2.2. Hence, the restriction of the oblique projector $P_{V/V_*^{\perp}}$ to $W' = \text{img } [1 \ \bar{Y} \ \bar{X}]$ equals the $\langle \cdot, \cdot \rangle_*$ -orthogonal projector $P_{V/V^{\perp*}}$.

In particular, proposition 4.3 becomes applicable to $P_{V/V^{\perp*}}$ and is relevant as

$$\begin{aligned} & \|x_{t,j} - [r_{t,j} + P_{V/V^{\perp*}}(x_{t,j} - r_{t,j})]\|^2 = \|\bar{x}_{t,j} - P_{V/V^{\perp*}} \bar{x}_{t,j}\|^2 \\ & = \|P_{V^{\perp}} \bar{x}_{t,j}\|^2 + \|(P_V - P_{V/V_*^{\perp}}) \bar{x}_{t,j}\|^2 \leq [1 + \tan^2 \theta_{\max}(V^{\perp}, V_*^{\perp})] \|P_{V^{\perp}} \bar{x}_{t,j}\|^2, \end{aligned} \quad <4.14>$$

wherein the final inequality is due to <4.4>. Corollary 4.5 allows the quantification of the multiplier $1 + \tan^2 \theta_{\max}(V^{\perp}, V_*^{\perp})$ in <4.14>. A proof of this result starts on page 124 in appendix 4.a. Its statement uses the notion of a variance matrix $\text{var}(z)$ of a random vector $z = (z_1, \dots, z_j)$, $j \in \mathbb{N}$, which is defined in example (e) in section 2.4.1.

Corollary 4.5. *In the above setting, in particular, with $I_{\text{obs}} \neq \emptyset$ and $\beta_b \notin B_b$, one has $\text{img } A^{\top} \neq \{0\}$. Moreover, the kernel of the variance matrix $\text{var}(\hat{x})$ of the random vector \hat{x} , which contains the columns of the linear map \hat{X} (in the same order), equals $\{0\}$. Finally, the two subspaces V and V_* differ and the inequalities*

$$\begin{aligned} \sup_{c \in \text{img } A, c \neq 0} \frac{\langle A^{\top} c, \text{var}(\tilde{x}) A^{\top} c \rangle}{\langle A^{\top} c, (\text{var}(\hat{x}) + S^2) A^{\top} c \rangle} & \leq \tan^2 \theta_{\max}(V^{\perp}, V_*^{\perp}) \leq \sup_{a \in \text{img } A^{\top}, a \neq 0} \frac{\langle a, \text{var}(\tilde{x}) a \rangle}{\langle a, \text{var}(\hat{x}) a \rangle}, \\ \inf_{c \in \text{img } A, c \neq 0} \frac{\langle A^{\top} c, \text{var}(\tilde{x}) A^{\top} c \rangle}{\langle A^{\top} c, (\text{var}(\hat{x}) + S^2) A^{\top} c \rangle} & \leq \tan^2 \theta_{\min, \neq 0}(V^{\perp}, V_*^{\perp}) \leq \inf_{a \in \text{img } A^{\top}, a \neq 0} \frac{\langle a, \text{var}(\tilde{x}) a \rangle}{\langle a, \text{var}(\hat{x}) a \rangle} \end{aligned}$$

hold, wherein \tilde{x} is defined in analogy with \hat{x} but with respect to $\tilde{X} = P_{U^{\perp}} \bar{X}$.

If $\ker A^{\top} \subset \ker S$, in particular, if $\ker A^{\top} = \{0\}$, then the lower bounds for the squared tangents $\tan^2 \theta_{\max}(V^{\perp}, V_*^{\perp})$ and $\tan^2 \theta_{\min, \neq 0}(V^{\perp}, V_*^{\perp})$ hold with equality.

A few comments on corollary 4.5 are in order. Firstly, the interest lies in the predictions, that is, the values of the translated—by $r_{t,j}$ —projections of $\bar{x}_{t,j}$, $(t, j) \in I_x$, at a

single $\omega \in \Omega$. However, the bounds resulting from corollary 4.5 refer to the norm $\|\cdot\|$, which merely provides an average (across ω) distance measure. Secondly, the findings of section 4.2.3 apply as solely projections in W' are of concern, that is, the bounds may be rather “conservative”. Finally, if $\ker A^\top = \{0\}$, then an increase of the error variances $\rho_{t,i}^2$, $(t, i) \in I_{\text{obs}}$, that is, the diagonal elements of S , decreases the first factor in the final term in <4.14>, but also increases the $\|\cdot\|$ -length of the residuals resulting from orthogonal projection onto V , that is, the second factor of the final term in <4.14>.

4.4. Prediction algorithms

4.4.1. Recursive computations via basis changes

This section considers nontrivial sequences of real-valued functions $x_{t,j}$, $(t, j) \in I_x = \{1, \dots, n\} \times \{1, \dots, m\}$ with $m, n \in \mathbb{N}$, and $v_{t,i}$, $(t, i) \in I_{\text{obs}}$, defined on a common set Ω , wherein the second index set $I_{\text{obs}} \subset \mathbb{N} \times \mathbb{N}$ is nonempty and finite. An inner product $\langle \cdot, \cdot \rangle$ equips their joint span W with a Euclidean space structure. This structure is such that the functions $v_{t,i}$, $(t, i) \in I_{\text{obs}}$, form an orthonormal basis of their span U_v . In addition, the subspace U_v equals the orthogonal complement of $U_x = \text{span}\{x_{t,j} \mid (t, j) \in I_x\}$ in W . Further relevant elements of the linear space W come in the form of

$$\underbrace{[y_{t,1} \ \cdots \ y_{t,k_t}]}_{Y_t \in W^{\times k_t}} = X \underbrace{[a_{t,1} \ \cdots \ a_{t,k_t}]}_{A_t^\top \in \mathbb{R}^{nm \times k_t}} + \underbrace{[v_{t,1} \ \cdots \ v_{t,k_t}]}_{V_t \in W^{\times k_t}} S_t^\top, \quad t \leq n, \quad <4.15>$$

wherein $X = [X_1 \ \cdots \ X_n]$, $X_t = [x_{t,1} \ \cdots \ x_{t,m}]$, $S_t \in \mathbb{R}^{k_t \times k_t}$, and $k_t \in \mathbb{N} \cup \{0\}$. If $k_t = 0$, then the quantities Y_t , A_t , V_t , and S_t disappear. The index set I_{obs} equals the union of the sets $\{t\} \times \{1, \dots, k_t\}$ with $t \leq n$. If $k_t = 0$, then the t -th of the latter sets amounts to the empty set \emptyset ; however, the following discussion presupposes $I_{\text{obs}} \neq \emptyset$.

Lemma 2.2 allows the representation of the linear map X in form $Q_X R_X$, wherein Q_X denotes an unitary map from $\mathbb{R}^{\text{rk } X}$ to U_x , and $R_X \in \mathbb{R}^{\text{rk } X \times mn}$ exhibits row echelon form. This expression for X together with the specification in <4.15> implies

$$\begin{bmatrix} \overbrace{Y_1 \ \cdots \ Y_n}^Y \\ \underbrace{Q_X \ V_1 \ \cdots \ V_n}_V \end{bmatrix} = \underbrace{\begin{pmatrix} \overbrace{R_X A_1^\top}^{\text{coords. of } Y_1} & \cdots & \overbrace{R_X A_n^\top}^{\text{coords. of } Y_n} & \overbrace{R_X}^{\text{coords. of } X} \\ \hline S_1^\top & & & \\ \hline & \ddots & & \\ \hline & & S_n^\top & \end{pmatrix}}_{\text{coordinate matrix } B} \begin{cases} \text{coord. vec.} \\ \text{with resp.} \\ \text{to } Q_X \\ \hline \text{coord. vec.} \\ \text{with resp.} \\ \text{to } V_1 \\ \hline \text{coord. vec.} \\ \text{with resp.} \\ \text{to } V_n \end{cases}, \quad <4.16>$$

wherein $k_1, k_n \geq 1$ is assumed for the sake of presentation. The linear map $[Q_X \ V]$

If $k_2 \geq 1$ and $\text{img } Y_2 \not\subset \text{img } Y_1$, then the columns $q_{k'_1+1}, \dots, q_{k'_1+k'_2}$ of Q_2 in <4.17> provide an orthonormal basis of the k'_2 -dimensional subspace $\text{img } P_{(\text{img } Y_1)^\perp} Y_2$ of W . The evaluation of these columns uses the already available vector $Q_1(\omega)$ alongside the equality $Y_2 = Q_1 R_{1,2} + Q_2 R_{2,2}$, which is implied by the representation in <4.17>. More specifically, the resulting pointwise (with respect to ω) relation

$$R_{2,2}^\top \begin{pmatrix} q_{k'_1+1}(\omega) \\ \vdots \\ q_{k'_1+k'_2}(\omega) \end{pmatrix} = \begin{pmatrix} y_{2,1}(\omega) \\ \vdots \\ y_{2,k_2}(\omega) \end{pmatrix} - R_{1,2}^\top \begin{pmatrix} q_1(\omega) \\ \vdots \\ q_{k'_1}(\omega) \end{pmatrix},$$

wherein $k'_1 > 0$ is assumed, uniquely determines $Q_2(\omega)$ due to the row echelon form of $R_{2,2}$. The evaluation of further basis elements proceeds in analogy. Finally, combining the images $q_i(\omega)$, $i \leq k' = \text{rk } Y$, with the coordinates of $x_{t,j}$, $(t, j) \in I_x$, with respect to columns $q_1, \dots, q_{k'}$ of $[Q_1 \cdots Q_n]$ in <4.17> yields the required images $(P_V x_{t,j})(\omega)$ of w under the orthogonal projections of $x_{t,j}$, $(t, j) \in I_x$, onto the subspace V .

Section 4.4.3 shows how a more refined recursive strategy allows to exploit a special and complementary structure of $\langle X, X \rangle$ and the aggregation matrix A .

4.4.2. Basis changes by reflection

The section considers a finite sequence z_1, \dots, z_ℓ , $\ell \in \mathbb{N}$, of linearly independent elements of a Euclidean space W with inner product $\langle \cdot, \cdot \rangle$. The sequence v_1, \dots, v_ℓ forms an orthonormal basis of $W' = \text{span}\{z_j \mid j \leq \ell\}$, and the i, j -th entry of $B \in \mathbb{R}^{\ell \times \ell}$ equals $b_{i,j} = \langle v_i, z_j \rangle$. Consequently, one has the equality

$$[z_1 \cdots z_\ell] = \underbrace{[v_1 \cdots v_\ell]}_{\text{unitary map } V} \begin{pmatrix} \overbrace{b_{1,1} \cdots b_{1,\ell}}^{\text{coords. of } z_1} \\ \vdots \\ \overbrace{b_{\ell,1} \cdots b_{\ell,\ell}}^{\text{coords. of } z_\ell} \end{pmatrix} \begin{matrix} \left. \begin{matrix} \langle v_1, z_1 \rangle \\ \vdots \\ \langle v_\ell, z_1 \rangle \end{matrix} \right\} \text{ with resp.} \\ \left. \begin{matrix} \langle v_1, z_\ell \rangle \\ \vdots \\ \langle v_\ell, z_\ell \rangle \end{matrix} \right\} \text{ with resp.} \end{matrix} \quad \text{<4.19>}$$

Therein, the columns of B provide a representation of z_1, \dots, z_ℓ in the sense of section 2.2.1. In this setting, the goal is to obtain a representation as in lemma 2.2, that is,

$$[z_1 \cdots z_\ell] = \underbrace{[q_1 \cdots q_\ell]}_{\text{unitary map from } \mathbb{R}^\ell \text{ to } W' = \text{span}\{z_1, \dots, z_\ell\}} \begin{pmatrix} \boxed{r_{1,1}} & \cdots & \boxed{r_{1,\ell}} \\ & \ddots & \\ & & \boxed{r_{\ell,\ell}} \end{pmatrix} \begin{matrix} \text{triangular coordinate} \\ \text{matrix } R \text{ with} \\ \text{nonzero diagonal elements} \end{matrix}.$$

Reflections about hyperplanes in \mathbb{R}^ℓ of the form $\{\langle u, \bullet \rangle = 0\}$ with nonzero u , that is, orthogonal complements of subspaces $\text{span}\{u\}$, are key to an alternative computational strategy, which yields the same output as the Gram-Schmidt orthogonalization <2.3> when applied to z_1, \dots, z_ℓ with suitable sign choices. Such reflections amount to unitary maps from \mathbb{R}^ℓ to \mathbb{R}^ℓ . Thus, their columns form orthonormal bases of \mathbb{R}^ℓ , which implies that the Gramian $\langle O, O \rangle = O^\top O$ of such a reflection O equals the $\ell \times \ell$ identity matrix I .

Lemma 4.6 asserts that the alternative strategy leads to the same output as <2.3>. More specifically, a sequence of reflections O_1, \dots, O_ℓ recovers the upper triangular coordinate matrix R produced by a Gram-Schmidt orthogonalization via $R = O_\ell \cdots O_1 B$.

Lemma 4.6. *If a sequence z_1, \dots, z_ℓ of linearly independent elements of a Euclidean space W exhibits a representation as in <4.19>, wherein v_1, \dots, v_ℓ amounts to an orthonormal basis of $W' = \text{span}\{z_i \mid i \leq \ell\}$, then there exist reflections O_1, \dots, O_ℓ such that*

$$VB = \underbrace{VO_1^\top O_2^\top \cdots O_\ell^\top}_{Q=[q_1 \cdots q_\ell]} \underbrace{O_\ell \cdots O_1 B}_R = QR,$$

wherein Q, R denote a unitary map and an upper triangular matrix, respectively, which equal the output of <2.3> when applied to z_1, \dots, z_ℓ with suitable sign choices.

If the coordinate matrix R with respect to q_1, \dots, q_ℓ of z_1, \dots, z_ℓ is available in addition to the images $z_1(\omega), \dots, z_\ell(\omega)$, then $q_1(\omega), \dots, q_\ell(\omega)$ may be obtained based on

$$\begin{pmatrix} r_{1,1} & & & \\ r_{1,2} & r_{2,2} & & \\ \vdots & \vdots & \ddots & \\ r_{1,\ell} & r_{2,\ell} & \cdots & r_{\ell,\ell} \end{pmatrix} \begin{pmatrix} q_1(\omega) \\ q_2(\omega) \\ \vdots \\ q_\ell(\omega) \end{pmatrix} = \begin{pmatrix} z_1(\omega) \\ z_2(\omega) \\ \vdots \\ z_\ell(\omega) \end{pmatrix},$$

wherein $\ell > 2$ is assumed for the sake of presentation. More specifically, $r_{j,j} \neq 0, j \leq \ell$, implies the equalities $q_1(\omega) = z_1(\omega)/r_{1,1}$, $q_2(\omega) = (z_2(\omega) - r_{1,2}q_1(\omega))/r_{2,2}$, and so forth. The computational recipe <4.20> exploits these relations to calculate $q'_i = q_i(\omega)$.

$$\begin{array}{l|l} 1 & z_j^{(0)} = z_j(\omega), \quad j \leq \ell \\ 2 & \mathbf{for} \ i = 1, \dots, \ell \\ 3 & \quad q'_i = z_i^{(i-1)}/r_{i,i} \\ 4 & \quad \mathbf{for} \ j = i + 1, \dots, \ell \\ 5 & \quad \quad z_j^{(i)} = z_j^{(i-1)} - r_{i,j}q'_i \end{array} \quad \langle 4.20 \rangle$$

The remainder of this section justifies the assertion of lemma 4.6

Panel (A) of figure 4.7 illustrates the *reflection* of a nonzero element $x \in \mathbb{R}^3$ about a hyperplane $H = \{\langle u, \bullet \rangle = 0\}$ into its mirror image x' . This transformation amounts to first projecting x onto H to obtain the orthogonal projection $\hat{x} = P_H x$ and thereby the corresponding residual $\tilde{x} = x - P_H x = P_{H^\perp} x$. Next, this residual is subtracted from \hat{x} to reach x' on the opposite side of—but with equal distance $\inf_{y \in H} \|x - y\| = \|\tilde{x}\|$ to—the

reflection




Figure 4.7

The figure illustrates the geometry of a reflection and shows that the image of a given $b_1 \in \mathbb{R}^2$ under a suitable reflection lies on the first coordinate axis. Panel (A) illustrates the reflection of a nonzero element $x \in \mathbb{R}^3$ about a hyperplane $H = (\text{span}\{u\})^\perp$. Panel (B) shows the transformation of a nonzero $b_1 \in \mathbb{R}^2$ into a multiple $-\|b_1\|e_1$ of the first standard basis element. The transition from b_1 to $-\|b_1\|e_1$ proceeds by reflection about $H = (\text{span}\{b_1 + \|b_1\|e_1\})^\perp$.

subspace H . Linearity of projectors and Pythagoras's theorem ensure that the map given by $x \mapsto x' = \hat{x} - \tilde{x}$, called *Householder transform*, is linear and length preserving. Moreover, the equalities $P_H(\hat{x} - \tilde{x}) = \hat{x}$ and $P_{H^\perp}(\hat{x} - \tilde{x}) = -\tilde{x}$ show that the reflection of $x' = \hat{x} - \tilde{x}$ about H equals $(x')' = \hat{x} - (-\tilde{x}) = x$. Consequently, Householder transforms are also bijective and therefore provide unitary maps from \mathbb{R}^ℓ to \mathbb{R}^ℓ .

Householder transform

The utility of reflections comes from their ability to transform a point x into any target x' of the same length by adequate choice of H . In fact, a reflection modifies x merely with respect to its component \tilde{x} in H^\perp . Thus, as shown in panel (A) of figure 4.7 the difference $x - x'$ lies in the one dimensional subspace H^\perp . Unit dimension ensures that if $x \neq x'$, then scaling the difference $x - x'$ leads to an orthonormal basis of H^\perp in form of $u = (x - x')/\|x - x'\|$. The residual $x - u\langle u, x \rangle$ from orthogonally projecting x onto H^\perp equals the orthogonal projection of x onto H . Thus, a reflection transforming x into x' is given by $\text{id} - 2u\langle u, \cdot \rangle$, wherein id symbolizes the identity map on \mathbb{R}^ℓ . If $x = x'$, then id maps x to x' and setting u to zero generalizes the previous construct.

The situation of lemma 4.6 is given by $Z = VB$, wherein the kernel of $Z = [z_1 \cdots z_\ell]$ equals $\{0\}$, $V = [v_1 \cdots v_\ell]$ is unitary, and $b_{i,j} = \langle v_i, z_j \rangle$. Linear independence of z_1, \dots, z_ℓ , $\ell \geq 1$, implies that z_1 and thus its length $\|z_1\|$ is nonzero. The two elements $\pm\|b_1\|e_1 - e_1$ being the first standard basis element of \mathbb{R}^ℓ as in example (a) of section 2.1—share their length with $b_1 = (b_{1,1}, \dots, b_{\ell,1})$ and hence z_1 . Therefore, a suitable Householder transform O_1 yields $O_1b_1 = \pm\|b_1\|e_1$. Panel (B) of figure 4.7 illustrates the construction of O_1 for the choice of $-\|b_1\|e_1$ and a nonzero element b_1 of \mathbb{R}^2 . Therein, the element needed for constructing O_1 —previously denoted by u —follows by scaling $b_1 + \|b_1\|e_1$.

The composition VO_1 of two unitary maps is itself of that kind, and consequently its images of the standard basis elements e_1, \dots, e_m form another orthonormal basis of $W' = \text{span}\{z_1, \dots, z_\ell\}$. Moreover, the Householder transform O_1 represents the multiplication with the symmetric matrix $I - 2uu^\top$, wherein u denotes an orthonormal basis of the respective H^\perp . The resulting equality $O_1^\top = O_1$ proves—once more—that $O_1O_1 = I$ with I being the identity matrix. This equality underlies the first basis change given by

$$Z = VB = VO_1 \overbrace{O_1B}^{\substack{\text{coords.} \\ \text{of } z_1 \\ \text{coords.} \\ \text{of } z_2 \\ \dots \\ \text{coords.} \\ \text{of } z_\ell}} \left(\begin{array}{c|c|c|c} \overbrace{r_{1,1}} & \overbrace{r_{1,2}} & \dots & \overbrace{r_{1,\ell}} \\ \hline & b_{2,2}^{(1)} & \dots & b_{2,\ell}^{(1)} \\ \hline & \vdots & \ddots & \vdots \\ \hline & b_{\ell,2}^{(1)} & \dots & b_{\ell,\ell}^{(1)} \end{array} \right) \left. \begin{array}{l} \text{with respect to } q_1 \\ \\ \text{with respect} \\ \text{to } v_2^{(1)}, \dots, v_\ell^{(1)} \end{array} \right\}$$

This display assumes $\ell > 1$. In any case, one has $z_1 = r_{1,1}q_1$, that is, $q_1 = \pm z_1/\|z_1\|$ and $r_{1,1} = \pm\|z_1\|$. If $\ell \geq 2$, then the orthogonal projection of z_j , $j \geq 2$, onto $\text{span}\{z_1\}$ equals $r_{1,j}q_1$. Hence, the quantities q_1 and $r_{1,j}$, $j \leq \ell$, coincide with those generated by the Gram-Schmidt orthogonalization <2.3>. Moreover, the sequence $v_2^{(1)}, \dots, v_\ell^{(1)}$ forms an orthonormal basis of $(\text{span}\{z_1\})^\perp$ (in W'). Consequently, $b_j^{(1)} = (b_{2,j}^{(1)}, \dots, b_{\ell,j}^{(1)})$ equals the coordinate vector with respect to this basis of the residual $z_j^{(1)} = z_j - r_{1,j}q_1$ for all $j \geq 2$. Linear independence guarantees that these residuals and $b_j^{(1)}$ are nonzero.

If $\ell \geq 2$, then a suitable Householder transform O'_2 reflects the coordinate vector $b_2^{(1)}$ with respect to $v_2^{(1)}, \dots, v_\ell^{(1)}$ into one of the two vectors $\pm\|b_2^{(1)}\|e_1$. The relevant properties of O'_2 carry over to $O_2 = \begin{pmatrix} 1 & \\ & O'_2 \end{pmatrix}$, which drives the next basis change. If $\ell > 2$, then

$$Z = (VO_1O_2)(O_2O_1B) = \left[\begin{array}{c|c|c|c} q_1 & q_2 & v_3^{(2)} & \dots & v_\ell^{(2)} \end{array} \right] \left(\begin{array}{c|c|c|c|c} \overbrace{r_{1,1}} & \overbrace{r_{1,2}} & \overbrace{r_{1,3}} & \dots & \overbrace{r_{1,\ell}} \\ \hline & r_{2,2} & r_{2,3} & \dots & r_{2,\ell} \\ \hline & & b_{3,3}^{(2)} & \dots & b_{3,\ell}^{(2)} \\ \hline & & \vdots & \ddots & \vdots \\ \hline & & b_{\ell,3}^{(2)} & \dots & b_{\ell,\ell}^{(2)} \end{array} \right) \left. \begin{array}{l} \text{with respect} \\ \text{to } q_1 \text{ and } q_2 \\ \\ \text{with respect} \\ \text{to } v_3^{(2)}, \dots, v_\ell^{(2)} \end{array} \right\}$$

wherein the presentation presupposes $\ell > 3$. Proceeding in this fashion verifies the assertion of lemma 4.6 for an arbitrary sequence length $\ell \in \mathbb{N}$.

4.4.3. Recursive processing

This section designs a recursive procedure to evaluate orthogonal projections. The setting amounts to an extended special case of the scenario in section 4.4.1. More specifically, two sequences of real-valued functions $w_{t,j}$, $I_x = \{1, \dots, n\} \times \{1, \dots, m\}$, $m, n \in \mathbb{N}$, and $v_{t,i}$, $(t, i) \in I_{\text{obs}} \neq \emptyset$, on a common set Ω form an orthonormal basis of a Euclidean space $(W, \langle \cdot, \cdot \rangle)$. In particular, the second index set I_{obs} is a finite and nonempty subset of $\mathbb{N} \times \mathbb{N}$. The discussion mostly focuses on the additional elements $x_{t,j}$, $(t, j) \in I_x$, and $y_{t,i}$, $(t, i) \in I_{\text{obs}}$, of W . In fact, the task is to evaluate at ω the orthogonal projections $P_V x_{t,j}$ of $x_{t,j}$ onto $V = \text{span}\{y_{t,i} \mid (t, i) \in I_{\text{obs}}\}$ based on coordinate information and the observations $y_{t,i}(\omega)$, $(t, i) \in I_{\text{obs}}$. The functions $x_{t,j}$ and $y_{t,i}$ are given by

$$\begin{aligned} X_1 &= [x_{1,1} \cdots x_{1,m}] = W_1 L_1^\top, \\ X_t &= [x_{t,1} \cdots x_{t,m}] = X_{t-1} \Theta^\top + \rho W_t, \quad 2 \leq t \leq n, \quad \text{and} \end{aligned} \tag{4.21a}$$

$$Y_t = [y_{t,1} \cdots y_{t,k_t}] = \overbrace{[X_{t-1} B_t^\top \ X_t]}^{Z_t} \overbrace{\begin{bmatrix} -I \\ B_t^\top & J_t^\top \end{bmatrix}}^{A_t^\top} + V_t S_t^\top, \quad t \leq n. \tag{4.21b}$$

The first part <4.21a> defines the quantities $x_{t,j}$. Therein, the linear maps W_t , $t \leq n$, amount to $[w_{t,1} \cdots w_{t,m}]$, the kernel of the $m \times m$ matrix L_1^\top equals $\{0\}$, and $\rho > 0$. These restrictions guarantee linear independence of $x_{t,j}$, $(t, j) \in I_x$. Moreover, $\Theta \in \mathbb{R}^{m \times m}$.

The second part <4.21b> specifies the observables $y_{t,i}$. Here, the first summand equals

$$Z_t A_t^\top = [X_{t-1} B_t^\top \ X_t] \begin{bmatrix} -I \\ B_t^\top & J_t^\top \end{bmatrix} = [X_{t-1} \ X_t] \begin{bmatrix} -B_t^\top \\ B_t^\top & J_t^\top \end{bmatrix}$$

with $J_t^\top \in \mathbb{R}^{m \times k_t'}$, $B_t^\top \in \mathbb{R}^{m \times k_t''}$, $k_t' + k_t'' = k_t > 0$ being the number of observations at t , and $k_t', k_t'' \in \mathbb{N} \cup \{0\}$. If either of k_t' and k_t'' equals zero, then the other coincides with k_t , and A_t^\top consists of only one of the shown block columns. More specifically, the equality $k_t'' = 0$ implies $Z_t = X_t$ and $A_t^\top = J_t^\top$. This case holds for $t = 1$, and hence there is no need to ponder the meaning of X_0 . If $k_t' > 0$, then the columns of J_t^\top form a sequence of linearly independent elements of \mathbb{R}^m . Finally, the second summand of Y_t equals the composition $V_t S_t^\top$ of the linear map $V_t = [v_{t,1} \cdots v_{t,k_t}]$ with the matrix $S_t^\top \in \mathbb{R}^{k_t \times k_t}$.

The subsequent discussion assumes linear independence of the observables $y_{t,i}$, $(t, i) \in I_{\text{obs}} = \cup_{t \leq n} (\{t\} \times \{1, \dots, k_t\})$. The latter amounts to linear independence of the columns of the matrix shown in figure 4.8 as $x_{t,j}$, $v_{t,i}$, $(t, j) \in I_x$, $(t, i) \in I_{\text{obs}}$, are linearly independent. The equalities $\ker S_t^\top = \{0\}$, $t \leq n$, or alternatively $k_t'' = 0$, $t \leq n$, guarantee this condition. Otherwise, if $k_t'' > 0$ and $\ker S_t^\top \neq \{0\}$ for some $t \leq n$, then the requirement of linearly independent observables restricts consecutive matrices J_t , J_{t-1} and $\ker B_t^\top$, $t \geq 2$.

Changing the focus from $x_{t,j}$, $(t, j) \in I_x$, to the functions $z_{t,j}$, $t \leq n$, $j \leq k_t'' + m$, lowers the complexity of the notation when designing a computational strategy to evaluate the mentioned predictions. In fact, the equalities $x_{t,j} = z_{t,j}$, $j \geq k_t'' + 1$, imply that the computation of the rows $(P_V Z_t)(\omega)$ of $P_V Z_t$, $t \leq n$, settles the original prediction task.

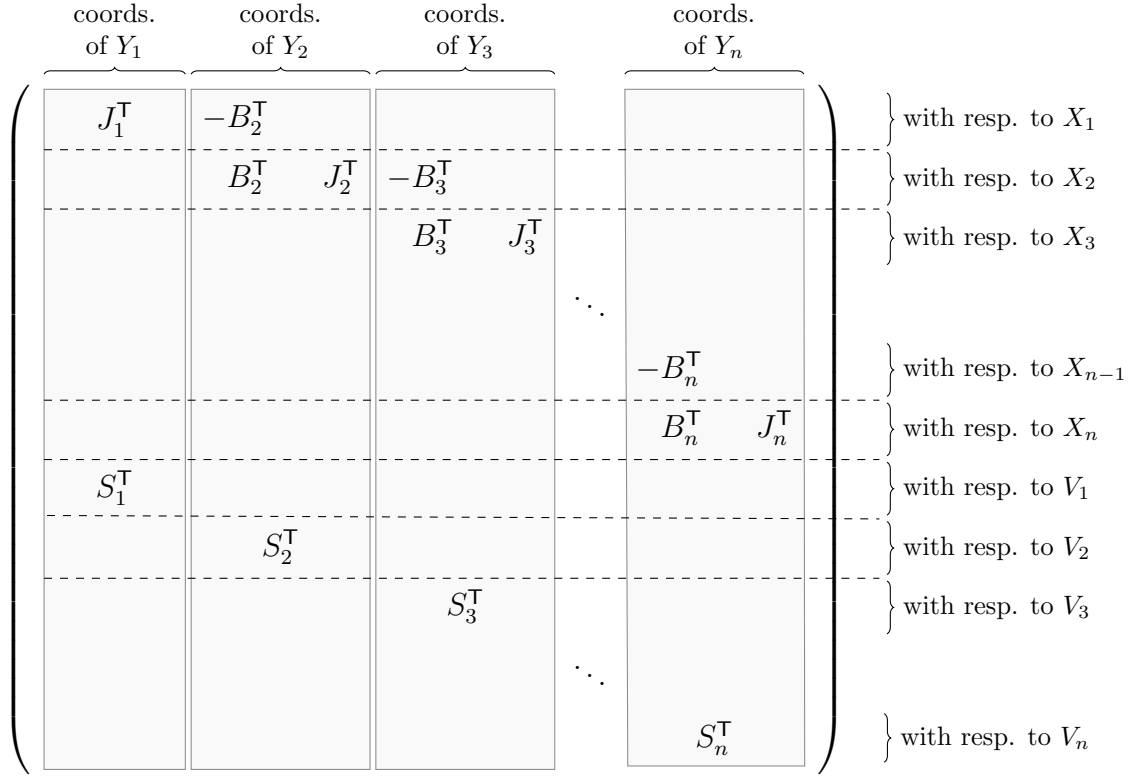


Figure 4.8

The figure shows the matrix of coordinates of the columns of $[Y_1 \cdots Y_n]$ with respect to the columns of $[X_1 \cdots X_n V_1 \cdots V_n]$ as specified in <4.21a> and <4.21b> and for the case $n > 3$. If the equality $k_t'' = 0$ holds—as in case $t = 1$ —or $k_t' = 0$ for some $t \leq n$, then $k_t' = k_t > 0$ or $k_t'' = k_t > 0$ and the block column containing B_t^\top and J_t^\top , respectively, disappears.

In addition, the specification in <4.21a> may be replaced by the equalities

$$Z_1 = W_1 L_1^\top, \quad Z_t = Z_{t-1} \left(\overbrace{B_t^\top \quad \Theta^\top}^{T_t^\top} \right) + W_t \left[\overbrace{\rho I}^{K^\top} \right], \quad 2 \leq t \leq n, \quad <4.22>$$

wherein I denotes the $m \times m$ identity matrix, and $n \geq 2$ is assumed. If the number k_{t-1}'' of columns of B_{t-1}^\top equals zero—as in case $t = 2$, then the first block row of T_t^\top disappears. The same applies to the first block column of T_t^\top and K^\top if $k_t'' = 0$. If the equalities $k_{t-1}'' = 0$ and $k_t'' = 0$ hold simultaneously, then $Z_t = X_t$, $T_t = \Theta$, and $K = \rho I$.

The following presentation considers the case $n > 4$. In this setting, the strategy of section 4.4.1 is applicable. However, adapting the computations to the specification <4.22> and <4.21b> allows to exploit this particular structure. The rearranged equivalent

$$[Y_1 \ Z_1 \ Y_2 \ Z_2 \ \dots \ Y_n \ Z_n] = [W_1 \ V_1 \ W_2 \ V_2 \ \dots \ W_n \ V_n] B \quad <4.23>$$

to <4.16>, wherein the matrix B is as shown in figure 4.9, facilitates this endeavor.

The computational strategy which is developed here proceeds in two stages. The first stage—referred to as *filtering*—exchanges the orthonormal basis consisting of columns of W_t and V_t , $t \leq n$, by another orthonormal basis which contains an orthonormal basis of $V = \text{span}\{y_{t,i} \mid (t,i) \in I_{\text{obs}}\}$ as a subsequence. Successive evaluation at ω of the orthogonal projections of the columns of Z_1 onto $V^{(1)} = \text{img } Y_1$, the orthogonal projections of the columns of Z_2 onto $V^{(2)} = \text{img } [Y_1 \ Y_2]$, and so forth allows a considerable reduction of the computations needed to evaluate the elements of this subsequence. The second step—called *smoothing*—calculates the required images $(P_V x_{t,j})(\omega)$, $(t,j) \in I_x$, based on the output of the first step. Herein, successive calculation of the entries of $(P_V Z_{n-1})(\omega)$, the entries of $(P_V Z_{n-2})(\omega)$, and so forth further allows to avoid redundant calculations.

The first step of the filtering stage replaces the columns of W_1 and V_1 . More specifically, a suitable sequence of k_1 Householder transforms—chosen with respect to the columns of Y_1 , but applied to the relevant rows of all columns of B —yields

$$\begin{array}{c}
 \text{modified by} \\
 \text{Householder} \\
 \text{transforms} \\
 \left. \begin{array}{l} [Y_1 \ Z_1 \ Y_2 \ Z_2 \ \cdots] \\ = [Q_1 \ V'_1 \ W_2 \ V_2 \ \cdots] \end{array} \right\} \left(\begin{array}{c|c|c|c|c}
 R_{1,Y_1} & R_{1,Z_1} & R_{1,Z_1} T_2^\top A_2^\top & R_{1,Z_1} T_2^\top & \cdots \\
 \hline
 & [L_{2,*}^\top] & [L_{2,*}^\top T_2^\top] A_2^\top & [L_{2,*}^\top T_2^\top] & \cdots \\
 & & \begin{array}{c} K^\top \\ S_2^\top \end{array} & \underbrace{\begin{array}{c} L_2^\top \\ K^\top \end{array}}_{L_2^\top} & \cdots \\
 & & & & \ddots
 \end{array} \right) \left. \begin{array}{l} \text{coords. of proj.} \\ \text{onto } V^{(1)} = \text{img } Y_1 \\ \text{coords. of proj.} \\ \text{onto } [V^{(1)}]^\perp \end{array} \right\}
 \end{array}$$

$$\begin{array}{c}
 W'_2 = [V'_1 \ W_2] \\
 V_2 \\
 W_3 \\
 V_3 \\
 \vdots
 \end{array}
 \left| \begin{array}{c}
 \begin{array}{c} P_{(V^{(1)})^\perp} Y_2 \\ P_{(V^{(1)})^\perp} Z_2 \\ P_{(V^{(1)})^\perp} Y_3 \\ P_{(V^{(1)})^\perp} Z_3 \\ \cdots \end{array} \\
 L_2^\top A_2^\top \quad L_2^\top \quad L_2^\top T_3^\top A_3^\top \quad L_2^\top T_3^\top \quad \cdots \\
 S_2^\top \\
 K^\top A_3^\top \quad K^\top \quad \cdots \\
 S_3^\top \\
 \ddots
 \end{array}
 \right.$$

Therein, the use of solely k_1 Householder transforms does not ensure a particular structure of $L_{2,*}^\top$, which is therefore treated as a general $m \times m$ matrix. In contrast, linear independence of $y_{1,1}, \dots, y_{1,k_1}$ implies that R_{1,Y_1} is a $k_1 \times k_1$ upper triangular matrix with nonzero diagonal elements. This structure facilitates the evaluation of the columns of Q_1 , which form an orthonormal basis of $V^{(1)}$, via the equality $Y_1 = Q_1 R_{1,Y_1}$. In particular, the recursive strategy in <4.20> is applicable to the calculation of the entries $q_{1,1}(\omega), \dots, q_{1,k_1}(\omega)$ of $Q_1(\omega)$. Once the row $Q_1(\omega) \in \mathbb{R}^{k_1}$ is available, the images of ω under the columns of $P_{V^{(1)}} Z_1$, $P_{V^{(1)}} Z_2$, and $P_{(V^{(1)})^\perp} Y_2$ may be obtained via

$$P_{V^{(1)}} Z_1 = Q_1 R_{1,Z_1}, \quad P_{V^{(1)}} Z_2 = (P_{V^{(1)}} Z_1) T_2^\top, \quad \text{and} \quad P_{(V^{(1)})^\perp} Y_2 = Y_2 - (P_{V^{(1)}} Z_2) A_2^\top.$$

Herein, linear independence of $y_{1,1}, \dots, y_{1,k_1}, y_{2,1}, \dots, y_{2,k_2}$ ensures that $R_{2,Y_2} \in \mathbb{R}^{k_2 \times k_2}$ is upper triangular with nonzero diagonal elements. Hence, the equality $P_{(V^{(1)})^\perp} Y_2 = Q_2 R_{2,Y_2}$ allows the computation of the row $Q_2(\omega) \in \mathbb{R}^{k_2}$ via <4.20>. The columns of the corresponding linear map Q_2 form an orthonormal basis of the subspace $\text{img } P_{(V^{(1)})^\perp} Y_2$, which equals the orthogonal complement of $V^{(1)}$ in $V^{(2)}$. The images under the relevant projections onto $V^{(2)}$ and $(V^{(2)})^\perp$, respectively, follow from

$$\begin{aligned} P_{V^{(2)}} Z_2 &= P_{V^{(1)}} Z_2 + Q_2 R_{2,Z_2} , \\ P_{V^{(2)}} Z_3 &= (P_{V^{(1)}} Z_2) T_3^\top + (Q_2 R_{2,Z_2}) T_3^\top = (P_{V^{(2)}} Z_2) T_3^\top , \quad \text{and} \\ P_{(V^{(2)})^\perp} Y_3 &= Y_3 - (P_{V^{(2)}} Z_3) A_3^\top . \end{aligned}$$

Finally, the linear map $W'_3 = [V'_2 \ W_3]$ exhibits $3m$ columns and thus $L'_3 \in \mathbb{R}^{3m \times (k'_3 + m)}$.

The next steps of the filtering stage proceed in analogy. Display <4.24> contains a complete description. The notation used therein is in accordance with the previous discussion. In addition, the symbols $y_t, \tilde{y}_{t+1|t}, \hat{z}_{t|j}, t \in \{t-1, t\}$, and q_t refer to the vectors $Y_t(\omega), (P_{(V^{(t)})^\perp} Y_{t+1})(\omega), (P_{V^{(j)}} Z_t)(\omega)$, and $Q_t(\omega)$, respectively. Then,

$$\begin{array}{l|l} 1 & \tilde{y}_{1|0} = y_1 \\ 2 & \hat{z}_{1|0} = 0 \\ 3 & \text{for } t = 1, \dots, n \\ 4 & D_t = \begin{bmatrix} L_t^\top A_t^\top & L_t^\top \\ S_t^\top & \end{bmatrix} \xrightarrow{\text{Householder}} D'_t = \begin{bmatrix} R_{t,Y_t} & R_{t,Z_t} \\ & L_{t+1,*}^\top \end{bmatrix} \\ 5 & \text{solve } R_{t,Y_t}^\top q_t = \tilde{y}_{t|t-1} , & \text{<4.24>} \\ 6 & \hat{z}_{t|t} = \hat{z}_{t|t-1} + R_{t,Z_t}^\top q_t \\ 7 & \text{if } t < n \\ 8 & L_{t+1}^\top = \begin{bmatrix} L_{t+1,*}^\top & T_{t+1}^\top \\ & K^\top \end{bmatrix} \\ 9 & \hat{z}_{t+1|t} = T_{t+1} \hat{z}_{t|t} \\ 10 & \tilde{y}_{t+1|t} = y_{t+1} - A_{t+1} \hat{z}_{t+1|t} \end{array}$$

wherein the number of rows of the remainder $L_{t+1,*}^\top$ equals tm . In fact, the representation of $Z_{t'+1}$ after processing line 8 of <4.24> with $t = t' < n$ has the form

$$\begin{array}{l} \text{orthon. basis of the } (\sum_{t \leq t'} k_t)\text{-dim.} \\ \text{lin. space } \text{img } [Y_1 \ \dots \ Y_{t'} \quad W'_{t'+1}] \\ \left[\underbrace{Q_1 \ Q_2 \ \dots \ Q_{t'}}_{\text{orthon. basis of the } (\sum_{t \leq t'} k_t + t'm + m)\text{-dim.}} \quad \underbrace{[V'_{t'} \ W'_{t'+1}]}_{\text{lin. space } \text{img } [W_1 \ V_1 \ \dots \ W_{t'} \ V_{t'} \ W'_{t'+1}]} \quad \dots \right] \left[\begin{array}{c} \dots \ R_{1,Z_1} T_2^\top \ \dots \ T_{t'+1}^\top \quad \dots \\ \dots \ R_{2,Z_2} T_3^\top \ \dots \ T_{t'+1}^\top \quad \dots \\ \vdots \\ \dots \ R_{t',Z_{t'}} T_{t'+1}^\top \quad \dots \\ \quad \quad \quad L_{t'+1}^\top \end{array} \right] \left. \begin{array}{l} \dots \\ \dots \\ \text{coords.} \\ \text{of } Z_{t'+1} \\ \dots \end{array} \right\} , \end{array}$$

wherein $L_{t'+1} = [T_{t'+1}L_{t'+1,*} \ K]$, and the presentation considers the case $t' > 2$. However, the rank of $Z_{t'+1}$ and therefore the rank of its coordinate block column in the previous display equals at most $k''_{t'+1} + m$. If $(t' + 1)m > \text{rk } P_{(V^{(t')})^\perp} Z_{t'+1} = \text{rk } L_{t'+1}^\top$, then an intermediate auxiliary basis change allows a reduction of the number of rows of $L_{t'+1}^\top$. This additional transformation is also needed to develop the smoothing stage.

The output generated during the filtering stage is comparable to that of the Gram-Schmidt orthogonalization <2.3>. Each iteration—indexed by t —of <4.24> considers a sequence $y_{t,1}, \dots, y_{t,k_t}$ instead of a single y_t as in <2.3>, but also involves a orthogonalization part (line 4 and 10) and a scaling step (line 5). However, the complexity of the orthogonalization part—measured as the number of orthogonalization steps k_t —does not increase systematically with t in <4.24>—or equivalently j in <2.3>. This reduction in computational effort provides the gain from exploding the present structure.

The recursion <4.24> does not yield the coordinates with respect to Q_{t+1}, \dots, Q_n of the columns of $P_V X_t$, $t < n$. These coordinates are however needed in the smoothing stage. A reconsideration of the first step of the filtering stage yields a suitable extension of <4.24> in form of the above mentioned auxiliary basis change. More specifically, the first step of filtering yields the modified representation

$$\begin{bmatrix} Y_1 & Z_1 & Y_2 & Z_2 & \cdots \\ Q_1 & V'_1 & W_2 & V_2 & \cdots \end{bmatrix} = \begin{pmatrix} R_{1,Y_1} & R_{1,Z_1} & R_{1,Z_1} T_2^\top A_2^\top & R_{1,Z_1} T_2^\top & R_{1,Z_1} T_2^\top T_3^\top A_3^\top & \cdots \\ & \begin{bmatrix} L_{2,*}^\top \\ \vdots \end{bmatrix} & L_2^\top A_2^\top & L_2^\top & L_2^\top T_3^\top A_3^\top & \cdots \\ & & S_2^\top & & \vdots & \ddots \end{pmatrix}.$$

Therein, the $2m \times (k''_2 + m)$ matrix L_2^\top exhibits an extended singular value decomposition

$$L_2^\top = \begin{bmatrix} L_{2,*}^\top & T_2^\top \\ & K^\top \end{bmatrix} = \underbrace{\begin{bmatrix} \bar{u}_1 & \cdots & \bar{u}_{\text{rk } L_2^\top} & \bar{u}_{\text{rk } L_2^\top + 1} & \cdots & \bar{u}_{2m} \end{bmatrix}}_{\substack{\text{left singular} \\ \text{vectors of } L_2^\top} \quad \underbrace{\text{extension to} \\ \text{orthon. basis of } \mathbb{R}^{2m}}} \bar{U}_2 \underbrace{\begin{pmatrix} \sigma_1(L_2^\top) & & & & \\ & \ddots & & & \\ & & \sigma_{\text{rk } L_2^\top}(L_2^\top) & & \\ & & & \ddots & \\ & & & & 0 \end{pmatrix}}_{\bar{D}_2} \underbrace{\begin{bmatrix} \bar{v}_2^\top \\ \vdots \end{bmatrix}}_{\substack{\text{right singular} \\ \text{vectors of } L_2^\top}},$$

nonzero
singular
values of L_2^\top

wherein $\text{rk } L_2^\top < 2m$ is assumed for the sake of presentation, but $\text{rk } L_2^\top = 2m$ is possible. In the latter case, the lower block of zeros in the singular value matrix as well as the extension $\bar{u}_{\text{rk } L_2^\top + 1}, \dots, \bar{u}_{2m}$ disappears. In any case, one has (in)equalities $m \leq \text{rk } L_2^\top = \text{rk } P_{(V^{(1)})^\perp} Z_2 \leq \text{rk } Z_2 \leq k''_2 + m$, wherein the first inequality is due to the presence of K^\top . This extended singular value decomposition leads to the equalities

$$W_2' \bar{U}_2 = [V'_1 \ W_2] \bar{U}_2 = [W_2'' \ W_{2,*}''] \quad \text{and} \quad \bar{U}_2^\top \begin{bmatrix} L_{2,*}^\top \\ \vdots \end{bmatrix} = \begin{bmatrix} R_{2,Z_1} \\ L_{2,**}^\top \end{bmatrix}. \quad <4.25>$$

Therein, the leftmost term of the first equality amounts to the composition of two unitary maps, thus, is itself unitary. The righthand side of the previous display shows the coordinates of $P_{(V^{(1)})^\perp} Z_1$ with respect to the columns of $[W_2'' \ W_{2,*}'']$ as $\bar{U}_1 \bar{U}_1^\top = \sum_{i \leq 2m} \bar{u}_i \langle \bar{u}_i, \bullet \rangle \in \mathbb{R}^{m \times m}$ represents the orthogonal projector $P_{\mathbb{R}^{2m}}$, hence, equals the identity matrix I . The partition of this coordinate matrix corresponds to that of the matrix \bar{U}_2 , that is, R_{2,Z_1} is the $\text{rk} L_2^\top \times (k_2'' + m)$ matrix containing the inner products of the original coordinate vectors with $\bar{u}_1, \dots, \bar{u}_{\text{rk} L_2^\top}$. This notation leads to

$$\begin{array}{l}
 [Y_1 \ Z_1 \ Y_2 \ Z_2 \ \cdots] = \\
 [Q_1 \ W_2'' \ V_2 \ \cdots \ W_{2,*}''] \\
 \end{array}
 \left(
 \begin{array}{cccccc}
 R_{1,Y_1} & R_{1,Z_1} & R_{1,Z_1} T_2^\top A_2^\top & R_{1,Z_1} T_2^\top & R_{1,Z_1} T_2^\top T_3^\top A_3^\top & \cdots \\
 & \begin{bmatrix} \bar{L}_2^\top \\ B_2^\top \end{bmatrix} & \begin{bmatrix} \bar{L}_2^\top A_2^\top \\ S_2^\top \end{bmatrix} & \begin{bmatrix} \bar{L}_2^\top \end{bmatrix} & \begin{bmatrix} \bar{L}_2^\top \end{bmatrix} T_3^\top A_3^\top & \cdots \\
 & & & & \vdots & \ddots \\
 & & L_{2,**}^\top & & &
 \end{array}
 \right),$$

wherein $\bar{L}_2^\top = \bar{D}_2 \bar{V}_2^\top = [\bar{u}_1 \ \cdots \ \bar{u}_{\text{rk} L_2^\top}]^\top L_2^\top$, $B_2^\top = \bar{V}_2 \bar{D}_2^{-1} R_{2,Z_1}$, and therefore $\bar{L}_2^\top B_2^\top = R_{2,Z_1}$. Herein, the case of $L_{2,*}^\top$ being an $m \times (k_2'' + m)$ zero matrix is possible and indicates $\text{img} Z_1 \subset \text{img} Y_1$. Then, all entries of the matrices R_{2,Z_1} , $L_{2,**}^\top$, and B_2^\top equal zero.

The coordinate matrix in the previous display—ignoring its final block row—exhibits the same structure as before the auxiliary basis change but with L_2^\top replaced by \bar{L}_2^\top . Consequently, the second filtering step may proceed as above to obtain an orthonormal basis of $\text{img} P_{(V^{(1)})^\perp} Y_2$ by transformation of the columns of the linear map $[W_2'' \ V_2]$.

Inserting an auxiliary basis change of the form <4.25> and some rearrangements of the basis elements at the end of every filtering step yields the matrices $B_2^\top, B_3^\top, \dots, B_n^\top$, which are needed in the smoothing stage. Exchanging line 8 of <4.24> with

$$\begin{array}{l}
 8 \\
 9 \\
 10
 \end{array}
 \left|
 \begin{array}{l}
 L_{t+1}^\top = \begin{bmatrix} L_{t+1,*}^\top T_{t+1}^\top \\ K^\top \end{bmatrix} = [\bar{u}_1 \ \cdots \ \bar{u}_{\text{rk}(L_{t+1}^\top)}]^\top \bar{D}_{t+1} \bar{V}_{t+1}^\top \\
 B_{t+1}^\top = \bar{V}_{t+1} \bar{D}_{t+1}^{-1} [\bar{u}_1 \ \cdots \ \bar{u}_{\text{rk}(L_{t+1}^\top)}]^\top \begin{bmatrix} L_{t+1,*}^\top \\ \end{bmatrix} \\
 L_{t+1}^\top = \bar{D}_{t+1} \bar{V}_{t+1}^\top
 \end{array}
 \right.
 \quad <4.26>$$

provides a suitably extended filtering procedure, wherein the notation is adapted to that of <4.24>. A matrix $L_{t+1,*}^\top$ obtained using the extension <4.26> contains at most $k_t'' + m$ rows; thus, there is no systematic increase of the row count. Moreover, the orthonormal basis $q'_{1,1}, \dots, q'_{1,k_1}, \dots, q'_{n,k_n}$ of V evaluated when using the above extension coincides with the orthonormal basis $q_{1,1}, \dots, q_{n,k_n}$ considered by the unmodified recursions <4.24> up to sign changes. In fact, both bases lead to a representation—in the sense of section 2.2.1—of the columns of $Y = [Y_1 \ \dots \ Y_n]$ in form of an $k \times k$ upper triangular matrix with nonzero diagonal elements. As a consequence, the same equality assertion applies to the corresponding coordinates of the projections onto $V^{(t)}$, $t \leq n$.

The extended filtering stage concludes with a representations of the columns of the linear map $[Y_1 \ Z_1 \ Y_2 \ Z_2 \ \cdots \ Y_n \ Z_n]$ as shown in figure 4.10. Consequently, the smoothing stage may proceed as shown in <4.27>. The latter uses the same notation as <4.24>

$[Y_1 Z_1 \cdots Y_{n-1} Z_{n-1} Y_n Z_n] = [Q_1 \cdots Q_{n-1} Q_n V' W'_2 \cdots W'_{n,**}] R$, wherein R equals

$$\begin{array}{c}
 \left[\begin{array}{c|c|c|c|c|c}
 \text{coords. of } Y_1 & \text{coords. of } Z_1 & \text{coords. of } Y_{n-1} & \text{coords. of } Z_{n-1} & \text{coords. of } Y_n & \text{coords. of } Z_n \\
 \hline
 R_{1,Y_1} & R_{1,Z_1} & R_{1,Z_1} T_2^\top \cdots T_{n-1}^\top A_{n-1}^\top & R_{1,Z_1} T_2^\top \cdots T_{n-1}^\top & R_{1,Z_1} T_2^\top \cdots T_n^\top A_n^\top & R_{1,Z_1} T_2^\top \cdots T_n^\top \\
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
 R_{n-1,Z_{n-1}} B_{n-1}^\top \cdots B_2^\top & R_{n-1,Z_{n-1}} B_{n-1}^\top \cdots B_2^\top & R_{n-1,Y_{n-1}} & R_{n-1,Z_{n-1}} & R_{n-1,Z_{n-1}} T_n^\top A_n^\top & R_{n-1,Z_{n-1}} T_n^\top \\
 R_{n,Z_n} B_n^\top \cdots B_2^\top & R_{n,Z_n} B_n^\top \cdots B_2^\top & R_{n-1,Y_{n-1}} & R_{n,Z_n} B_n^\top & R_{n,Y_n} & R_{n,Z_n} \\
 \hline
 L_{n+1,*} B_n^\top \cdots B_2^\top & L_{n+1,*} B_n^\top \cdots B_2^\top & L_{n+1,*} B_n^\top & L_{n+1,*} B_n^\top & L_{n+1,*} & L_{n+1,*} \\
 L_{2,**} & L_{2,**} & & & & \\
 \vdots & \vdots & & & & \\
 L_{n,**} B_{n-1}^\top \cdots B_2^\top & L_{n,**} B_{n-1}^\top \cdots B_2^\top & & & & \\
 \vdots & \vdots & & & & \\
 & & & & & L_{n,**}
 \end{array} \right]
 \end{array}$$

Figure 4.10

The figure shows a representation—in the sense of section 2.2.1—of the columns of the linear map $[Y_1 Z_1 \cdots Y_n Z_n]$ obtained (implicitly) during the extended filtering stage described in <4.24> and <4.26> and if each filtering step is supplemented with a rearrangement of the basis elements as indicated in the display above <4.26>.

and assumes that the output of the extended filtering stage, in particular, B_2, \dots, B_n , q_1, \dots, q_n , and $\hat{z}_{1|1}, \dots, \hat{z}_{n-1|n-1}$, is available. The vector $\hat{z}_{n|n}$ generated during the final filtering step equals the row $(P_V Z_n)(\omega)$ of $P_V Z_n$, thus, requires no further treatment.

$$\begin{array}{l|l}
 1 & r_{n-1|n} = B_n R_{n,Z_n}^\top q_n \\
 2 & \mathbf{for } t = n - 1, \dots, 1 \\
 3 & \quad \hat{z}_{t|n} = \hat{z}_{t|t} + r_{t|t+1} \\
 4 & \quad \mathbf{if } t > 1 \\
 5 & \quad \quad r_{t-1|t} = B_t (r_{t|t+1} + R_{t,Z_t}^\top q_t)
 \end{array} \tag{4.27}$$

Comments and references

Section 4.1 Stewart and Sun (1990, sec. I.5, exercise 3) provide the definition of θ_{\max} in <4.2b>; the notation is borrowed from Böttcher and Spitkovsky (2010, ex. 3.5). The cosines of the angles $\theta_1, \dots, \theta_\ell$ are sometimes called *canonical correlations* (Anderson, 1958, sec. 12.2, def. 12.2.1). The content of lemma 4.2 can be found in Böttcher and Spitkovsky (2010) and Galántai (2008). Wedin (1983, sec. 1) serves as a role model for the discussion in section 4.1.1 and 4.1.2; his figure 4 closely resembles panel (A) of figure 4.2. The latter investigation implies $P_{V^\perp} P_U v'_i = (\cos^2 \theta_i) v'_i$ and $P_U P_{V^\perp} u_i = (\cos^2 \theta_i) u_i$ (Galántai, 2008, cor. 3), that is, v'_i and u_i provide *eigenvectors* of $P_{V^\perp} P_U$ and $P_U P_{V^\perp}$, respectively, associated with the *eigenvalue* $\cos^2 \theta_i$. In particular, U and V^\perp uniquely determine all of their principal angles—not just θ_{\max} and $\theta_{\min, \neq 0}$. Wedin (1983, app. 1, (A5)) generalizes the upper bound resulting from <4.1> and <4.3b>. Alternatively, Zhu and Knyazev (2013, thm. 4.1, rem. 4.1, tbl. 2 (1,2-entry)) show that if V and U are equal dimensional, then the i -th singular value of $P_V P_{U/V} = P_V - P_{V/U}$ equals $\tan \theta_i$, which implies <4.1>. Their figure 1 illustrates this phenomenon.

Section 4.2 Björck (1996, sec. 5.1.1) considers the *sequential least-squares problem* in <4.7> as a generalization of the classical *constrained least-squares problem*. Comparable convergence assertions to those at the end of section 4.2.2 are given by Lawson and Hanson (1974, ch. 22), Stewart (1997), Ansley and Kohn (1985), and Koopman (1997) amongst others. The related considerations in De Jong (1991) and Eubank (2006, sec. 6.2.2) (implicitly) utilize the $\langle \cdot, \cdot \rangle_*$ -related construct.

The discussion of the case $\text{img}[Y x] \cap U^\perp = \{0\}$ can be extended to $\text{img } Y \cap U^\perp = \{0\}$ by formal manipulation. If the latter holds, then $\ker[Y x] \subset \ker[\hat{Y} \hat{x}]$ is possible. Nonetheless, the bilinear map $\langle \cdot, \cdot \rangle_\circ$ may be defined—in analogy—on $W' \times W'$, however, does not generally provide an inner product. Then, the orthogonality considerations of the main text still apply. In particular, the $\langle \cdot, \cdot \rangle_\circ$ -orthogonal projector equals $P_{V/V^\perp_*} x$, but a zero $\langle \cdot, \cdot \rangle_\circ$ -residual length is possible even if x is not an element of $V = \text{img } Y$.

Section 4.3 Cressie (1991, sec. 3.4) refers to similar predictions—but based on an orthogonal projector instead of an oblique projector—as *universal kriging* predictions; this term is the usual one in spatial statistics (Sherman, 2011, sec. 2.4). Cressie (1991,

sec. 3.4.5) also mentions the geometric perspective taken here. *Best linear unbiased prediction* (BLUP) is another catchword for predictions based on an orthogonal projector (Robinson, 1991). Doran (1992) verifies that—in the setting considered in section 4.4.3—predictions based on orthogonal projections interpolate observed values.

The computation of the (matrix of) coordinates with respect to the columns of $P_{(\text{span}\{1\})^\perp} \bar{Y}$ of the $\langle \cdot, \cdot \rangle_*$ -orthogonal projections of the columns of $P_{(\text{span}\{1\})^\perp} \bar{X}$ onto $\text{img } P_{(\text{span}\{1\})^\perp} \bar{Y}$ provides an alternative to the approach of this text, which ultimately calculates weighted sums of $q_1^*(\omega), \dots, q_{k'}^*(\omega)$ instead of $(P_{(\text{span}\{1\})^\perp} \bar{y}_{t,i})(\omega)$, $(t, i) \in I_{\text{obs}}$.

The ratio $\langle a, \text{var}(\tilde{x})a \rangle / \langle a, \text{var}(\bar{x})a \rangle$ in the upper bound in corollary 4.5 equals

$$\frac{\langle a, \text{var}(\tilde{x})a \rangle}{\langle a, \text{var}(\bar{x})a \rangle - \langle a, \text{var}(\tilde{x})a \rangle} = \frac{\langle a, \text{var}(\tilde{x})a \rangle / \langle a, \text{var}(\bar{x})a \rangle}{1 - \langle a, \text{var}(\tilde{x})a \rangle / \langle a, \text{var}(\bar{x})a \rangle} = \frac{r(a)}{1 - r(a)}.$$

Hence, the first upper bound equals $\sup_{a \in \text{img } A^\top} (r(a) / [1 - r(a)]) = \sup_a r(a) / (1 - \sup_a r(a))$ as $r \mapsto r / (1 - r)$ is monotone increasing on $[0, 1)$. The latter equals $(1 - R_{\min}^2) / R_{\min}^2$, wherein $R_{\min}^2 = \inf_{a \in \text{img } A^\top} (1 - r(a))$ has the interpretation of a minimal (population) *coefficient of determination* across the random variables $\bar{x}^\top a$, $a \in \text{img } A^\top$.

Section 4.4 The discussion of section 4.4.1 resembles Morf and Kailath (1975, sec. IV); their equation (40) contains the equation <4.18>. Golub and Van Loan (2013, sec. 5.1.2, 5.2.2) derive the representation of Householder transforms given in section 4.4.2 as well as the associated triangularization algorithm. The latter considers all of z_1, \dots, z_ℓ from the start whereas the Gram-Schmidt process <2.3> introduces them one after the other. Reorganizing <2.3> to obey the former strategy amounts to orthogonalization of $\tilde{z}_{j+1}^{(j-1)}, \dots, \tilde{z}_\ell^{(j-1)}$ against q_j immediately following its calculation. Björck (1996, algorithm 2.4.3) calls this modification *row oriented modified Gram-Schmidt process*. Panel (A) and (B) of figure 4.7 are akin to Trefethen and Bau (1997, fig. 10.1, 10.2), respectively.

Section 4.4.2 requires linear independence of z_1, \dots, z_ℓ . A generalization as in section 2.2.2 is immediate but not needed here. In fact, an implementation using finite precision arithmetic requires more refined methods of handling linear dependence as identification of zero elements is nontrivial if rounding errors are present. Golub and Van Loan (2013, sec. 5.4.2) considers a popular (rearranging) technique—called *pivoting*.

Eubank (2006, ch. 2–5) develops the (*Kalman*) *filtering* and (*Kalman*) *smoothing* recursions by geometric arguments. Paige (1985) provides a similar presentation. In the usual terminology, lines 10 and 4–6 of <4.24> amount to the *measurement update*; lines 8 and 9 form the *time update*. The former constructs the matrix D_t and then modifies it by pre-multiplication with special matrices. Kailath et al. (2000, ch. 12) discuss such *array algorithms* in-depth including their geometry. Zhang and Li (1996) suggest using singular value decompositions for filtering and smoothing. Extending the algorithm consisting of <4.24>, <4.26>, and <4.27> to yield Gramians of the corresponding residuals leads to so-called *square-root algorithms* (Morf and Kailath, 1975).

Anderson, T. W. (1958). *An introduction to multivariate statistical analysis*. Wiley series in probability and mathematical statistics: Probability and mathematical statistics. New York: Wiley.

- Ansley, C. F. and R. Kohn (1985). Estimation, filtering, and smoothing in state space models with incompletely specified initial conditions. *The Annals of Statistics* 13(4), 1286–1316.
- Björck, Å. (1996). *Numerical methods for least squares problems*. SIAM.
- Böttcher, A. and I. M. Spitkovsky (2010). A gentle guide to the basics of two projections theory. *Linear Algebra and its Applications* 432(6), 1412–1459.
- Cressie, N. A. C. (1991). *Statistics for spatial data*. Wiley series in probability and mathematical statistics: Applied probability and statistics. New York: Wiley.
- De Jong, P. (1991). The diffuse Kalman filter. *The Annals of Statistics* 19(2), 1073–1083.
- Doran, H. E. (1992). Constraining Kalman filter and smoothing estimates to satisfy time-varying restrictions. *The Review of Economics and Statistics* 74(3), 568–572.
- Eubank, R. L. (2006). *A Kalman filter primer*, Volume 186 of *Statistics*. Chapman & Hall/CRC.
- Galántai, A. (2008). Subspaces, angles and pairs of orthogonal projections. *Linear and Multilinear Algebra* 56(3), 227–260.
- Golub, G. H. and C. F. Van Loan (2013). *Matrix computations* (4 ed.). Johns Hopkins studies in the mathematical sciences. Baltimore, Md.: Johns Hopkins Univ. Press.
- Kailath, T., A. H. Sayed, and B. Hassibi (2000). *Linear estimation*, Volume 1. Prentice Hall New Jersey.
- Koopman, S. J. (1997). Exact initial Kalman filtering and smoothing for nonstationary time series models. *Journal of the American Statistical Association* 92(440), 1630–1638.
- Lawson, C. L. and R. J. Hanson (1974). *Solving least squares problems*, Volume 161. SIAM.
- Morf, M. and T. Kailath (1975). Square-root algorithms for least-squares estimation. *Automatic Control, IEEE Transactions on* 20(4), 487–497.
- Paige, C. C. (1985). Covariance matrix representation in linear filtering. *Special Issue of Contemporary Mathematics on Linear Algebra and its Role in Systems Theory* 47, 309–321.
- Robinson, G. K. (1991, 2). That BLUP is a Good Thing: The Estimation of Random Effects. *Statistical Science* 6(1), 15–32.
- Sherman, M. (2011). *Spatial statistics and spatio-temporal data*. Wiley series in probability and statistics. Chichester: Wiley.
- Stewart, G. W. (1997). On the weighting method for least squares problems with linear equality constraints. *BIT Numerical Mathematics* 37(4), 961–967.
- Stewart, G. W. and J. Sun (1990). *Matrix perturbation theory*. Computer science and scientific computing. Boston: Acad. Press.
- Trefethen, L. N. and D. Bau (1997). *Numerical linear algebra*. Philadelphia: SIAM, Soc. for Industrial and Applied Math.
- Wedin, P. Å. (1983). On angles between subspaces of a finite dimensional inner product space. In *Matrix Pencils*, pp. 263–285. Springer.
- Zhang, Y. and X. R. Li (1996). Fixed-interval smoothing algorithm based on singular value decomposition. In *Proceedings of the IEEE International Conf. on Control Appl.*, pp. 916–921. IEEE.
- Zhu, P. and A. V. Knyazev (2013). Angles between subspaces and their tangents. *Journal of Numerical Mathematics* 21(4), 325–340.

Appendix

4.a. Proofs

Proof of proposition 4.3. The inclusions $\text{img } \hat{Y} \subset (\text{img } \tilde{Y})^\perp \subset (\text{img } \tilde{Y}_*)^\perp$ guarantee the final equality in $\ker Y = \ker \hat{Y} \cap \ker \tilde{Y} = \ker \hat{Y} \cap \ker \tilde{Y}_* = \ker Y_*$ and thereby $\dim V = \dim V_*$.

If $x \in V \cap V_*$, then $x = Yc = Y_*c'$ for some $c, c' \in \mathbb{R}^k$ with $c - c' \in \ker \hat{Y}$, $c - P_{\ker \hat{Y}} c' \in \ker \tilde{Y}$, and thereby $\tilde{Y}c'' \in \text{img } \tilde{Y}_*$, wherein $c'' = P_{(\ker \hat{Y})^\perp} c$. In particular, $x = Yc = \hat{Y}c'' + (\tilde{Y}c'' + \tilde{Y}_*c) \in V' = \text{img } \hat{Y}_\times + \text{img } \tilde{Y}_*$. Conversely, $V' \subset V_*$, and if $x = \hat{Y}c + \tilde{Y}c'$ for some $c \in \{\tilde{Y} \in \text{img } \tilde{Y}_*\}$ and

$c' \in \ker \hat{Y}$, then there exists $c'' \in \ker \hat{Y}$ with $\tilde{Y}c = \tilde{Y}_*c''$ and therefore $x = Yc + \tilde{Y}(c' - c'') = Y(c + c' - c'')$. Hence, $\text{img } \tilde{Y}_* = \text{img } \tilde{Y}$ implies $V \cap V_* \supset V, V_*$. Conversely, $V = V_*$ ensures that $\text{img } \hat{Y} \subset V$, which leads to $\text{img } \tilde{Y} \subset V = V_*$, and consequently $\text{img } \tilde{Y} = \text{img } \tilde{Y}_*$.

If c is an element of $\ker P_{(\text{img } \hat{Y}_\times)^\perp} \hat{Y}$, then $\hat{Y}c' \in \text{img } \hat{Y}_\times$, wherein c' denotes the orthogonal projection of c onto $(\ker \hat{Y})^\perp$. As a consequence, there exists $c'' \in \{\tilde{Y} \in \text{img } \tilde{Y}_*\}$ with $c' - c'' \in \ker \hat{Y}$. Furthermore, one may assume $c' = c''$ as $\tilde{Y}P_{(\ker \hat{Y})^\perp}c'' = \tilde{Y}c' - \tilde{Y}_*c'' \in \text{img } \tilde{Y}_*$. Thus, $\tilde{Y}c = \tilde{Y}P_{\ker \hat{Y}}c + \tilde{Y}P_{(\ker \hat{Y})^\perp}c = \tilde{Y}_*c + \tilde{Y}c' \in \text{img } \tilde{Y}_*$, that is, $P_{(\text{img } Y_*)^\perp}\tilde{Y}c = 0$. Conversely, if the latter holds, then $c \in \{\tilde{Y} \in \text{img } \tilde{Y}_*\}$, thus, $\hat{Y}c \in \text{img } \hat{Y}_\times$.

In particular, if $V \neq V_*$, then K is nontrivial and $\theta_{\min, \neq 0}(V^\perp, V_*^\perp) = \theta_{\min, \neq 0}(V, V_*)$ as well as $\theta_{\max}(V^\perp, V_*^\perp) = \theta_{\max}(V, V_*)$. Then, elements v of $V \cap (V \cap V_*)^\perp$ have the form $P_{(\text{img } \hat{Y}_\times)^\perp}\hat{Y}c + P_{(\text{img } \tilde{Y}_*)^\perp}\tilde{Y}c$ with $c \in K$. In fact, $(\text{img } \hat{Y}_\times)^\perp \ni P_{(\text{img } \hat{Y}_\times)^\perp}\hat{Y}c = \hat{Y}c - P_{\text{img } \hat{Y}_\times}\hat{Y}c \in (\text{img } \tilde{Y}_*)^\perp$ and likewise for $P_{(\text{img } \tilde{Y}_*)^\perp}\tilde{Y}c$. Moreover, $P_{(\text{img } \hat{Y}_\times)^\perp}\hat{Y}c + P_{(\text{img } \tilde{Y}_*)^\perp}\tilde{Y}c = Yc - (P_{\text{img } \hat{Y}_\times}\hat{Y}c + P_{\text{img } \tilde{Y}_*}\tilde{Y}c) \in \text{img } Y + (V \cap V_*) \subset \text{img } Y$. Conversely, if $Yc \in (\text{img } \hat{Y}_\times + \text{img } \tilde{Y}_*)^\perp$, then $\hat{Y}c \in (\text{img } \hat{Y}_\times)^\perp$ and $\tilde{Y}c = (\text{img } \tilde{Y}_*)^\perp$. Hence, if v is a nonzero element of $V \cap (V \cap V_*)^\perp$, then

$$\phi_{V_*}^2 \left(\frac{v}{\|v\|} \right) = \frac{\|(P_{\text{img } \hat{Y}} + P_{\text{img } \tilde{Y}_*})(P_{(\text{img } \hat{Y}_\times)^\perp}\hat{Y}c + P_{(\text{img } \tilde{Y}_*)^\perp}\tilde{Y}c)\|^2}{\|P_{(\text{img } \hat{Y}_\times)^\perp}\hat{Y}c\|^2 + \|P_{(\text{img } \tilde{Y}_*)^\perp}\tilde{Y}c\|^2} = \frac{1}{1 + \frac{\|P_{(\text{img } \tilde{Y}_*)^\perp}\tilde{Y}c\|^2}{\|P_{(\text{img } \hat{Y}_\times)^\perp}\hat{Y}c\|^2}}.$$

The latter implies the representation of $\tan \theta_{\min, \neq 0}$ and $\tan \theta_{\max}$ as $\phi_{V_*}(v') \in [0, 1)$ for all unit length v' , $\sqrt{\bullet}$ is a monotone function on $[0, 1)$, $\cos^2 \theta = 1/(1 + \tan^2 \theta)$, and $\tan \theta \geq 0$ for $\theta \in [0, \pi/2)$. In particular, $\theta_{\max} < \pi/2$ ensures that V and V_*^\perp are complementary. ■

Proof of the final equality in <4.9>. If $x \in V$, then $W' = \text{img } [Yx] = \text{img } Y = V$, and both projectors equal the identity map irrespective of the value of $\delta \in (0, 1]$. In the special case $\text{img } Y \subset U^\perp$, one has $Y_* = Y = Y_\delta$, and both projectors coincide for all $\delta \in (0, 1]$. Consequently, this special case requires no further consideration. If $x \notin V$, $\text{img } Y \not\subset U^\perp$, and $\delta \in (0, 1]$, then norm equivalence (on W') ensures the existence of $c, C > 0$ with $c\|z\|_* \leq \|z\| \leq C\|z\|_*$ for all $z \in W'$. Therefore, every unit $\|\bullet\|$ -length element $z' \in W'$ satisfies the inequality

$$\begin{aligned} \|(P_{V/V^{\perp*}} - P_{V/V^{\perp\delta}})z'\| &\leq C \sup_{\substack{z \in W' \\ z \neq 0}} \|(P_{V/V^{\perp*}} - P_{V/V^{\perp\delta}}) \frac{z}{\|z\|}\|_* \\ &\leq \frac{C}{c} \sup_{\substack{z \in W' \\ \|z\|_* = 1}} \|(P_{V/V^{\perp*}} - P_{V/V^{\perp\delta}})z\|_* \\ &\leq \frac{C}{c} \sup_{\substack{z \in W' \\ \|z\|_* = 1}} \left\| \sum_{i \leq \dim W'} \langle w_i, z \rangle (P_{V/V^{\perp*}} - P_{V/V^{\perp\delta}})w_i \right\|_* \\ &\leq \frac{C}{c} \underbrace{\left[\sup_{\substack{z \in W' \\ \|z\|_* = 1}} |\langle z, w_1 \rangle| \right]}_{\text{dual norm of } \|\bullet\|_* \text{ at } w_1} \|(P_{V/V^{\perp*}} - P_{V/V^{\perp\delta}})w_1\|_*, \end{aligned}$$

wherein $w_1, \dots, w_{\dim W'}$ denotes an $\langle \bullet, \bullet \rangle_*$ -orthonormal basis of W' such that $V^{\perp*} = \text{span}\{w_1\}$ and thereby $w_2, \dots, w_{\dim W'}$ form an orthonormal basis of V . By definition of W' , the basis element w_1 exhibits a representation $w_1 = c_1x + Yc_2$ with $c_1 \neq 0$ and $c_2 \in \mathbb{R}^k$. Therefore, it suffices to consider $\frac{1}{|c_1|} \|(P_{V/V^{\perp*}} - P_{V/V^{\perp\delta}})w_1\|_* = \|(P_{V/V^{\perp*}} - P_{V/V^{\perp\delta}})x\|_*$. To this end, let

$\tilde{Y}_\times = \tilde{Y}P_{(\ker \hat{Y})^\perp}$, $P_{V/V^\perp_*}x = Yb_*$, and $P_{V/V^\perp_\delta}x = Yb_\delta$, wherein the specific choice of b_* , $b_\delta \in \mathbb{R}^k$ —in case $\ker Y \neq \{0\}$ —is immaterial. Then, the inequality $\|(P_{V/V^\perp_*} - P_{V/V^\perp_\delta})x\|_* = \|Y_*(b_* - b_\delta)\| \leq \|\hat{Y}(b_* - b_\delta)\| + \|\tilde{Y}_*(b_* - b_\delta)\|$ reveals that separate consideration of the latter two summands suffices. The two vectors b_* and b_δ are minimizers of

$$b' \mapsto Q_*(b') = \frac{1}{\delta^2}\|\hat{x} - \hat{Y}b'\|^2 + \|\tilde{x} - \tilde{Y}_\times b - \tilde{Y}_*b'\|^2 \quad \text{and} \quad \langle \text{A4.1} \rangle$$

$$b' \mapsto Q_\delta(b') = \|x - Yb'\|_\delta^2 = \frac{1}{\delta^2}\|\hat{x} - \hat{Y}b'\|^2 + \|\tilde{x} - \tilde{Y}_\times b' - \tilde{Y}_*b'\|^2,$$

respectively, wherein $b \in (\ker \hat{Y})^\perp$ minimizes $b' \mapsto \|\hat{x} - \hat{Y}b'\|^2$, that is, $P_{\text{img } \hat{Y}}\hat{x} = P_{\text{img } \hat{Y}}x = \hat{Y}b$. The criterion $\langle \text{A4.1} \rangle$ differs from $b' \mapsto \|x - Yb'\|_*^2$; however, the alternative scaling of the first summand does not affect the set of minimizers as $(\ker \hat{Y})^\perp \subset \ker \tilde{Y}_*$. In fact, one has $\langle \hat{Y}, \hat{x} - \hat{Y}b_* \rangle = 0$, that is, $b_* - b \in \ker \hat{Y}$, and thereby $\|\hat{x} - \hat{Y}b_*\| \leq \|\hat{x} - \hat{Y}b_\delta\|$ as well as $Q_*(b_*) = Q_\delta(b_*) \geq Q_\delta(b_\delta)$. As a consequence, $\|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_*\| \geq \|\tilde{x} - \tilde{Y}_\times b_\delta - \tilde{Y}_*b_\delta\|$. Moreover, the above characterization of b_* leads to $\langle \tilde{Y}_*, (\tilde{x} - \tilde{Y}_\times b_*) - \tilde{Y}_*b_* \rangle = 0$. In addition, the inequality $\delta^2 Q_\delta(b_\delta) \leq \delta^2 Q_\delta(b_*)$ guarantees the inequalities

$$0 \leq \|\hat{Y}(b_* - b_\delta)\|^2 = \|\hat{x} - \hat{Y}b_\delta\|^2 - \|\hat{x} - \hat{Y}b_*\|^2 \leq \delta^2(\|\tilde{x} - \tilde{Y}b_*\|^2 - \|\tilde{x} - \tilde{Y}b_\delta\|^2) \leq \delta^2\|\tilde{x} - \tilde{Y}b_*\|^2.$$

The latter inequality handles the first of the two above mentioned summands and leads to

$$\begin{aligned} \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_*\| &\geq \|\tilde{x} - \tilde{Y}_\times b_\delta - \tilde{Y}_*b_\delta\| \geq \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_\delta\| - \|\tilde{Y}\|_{\text{op}}\|P_{(\ker \hat{Y})^\perp}(b_* - b_\delta)\| \\ &\geq \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_\delta\| - \frac{\|\tilde{Y}\|_{\text{op}}}{\sigma_{\min, \neq 0}(\hat{Y})}\|\hat{Y}(b_* - b_\delta)\| \geq \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_\delta\| - \delta C, \end{aligned}$$

wherein $C = \|\tilde{Y}\|_{\text{op}}\|\tilde{x} - \tilde{Y}b_*\|/\sigma_{\min, \neq 0}(\hat{Y})$ does not depend on $\delta \in (0, 1]$. The same applies to the interval $I = [0, \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_*\| + C]$, which—by virtue of the previous display—contains $\|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_\delta\|$ for all $\delta \in (0, 1]$. The map \bullet^2 is monotone increasing and Lipschitz continuous on the interval I with Lipschitz constant $L \leq 2(\|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_*\| + C)$. Thus,

$$\begin{aligned} \|\tilde{Y}_*(b_* - b_\delta)\|^2 + \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_*\|^2 &= \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_\delta\|^2 \\ &\leq (\|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_*\| + C\delta)^2 \leq \|\tilde{x} - \tilde{Y}_\times b_* - \tilde{Y}_*b_*\|^2 + LC\delta, \end{aligned}$$

wherein the first equality is due to the orthogonality condition $\langle \tilde{Y}_*, (\tilde{x} - \tilde{Y}_\times b_*) - \tilde{Y}_*b_* \rangle = 0$. In summary, one has $\|\tilde{Y}(b_* - b_\delta)\| \leq \sqrt{\delta}\sqrt{LC}$. \blacksquare

Proof of lemma 4.4. The equality $\ker[1 \hat{X}] = \{0\}$ is tantamount to $\ker P_{(\text{span}\{1\})^\perp} \hat{X} = \{0\}$. If the condition $\beta_b \notin B_b$ holds, and the matrix $C \in \mathbb{R}^{n \times m}$ is such that the linear combination $\sum_{(t,j) \in I_x} c_{t,j} P_{(\text{span}\{1\})^\perp} \hat{x}_{t,j}$ equals zero, then

$$0 = \sum_{(t,j) \in I_x} c_{t,j} P_{(\text{span}\{1\})^\perp} (\alpha_{t,j} + Z_{2,t,j} \beta_b) = P_{(\text{span}\{1\})^\perp} \left[\sum_{(t,j) \in I_x} c_{t,j} Z_{b,t,j} \right] \beta_b.$$

The requirement $\text{span}\{1\} \cap U_z = \{0\}$ implies that the equality between the leftmost and rightmost term in the previous display is equivalent to $[\sum_{(t,j) \in I_x} c_{t,j} Z_{b,t,j}] \beta_b = 0$. If $C \neq 0$, then the

latter equality continues to hold after division by $\|C\| > 0$. That is, $\beta_b \in \ker[\sum_{(t,j) \in I_x} c'_{t,j} Z_{b,t,j}]$ for some unit length $C' = \frac{1}{\|C\|}C \in \mathbb{R}^{n \times m}$, which contradicts the condition $\beta_b \notin B_b$. Conversely, $\beta_b \in B_b$ implies the existence of a unit length $C \in \mathbb{R}^{n \times m}$ such that $-(\sum_{(t,j) \in I_x} \alpha_{t,j} c_{t,j})1 + \sum_{(t,j) \in I_x} c_{t,j} \hat{x}_{t,j} = 0$, which reveals the linear dependence of $1, \hat{x}_{t,j}, (t,j) \in I_x$. ■

Proof of corollary 4.5. The inequality $k > 0$ and $a_{t,i} \neq 0$ implies $\text{img } A^\top \neq \{0\}$. The equality $\{0\} = \ker \text{var}(\hat{x}) = \langle P_{(\text{span}\{1\})^\perp} \hat{X}, P_{(\text{span}\{1\})^\perp} \hat{X} \rangle$ follows from $\ker[1 \hat{X}] = \ker[1 \bar{X}] = \{0\}$, wherein the first equality is due to $\beta_b \notin B_b$. The remainder of this proof uses the notation of proposition 4.3. The projection is onto $V = \text{img } Y'$ with $[1 \bar{Y}] = Y'$. The inclusion $1 \in U = U_{1,z} + U_{\bar{v}}$ implies that $\hat{Y}' = P_U [1 \bar{Y}] = [1 \hat{Y}]$. The equality $\ker [Y' \bar{X}] = \ker [\hat{Y}' \hat{X}]$ leads to $\ker \hat{Y}' = \ker Y' = \ker \hat{Y}' \cap \ker \tilde{Y}'$ with $\tilde{Y}' = [0 \tilde{Y}]$ and $\tilde{Y} = P_{U^\perp} \bar{Y}$. Thereby, $\{(c_1, c_2) \in \mathbb{R}^{1+k} \mid c_1 \in \mathbb{R}, c_2 \in \ker \tilde{Y}\} = \ker \tilde{Y}' \supset \ker \hat{Y}'$. Consequently, the image of $\hat{Y}'_* = \tilde{Y}' P_{\ker \hat{Y}'}$ equals $\{0\}$. Furthermore, $\{\hat{Y}' \in \text{img } \hat{Y}'_*\} = \ker \tilde{Y}'$, that is, $\hat{Y}'_\times = \hat{Y}' P_{\{\hat{Y}' \in \text{img } \hat{Y}'_*\}} = \hat{Y}' P_{\ker \tilde{Y}'}$. The equality $U^\perp = (U_{1,z} + U_{\bar{v}})^\perp = U_{\bar{v}}$ leads to $\tilde{X} = P_{U^\perp} \bar{X} = [V'_n \dots V'_1]$, $V'_t = [v_{t,1} \dots v_{t,m}]$, $1 \leq t \leq n$, and $\tilde{Y}' = P_{U^\perp} Y' = [0 \tilde{X} A^\top]$. In particular, linear independence of $v_{t,j}, (t,j) \in I_x$, ensures that $\ker P_{(\text{img } \tilde{Y}'_*)^\perp} \tilde{Y}' = \ker \tilde{Y}' = \mathbb{R} \times \ker A^\top = (\{0\} \times \text{img } A)^\perp$, that is, $K = \{0\} \times \text{img } A$, and $\text{img } \tilde{Y}' \neq \{0\} = \text{img } \tilde{Y}'_*$, which implies $V \neq V_*$. The equality $\hat{Y}'_\times = \hat{Y}' P_{\ker \hat{Y}'} = [1 (\hat{X} A^\top + \bar{V} S^\top) P_{\ker A^\top}] = [1 \bar{V} S^\top P_{\ker A^\top}]$ implies $\text{span}\{1\} \subset \text{img } \hat{Y}'_\times \subset \text{img}[1 \bar{V}] = \text{span}\{1\} + U_{\bar{v}}$. Hence,

$$\begin{aligned} \|P_{(\text{span}\{1\})^\perp} \hat{X} A^\top c\| &= \|P_{(\text{span}\{1\})^\perp \cap U_{\bar{v}}^\perp} \hat{X} A^\top c\| = \|P_{(\text{span}\{1\})^\perp \cap U_{\bar{v}}^\perp} \hat{Y}' \begin{pmatrix} 0 \\ c \end{pmatrix}\| \\ &\leq \|P_{(\text{img } \hat{Y}'_\times)^\perp} \hat{Y}' \begin{pmatrix} 0 \\ c \end{pmatrix}\| \leq \|P_{(\text{span}\{1\})^\perp} \hat{Y}' \begin{pmatrix} 0 \\ c \end{pmatrix}\| = \|P_{(\text{span}\{1\})^\perp} \hat{X} A^\top c + \bar{V} S^\top c\| \end{aligned}$$

holds for every $c \in \mathbb{R}^k$ as $P_{(\text{span}\{1\})^\perp} \hat{X} A^\top c = \hat{X} A^\top c - P_{\text{span}\{1\}} \hat{X} A^\top c - P_{U_{\bar{v}}} \hat{X} A^\top c = \hat{X} A^\top c - P_{\text{span}\{1\} + U_{\bar{v}}} \hat{X} A^\top c = P_{(\text{span}\{1\})^\perp \cap U_{\bar{v}}^\perp} \hat{X} A^\top c$ as well as $P_{(\text{span}\{1\})^\perp \cap U_{\bar{v}}^\perp} \hat{Y}' \begin{pmatrix} 0 \\ c \end{pmatrix} = \hat{X} A^\top c + \bar{V} S^\top c - P_{\text{span}\{1\} + U_{\bar{v}}} \hat{X} A^\top c - P_{\text{span}\{1\} + U_{\bar{v}}} \bar{V} S^\top c = P_{(\text{span}\{1\})^\perp \cap U_{\bar{v}}^\perp} \hat{X} A^\top c$. The equality $\ker \text{var}(\hat{x}) = \{0\}$ implies that the kernel $\ker P_{(\text{span}\{1\})^\perp} \hat{X} A^\top$ coincides with $\ker A^\top = (\text{img } A)^\perp$. Moreover, the inclusion $\tilde{x}_{t,j} \in (\text{span}\{1\})^\perp \supset U^\perp$ holds. Consequently, $c \in \text{img } A$ satisfies the inequalities

$$\begin{aligned} \left(\frac{\langle A^\top c, \text{var}(\tilde{x}) A^\top c \rangle}{\langle A^\top c, \text{var}(\hat{x}) A^\top c \rangle} \right)^{1/2} &= \frac{\|P_{(\text{span}\{1\})^\perp} \tilde{X} A^\top c\|}{\|P_{(\text{span}\{1\})^\perp} \hat{X} A^\top c\|} \geq \frac{\|P_{(\text{img } \tilde{Y}'_*)^\perp} \tilde{Y}' \begin{pmatrix} 0 \\ c \end{pmatrix}\|}{\|P_{(\text{img } \hat{Y}'_\times)^\perp} \hat{Y}' \begin{pmatrix} 0 \\ c \end{pmatrix}\|} \\ &\geq \frac{\|P_{(\text{span}\{1\})^\perp} \tilde{X} A^\top c\|}{\|P_{(\text{span}\{1\})^\perp} \hat{X} A^\top c + \bar{V} S^\top c\|} = \left(\frac{\langle A^\top c, \text{var}(\tilde{x}) A^\top c \rangle}{\langle A^\top c, [\text{var}(\hat{x}) + S^2] A^\top c \rangle} \right)^{1/2} \quad \langle \text{A4.2} \rangle \end{aligned}$$

If $\ker A^\top \subset \ker S = \ker S^\top$, then $\text{img } \bar{V} S^\top P_{\ker A^\top} = \{0\}$ and therefore $\text{img } \hat{Y}'_\times = \text{span}\{1\}$. Consequently, the final inequality in $\langle \text{A4.2} \rangle$ becomes an equality. ■

Global bibliography

- Anderson, T. W. (1958). *An introduction to multivariate statistical analysis*. Wiley series in probability and mathematical statistics: Probability and mathematical statistics. New York: Wiley.
- Ansley, C. F. and R. Kohn (1985). Estimation, filtering, and smoothing in state space models with incompletely specified initial conditions. *The Annals of Statistics* 13(4), 1286–1316.
- Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American Mathematical Society* 68(3), 337–404.
- Bach, F., R. Jenatton, J. Mairal, and G. Obozinski (2011). Convex optimization with sparsity-inducing norms. In S. Sra, S. Nowozin, and S. J. Wright (Eds.), *Optimization for Machine Learning* (1 ed.), Neural Information Processing, pp. 19–53. MIT Press.
- Bapat, R. B. (2014). *Graphs and Matrices* (2 ed.). Universitext. London: Springer.
- Basu, S. (2014). *Modeling and Estimation of High-dimensional Vector Autoregressions*. Ph. D. thesis, The University of Michigan.
- Basu, S. and G. Michailidis (2015). Regularized estimation in sparse high-dimensional time series models. *The Annals of Statistics* 43(4), 1535–1567.
- Bertsekas, D. P. (1999). *Nonlinear programming* (2 ed.). Athena Scientific.
- Bertsekas, D. P. (2014). *Supplementary Chapter 6 on Convex Optimization Algorithms*, Chapter 6, pp. 249–460. Athena Scientific.
- Bickel, P. J. and K. A. Doksum (2015). *Mathematical Statistics: Basic Ideas and Selected Topics, volume I*, Volume 117. CRC Press.
- Björck, Å. (1996). *Numerical methods for least squares problems*. SIAM.
- Blien, U., H. Gartner, H. Stüber, and K. Wolf (2009). Regional price levels and the agglomeration wage differential in western Germany. *The Annals of Regional Science* 43(1), 71–88.
- Borwein, J. M. and A. S. Lewis (2010). *Convex analysis and nonlinear optimization: theory and examples* (2 ed.), Volume 3. Springer.
- Böttcher, A. and I. M. Spitkovsky (2010). A gentle guide to the basics of two projections theory. *Linear Algebra and its Applications* 432(6), 1412–1459.
- Boyd, S. P. and L. Vandenberghe (2004). *Convex optimization* (Reprint. with corrections ed.). Cambridge [u.a.]: Cambridge University Press.
- Bunea, F., Y. She, and M. H. Wegkamp (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *The Annals of Statistics* 39(2), 1282–1309.
- Chandrasekaran, V., B. Recht, P. A. Parrilo, and A. S. Willsky (2012). The convex geometry of linear inverse problems. *Foundations of Computational Mathematics* 12(6), 805–849.
- Cressie, N. A. C. (1991). *Statistics for spatial data*. Wiley series in probability and mathematical statistics: Applied probability and statistics. New York: Wiley.
- De Jong, P. (1991). The diffuse Kalman filter. *The Annals of Statistics* 19(2), 1073–1083.
- Doran, H. E. (1992). Constraining Kalman filter and smoothing estimates to satisfy time-varying restrictions. *The Review of Economics and Statistics* 74(3), 568–572.
- Doz, C., D. Giannone, and L. Reichlin (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics* 164(1), 188–205.
- Drusvyatskiy, D., S. A. Vavasis, and H. Wolkowicz (2015). Extreme point inequalities and geometry of the rank sparsity ball. *Mathematical Programming* 152(1-2), 521–544.
- Egner, U. (2003). Umstellung des Verbraucherpreisindex auf Basis 2000. *Wirtschaft und Statistik* 5(2003), 423–432.

- Egner, U. (2013). Verbraucherpreisstatistik auf neuer Basis 2010. *Wirtschaft und Statistik* 5, 329–344.
- Elbel, G. (1995). Zur Neuberechnung des Preisindex für die Lebenshaltung auf Basis 1991. *Wirtschaft und Statistik* 11(1995), 801–809.
- Elbel, G. (1999). Die Berechnung der Wägungsschemata für die Preisindizes für die Lebenshaltung. *Wirtschaft und Statistik* 3, 171–178.
- Elbel, G. and U. Egner (2008). Verbraucherpreisstatistik auf neuer Basis 2005. *Wirtschaft und Statistik* 4(2008), 339–50.
- Eubank, R. L. (2006). *A Kalman filter primer*, Volume 186 of *Statistics*. Chapman & Hall/CRC.
- Fazel, M. (2002). *Matrix rank minimization with applications*. Ph. D. thesis, Department of Electrical Engineering, Stanford University.
- Fazel, M., H. Hindi, and S. P. Boyd (2001). A rank minimization heuristic with application to minimum order system approximation. In *American Control Conference 2001, Proceedings of the*, Volume 6, pp. 4734–4739. IEEE.
- Galántai, A. (2008). Subspaces, angles and pairs of orthogonal projections. *Linear and Multilinear Algebra* 56(3), 227–260.
- Golub, G. H. and C. F. Van Loan (2013). *Matrix computations* (4 ed.). Johns Hopkins studies in the mathematical sciences. Baltimore, Md.: Johns Hopkins Univ. Press.
- Grigorieff, R. D. (1991). A note on von Neumanns trace inequality. *Math. Nachr* 151, 327–328.
- Halmos, P. R. (1974). *Finite-dimensional vector spaces* (Repr. of the 2. ed. ed.). New York: Springer.
- Hill, R. J. (2016). A least squares approach to imposing within-region fixity in the international comparisons program. *Journal of Econometrics* 191, 407–413.
- Hiriart-Urruty, J.-B. and C. Lemaréchal (1996). *Convex analysis and minimization algorithms II: Advanced theory and bundle methods*, Volume 306 of *Grundlagen der mathematischen Wissenschaften*. Springer, Berlin.
- Ji, S. and J. Ye (2009). An accelerated gradient method for trace norm minimization. In *Proc. of the 26th annual international conference on machine learning*, pp. 457–464. ACM.
- Kailath, T., A. H. Sayed, and B. Hassibi (2000). *Linear estimation*, Volume 1. Prentice Hall New Jersey.
- Kawka, R. (2010). Regionale Preisunterschiede in den alten und neuen Ländern. *ifo Dresden berichtet* 17(02), 5–16.
- Koltchinskii, V. (2011). *Oracle inequalities in empirical risk minimization and sparse recovery problems*, Volume 2033 of *Lecture notes in mathematics; 2033: Ecole d’Eté de Probabilités de Saint-Flour*. Berlin: Springer.
- Koltchinskii, V. and P. Rangel (2013). Low rank estimation of smooth kernels on graphs. *The Annals of Statistics* 41(2), 604–640.
- Koopman, S. J. (1997). Exact initial Kalman filtering and smoothing for nonstationary time series models. *Journal of the American Statistical Association* 92(440), 1630–1638.
- Kosfeld, R., H.-F. Eckey, and J. Lauridsen (2008). Disparities in prices and income across German NUTS 3 regions. *Applied Economics Quarterly* 54(2), 123–141.
- Lawson, C. L. and R. J. Hanson (1974). *Solving least squares problems*, Volume 161. SIAM.
- Lütkepohl, H. (2007). *New introduction to multiple time series analysis* (1, corr. 2. print. ed.). Berlin: Springer.

- Mishra, B., G. Meyer, F. Bach, and R. Sepulchre (2013). Low-rank optimization with trace norm penalty. *SIAM Journal on Optimization* 23(4), 2124–2149.
- Morf, M. and T. Kailath (1975). Square-root algorithms for least-squares estimation. *Automatic Control, IEEE Transactions on* 20(4), 487–497.
- Negahban, S., P. Ravikumar, M. J. Wainwright, and B. Yu (2012). A unified framework for the analysis of regularized M -estimators. *Statistical Science* 27(4), 538–557.
- Negahban, S. and M. J. Wainwright (2011a). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics* 39(2), 1069–1097.
- Negahban, S. and M. J. Wainwright (2011b). Supplementary Material for: Estimation of (near) low-rank matrices with noise and high-dimensional scaling. available online.
- Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*, Volume 87 of *Applied Optimization*. Springer Science & Business Media, New York.
- Paige, C. C. (1985). Covariance matrix representation in linear filtering. *Special Issue of Contemporary Mathematics on Linear Algebra and its Role in Systems Theory* 47, 309–321.
- Parzen, E. (1961). An approach to Time Series Analysis. *The Annals of Mathematical Statistics* 32(4), pp. 951–989.
- Pollard, D. (2002). *A user's guide to measure theoretic probability*, Volume 8 of *Cambridge series in statistical and probabilistic mathematics*. Cambridge: Cambridge University Press.
- Pong, T. K., P. Tseng, S. Ji, and J. Ye (2010). Trace norm regularization: Reformulations, algorithms, and multi-task learning. *SIAM Journal on Optimization* 20(6), 3465–3489.
- Rao, P. D. S., A. Rambaldi, and H. Doran (2010). Extrapolation of purchasing power parities using multiple benchmarks and auxiliary information: A new approach. *Review of Income and Wealth* 56, 59–98.
- Recht, B., M. Fazel, and P. A. Parrilo (2010). Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review* 52(3), 471–501.
- Robinson, G. K. (1991, 2). That BLUP is a Good Thing: The Estimation of Random Effects. *Statistical Science* 6(1), 15–32.
- Rockafellar, R. T. (1970). *Convex analysis*, Volume 28 of *Princeton mathematical series*. Princeton, NJ: Princeton University Press.
- Roos, M. (2006). Regional price levels in Germany. *Applied Economics* 38(13), 1553–1566.
- Rudelson, M. and R. Vershynin (2013). Hanson-Wright inequality and sub-gaussian concentration. *Electron. Commun. Probab* 18(82), 1–9.
- Schölkopf, B. and A. J. Smola (2002). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. Adaptive computation and machine learning. Cambridge, Mass.: MIT Press.
- Sherman, M. (2011). *Spatial statistics and spatio-temporal data*. Wiley series in probability and statistics. Chichester: Wiley.
- Shumway, R. H. and D. S. Stoffer (2006). *Time series analysis and its applications with R-examples* (3 ed.). Springer texts in statistics. New York: Springer.
- So, W. (1990). Facial structures of Schatten p -norms. *Linear and Multilinear Algebra* 27(3), 207–212.
- Statistisches Bundesamt (2012, July). *Consumer Price Index: Quality Report*. Wiesbaden: Statistisches Bundesamt.
- Statistisches Landesamt Berlin (1997). *Statistisches Jahrbuch 1997*. Statistisches Landesamt Berlin.

- Statistisches Landesamt Berlin (1998). *Statistisches Jahrbuch 1998*. Statistisches Landesamt Berlin.
- Stewart, G. W. (1973). Error and perturbation bounds for subspaces associated with certain eigenvalue problems. *SIAM review* 15(4), 727–764.
- Stewart, G. W. (1997). On the weighting method for least squares problems with linear equality constraints. *BIT Numerical Mathematics* 37(4), 961–967.
- Stewart, G. W. and J. Sun (1990). *Matrix perturbation theory*. Computer science and scientific computing. Boston: Acad. Press.
- Stock, J. H. and M. W. Watson (2011). Dynamic factor models. *Oxford Handbook of Economic Forecasting* 1, 35–59.
- Strang, G. (2005). *Linear algebra and its applications* (4 ed.). Cengage Learning.
- Ströhl, G. (1994). Zwischenörtlicher Vergleich des Verbraucherpreisniveaus in 50 Städten. *Wirtschaft und Statistik* 6(1994), 415–434.
- Toh, K.-C. and S. Yun (2010). An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems. *Pacific Journal of Optimization* 6(615-640), 15.
- Trefethen, L. N. and D. Bau (1997). *Numerical linear algebra*. Philadelphia: SIAM, Soc. for Industrial and Applied Math.
- Vaart, A. W. v. d. and J. A. Wellner (2000). *Weak convergence and empirical processes* (Corr. 2. print. ed.). Springer series in statistics. New York: Springer.
- Vershynin, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In Y. C. Eldar and G. Kutyniok (Eds.), *Compressed Sensing*, pp. 210–268. Cambridge University Press.
- Wedin, P. Å. (1983). On angles between subspaces of a finite dimensional inner product space. In *Matrix Pencils*, pp. 263–285. Springer.
- Yuan, M., A. Ekici, Z. Lu, and R. Monteiro (2007). Dimension reduction and coefficient estimation in multivariate linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69(3), 329–346.
- Zhang, Y. and X. R. Li (1996). Fixed-interval smoothing algorithm based on singular value decomposition. In *Proceedings of the IEEE International Conf. on Control Appl.*, pp. 916–921. IEEE.
- Zhu, P. and A. V. Knyazev (2013). Angles between subspaces and their tangents. *Journal of Numerical Mathematics* 21(4), 325–340.
- Zi̇etak, K. (1988). On the characterization of the extremal points of the unit sphere of matrices. *Linear Algebra and its Applications* 106, 57–75.
- Zi̇etak, K. (1993). Subdifferentials, faces, and dual matrices. *Linear algebra and its applications* 185, 125–141.

Index

- ε -covering number ($N(S, \|\cdot\|, \varepsilon)$), 18
- ε -net, 17
- affine
 - function, 48
- alignment
 - perfect, 64
- angle
 - between a point and a subspace, 91
 - between two points, 19
 - largest principal, between two eq.-dim. subspaces (θ_{\max}), 89
 - least nonzero principal, between two eq.-dim. subspaces ($\theta_{\min, \neq 0}$), 89
 - principal, between two eq.-dim. subspaces, 89
- autoregressive approximation factor
 - uniform, 71
- basis, 15
 - standard basis of $\mathbb{R}^{m \times k}$ ($B_{i,j}$), 15
 - standard basis of \mathbb{R}^k (e_i), 15
 - standard basis of \mathbb{S}^m ($\bar{B}_{i,j}$), 16
- biconjugate (function) ($g_{\star\star}$), 48
- boundary point, 43
- Cauchy-Schwarz inequality, *see* inequality
- Cholesky decomposition, 28
- Cholesky factor
 - of a Gramian, 28
- codimension (codim), 25
- column space (img), 16
- compatibility constant
 - lower ($c(W, \|\cdot\|, \|\cdot\|')$), 17
 - upper ($C(W, \|\cdot\|, \|\cdot\|')$), 17
- complement, 23, 25
 - orthogonal ($^\perp$), 23, 24
- cone, 43
- conjugate (function) (g_\star), 48
- contraction, *see* linear map
- convex
 - combination, 42
 - envelope, 48
 - function, *see* convex function
 - hull (conv), 42
 - set, *see* convex set
- convex function, 45
 - curvature constant, strongly convex function, 48
 - epigraph, *see* epigraph
 - minimizer, *see* minimizer
 - strictly, 48
 - strongly, 48
 - subdifferential, *see* subdifferential
 - subgradient, *see* subgradient
- convex set, 42
 - face, 44
 - exposed, 44
 - exposing element, 44
 - proper, 44
 - normal cone (ncone), 43
 - supporting hyperplane, 43
- coordinates, 15
- curvature constant
 - strongly convex function, *see* convex function
- decay rate
 - uniform, 70
- density, 78
- diagonal entries, *see* matrix
- differentiable, 46
- dimension (dim), 15
- distinct singular value, *see* singular value
- distribution
 - density, *see* density
 - empirical, *see* empirical distribution
 - Gaussian, *see* Gaussian distribution
 - of a random variable, 68, 78
 - of a random vector, 66
 - subgaussian, *see* subgaussian
- dual norm, 33

dual problem, *see* least-squares
 empirical distribution
 of a random vector, 66
 epigraph
 of a convex function (epi), 45
 of a function (epi), 48
 Euclidean space, 19
 expectation
 of a random matrix (\mathbb{E}), 27
 of a random variable (\mathbb{E}), 27
 exposed face, *see* convex
 exposing element, *see* convex
 extended model space, *see* model space

 face, *see* convex
 factor, 57
 Fenchel-Young inequality, *see* inequality
 filtering, 113
 finite dimensional, 15
 Frobenius norm, *see* norm

 Gaussian distribution, 78
 gradient, 46
 Gram-Schmidt orthogonalization, 21
 Gramian ($\langle Y, Y \rangle$), 26
 Cholesky factor, *see* Cholesky factor

 halfspace, 42
 Hanson-Wright inequality, *see* inequality
 Householder transform, 109
 hyperplane, 42
 supporting (of a convex set), *see* convex

 idempotence, *see* linear map
 identity
 matrix, *see* matrix
 image (img)
 of a linear map on a(n abstract) linear space W , 26
 of a linear map on \mathbb{R}^k , 16
 image measure, 39
 inequality
 Cauchy-Schwarz, 33
 Fenchel-Young, 48
 Hanson-Wright, 79
 von Neumann trace, 33
 inner product
 matrix of two linear maps ($\langle Y, X \rangle$), 26
 interior, 43

 kernel (ker)
 of a linear map on a(n abstract) linear space W , 26
 of a linear map on \mathbb{R}^k , 16

 least-squares
 (primal) problem, 49
 dual problem, 50
 left singular vector, *see* singular vector
 linear combination, 15
 linear map
 block/bracket notation ($[y_1 \cdots y_k]$), 16
 column, 16
 column space, *see* column space
 contraction, 25
 idempotent, 25
 image, *see* image
 kernel, *see* kernel
 rank, *see* rank
 reflection, *see* reflection
 row, 28
 singular subspace, *see* subspace
 singular value, *see* singular value
 singular value decomp., *see* singular value
 singular vector, *see* singular vector
 unitary, *see* unitary map
 linearly independent, 15
 Lipschitz constant (L), 18

 matrix
 column, 16
 column space, *see* column space
 diagonal, 35
 diagonal entries, 19
 entries of, 15

Gramian, *see* Gramian
 identity, 45, 58
 image, *see* image
 kernel, *see* kernel
 positive definite, 26
 positive semidefinite, 26
 quadratic, 16
 rank, *see* rank
 representation, 15
 row echelon, *see* triangular matrix
 singular subspace, *see* subspace
 singular value, *see* singular value
 singular value decomp., *see* singular value
 singular vector, *see* singular vector
 skew-symmetric, *see* skew-symm.
 standard basis, *see* basis
 symmetric, 16
 spectral decomposition, *see* spectral decomposition
 standard basis, *see* basis
 trace, *see* trace
 transition, *see* transition matrix
 transpose, *see* transpose
 upper triangular, *see* triangular matrix

mean, *see* expectation
 minimizer, 47
 ε -almost, 48
 set of (argmin), 47
 model space
 k -model space (V_k), 61
 extended k -model space (\bar{V}_k), 63
 multiplicity
 singular value, *see* singular value

nontrivial
 linear space, 15
 sequence in a linear space, 21
 norm, 17
 L^2 ($\|\cdot\|$), 17
 dual, 33
 Euclidean ($\|\cdot\|$), 17
 Frobenius ($\|\cdot\|$), 17
 Frobenius, on $W^{\times k}$ ($\|\cdot\|$), 32
 lower subspace) comp. constant, *see* compatibility constant
 nuclear, on $W^{\times k}$ ($\|\cdot\|_{\text{nuc}}$), 32
 operator, on $W^{\times k}$ ($\|\cdot\|_{\text{op}}$), 32
 subgaussian ($\|\cdot\|_{\psi_2}$), 79
 topology, 17
 unitarily invariant, 32
 upper comp. constant, *see* compatibility constant
 normal cone, *see* convex
 nuclear norm, *see* norm
 null space, *see* kernel

oblique
 projection, *see* projection
 projector, *see* projector
 operator norm, *see* norm
 orthogonal, 20
 complement, *see* complement
 projection, *see* projection
 projector, *see* projector
 orthogonal complement, *see* complement
 orthonormal
 basis, 21

polarization identity, 20
 prediction, 100
 principal angle, *see* angle
 probability
 with, at least, 67
 projection, 23, 24
 oblique
 composition with a linear map X
 ($\hat{X}_{V/U}, \hat{X}_I$), 26
 oblique ($\hat{\cdot}_I$), 26
 orthogonal ($\hat{\cdot}$), 23, 24
 prediction, *see* prediction
 projector
 oblique ($\hat{\cdot}_I, P_I$), 26
 orthogonal
 composition with a linear map X
 (\hat{X}_V, \hat{X}), 26
 orthogonal ($\hat{\cdot}, P$), 25

- relation between oblique and orthogonal projectors, 26
- random
 - field, 59
 - matrix, 27
 - expectation, *see* expectation
 - variable, 16
 - expectation, *see* expectation
 - variance, *see* variance
 - vector, 27
 - variance matrix, *see* variance
- rank (rk), 16
- reflection, 108
- representation
 - of a finite sequence, 21
- residual, 23, 24
 - of oblique projection ($\tilde{\cdot}$), 26
 - of orthogonal projection ($\tilde{\cdot}$), 23, 24
- right singular vector, *see* singular vector
- sign function (*sign*), 38
- singular subspace, 34
- singular value ($\sigma_i(\cdot)$), 32
 - decomposition, 32
 - distinct, nonzero ($\bar{\sigma}_i(\cdot)$), 34
 - least nonzero ($\sigma_{\min, \neq 0}(\cdot)$), 32
 - maximal (σ_{\max}), 32
 - multiplicity, 34
- singular vector
 - left, 31
 - right, 31
- skew-symmetric ($A^T = -A$), 21
- smoothing, 113
- span ($\text{span}\{\dots\}$), 15
- spanning sequence, 15
- spectral decomposition
 - of a symmetric matrix, 34
- strictly convex, *see* convex function
- strongly convex, *see* convex function
- subdifferential
 - of a convex function ($\partial f(\cdot)$), 46
 - of a function ($\partial g(\cdot)$), 48
- subgaussian, 68, 78
 - norm, *see* norm
- subgradient
 - of a convex function, 46
 - of a function, 48
- sublevel sets, 49
- subspace, 16
 - codimension, *see* codimension
 - complement, *see* complement
 - complementary, *see* complementary
 - sum of a sequence of subspaces, 34
 - sum of two subspaces, 25
- superspace, 16
- supporting hyperplane, *see* convex
- time series, 59
- trace
 - cyclic property, 19
- transition matrix, 59
- transpose (T), 19
- triangular matrix
 - row echelon, 21
 - upper, 21
- uniform autoregressive approximation
 - factor, *see* autoregressive
 - approximation factor
- uniform decay rate, 70
- union bound, 69
- unit sphere, 17
- unitarily invariant, *see* norm
- unitary map, 20
- variance
 - matrix ($\text{var}(\cdot)$), 27
 - of a random variable ($\text{var}(\cdot)$), 27
- vector, 15
 - i -th entry of, 15
 - as matrix, 15
 - representation, 15
 - standard basis, *see* basis
- von Neumann trace inequality, *see* inequality
- with probability at least, 67