Nonlinear and Robust Independent Component Analysis for Stochastic Processes

Alexander Schell
St John's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2022

Für meine Eltern

Acknowledgements

I would like to thank my supervisor Harald Oberhauser for his guidance and support and for his continued interest in my work.

Many thanks are also due to Markus Reiß and Zhongmin Qian for their assessment of this thesis and constructive feedback. Further, I am thankful to Aapo Hyvärinen for his advice and helpful comments on an early version of this work.

I would also like to extend my gratitude to Sam Cohen and Terry Lyons for their feedback on my confirmation thesis, and I am especially grateful to Sam Cohen and Philip Maini for their much appreciated academic and personal advice.

Further thanks are due to Christina Zou for her entertainment and counsel in virtually all matters of life, and to my office mates Patric Bonnier and Csaba Toth for our various discussions and for creating a friendly and productive working atmosphere. I would also like to thank the past and current members of Mathematical Institute, and especially those of the Stochastic Analysis and Mathematical Finance Group, for many refreshing and inspiring encounters; in particular I thank Melanie Beckerleg, Ching Lok Chong, James Foster, Patrick Kidger, Saad Labyad, Jiawei Li, Vlad Margarint, Abraham Ng, and Cris Salvi. I would further like to acknowledge Mazen Ali, Bogdan Dina, Maxim Efremov, Daniel Karrer, Simon-Philipp Merz, Martin Rasmussen, André Steck, and Rico Zacher for many enjoyable and constructive conversations.

I am especially thankful to Tara, for our time together and for making my life in Oxford so much more enjoyable than it would have been without her.

This work was made possible thanks to funding from the Oxford-Cocker Graduate Scholarship and the Mathematical Institute Scholarship. I am further grateful to St John's College for accommodation, hospitality, and good cheer.

Finally and above all, my deepest gratitude goes to my family, especially to my parents Lilly and Erich, to my brother Maximilian and to my grandparents Lilly and Nikolaus, whom I cannot thank enough for their extraordinary support and unending encouragement throughout the past years and at every previous stage of my life. This thesis is dedicated to them.

Abstract

This thesis is organised into two main parts, which are both preceded by a joint introduction and a brief chapter on the signature analysis of time-ordered data.

In the first part, we study the classical problem of recovering a multidimensional source signal from observations of nonlinear mixtures of this signal. We show that this recovery is possible (up to a permutation and monotone scaling of the source's original component signals) if the mixture is due to a sufficiently differentiable and invertible but otherwise arbitrarily nonlinear function and the component signals of the source are statistically independent with 'non-degenerate' second-order statistics. The latter assumption requires the source signal to meet one of three regularity conditions which essentially ensure that the source is sufficiently far away from the non-recoverable extremes of being deterministic or constant in time. These assumptions, which cover many popular time series models and stochastic processes, allow us to reformulate the initial problem of nonlinear blind source separation as a simple-to-state problem of optimisation-based function approximation. We propose to solve this approximation problem by minimizing a novel type of objective function that efficiently quantifies the mutual statistical dependence between multiple stochastic processes via cumulant-like statistics. This yields a scalable and direct new method for nonlinear Independent Component Analysis with widely applicable theoretical guarantees and for which our experiments indicate good performance.

In the second part, we revisit the problem of blind source separation from the perspective of statistical robustness. Blind source separation (BSS) aims to recover an unobserved signal S from its mixture $X = f(S)$ under the condition that the effecting transformation f is invertible but unknown. This being a basic problem with numerous practical applications, a fundamental issue is to understand how the solutions to this problem behave when their supporting statistical prior assumptions are violated. In the classical context of linear mixtures, we present a general framework to analyse such violations and quantify the effect they have on the blind recovery of S from X. Modelling S as a multidimensional stochastic process, we introduce an informative topology on the space of possible causes underlying a mixture X and show that the behaviour of a generic BSS-solution in response to general deviations from its defining structural assumptions can be profitably analysed in the form of explicit continuity guarantees with respect to this topology. This enables a flexible and convenient quantification of general model uncertainty scenarios and amounts to the first comprehensive robustness framework for BSS. Our theory is entirely constructive, and we demonstrate its utility with a number of statistical applications.

Contents

List of Figures

Table of Symbols

Symbol	Meaning	Page
$[k]$	$:= \{1,\dots,k\}$, and $[k]_0 := [k] \cup \{0\}$ $(k \in \mathbb{N})$.	vi
S_d	$:= \{\tau : [d] \to [d] \mid \tau$ is bijective$\}$; the group of all permutations of $[d]$.	vi
$\lvert \cdot \rvert$	$: \mathbb{R}^n \to \mathbb{R}_+$, Euclidean norm on $\mathbb{R}^n$; in partic., $\lvert \cdot \rvert$ is the absolute value on $\mathbb{R} = \mathbb{R}^1$.	vi
$\mathcal{M}_1(E)$	$:= \{\mu : \mathcal{B}(E) \to [0,1]\}$; set of all Borel probability measures on a (topological) space E.	vi
$\varphi\big\rvert_{\tilde{A}}$	$: \tilde{A} \to B$; the restriction of a map $\varphi : A \to B$ to a subset $\tilde{A} \subseteq A$.	15
$[d]^\star$	$:= \bigcup_{m\geq 0}[d]^{\times m} = \{\emptyset, 1, 2, \dots, 11, 12, 112, \dots\}$; the set of all multi-indices with entries in $[d]$, where $[d]^0 := \{\emptyset\}$ and $i_1 \cdots i_m \equiv (i_1, \dots, i_m)$.	21
$\delta_{u,v}$	$\equiv \delta_{uv} := \mathbb{1}_{\{v\}}(u)$; the Kronecker delta over two elements u and v.	26
Δ_d	$:= \{\Lambda = (\lambda_i \cdot \delta_{ij}) \in \mathrm{GL}_d \mid \lambda_1, \dots, \lambda_d \in \mathbb{R} \setminus \{0\}\}$; the group of (real) invertible diagonal $d \times d$ matrices.	35
P_d	$:= \{(\delta_{\sigma(i),j})_{i,j\in[d]} \in \mathrm{GL}_d \mid \sigma \in S_d\}$; the $d \times d$ permutation matrices.	35
$\mathfrak{M} \cdot Y$	$:= \{M(Y) \equiv (M(Y_t)) \mid M \in \mathfrak{M}\}$; the image of a process $Y = (Y_t)$ under a set $\mathfrak{M}$ of transformations from $\mathbb{R}^d$ to $\mathbb{R}^d$(resp. from the spatial support (3.7) of Y to $\mathbb{R}^d$).	35
$\arg\min_{a\in A} \phi(a)$	$:= \phi^{-1}\big(\min_{a\in A} \phi(a)\big)$; the set of all points in A at which the function $\phi : A \to \mathbb{R}$ attains its global minimum.	36
GL_d	$:= \{A \in \mathbb{R}^{d\times d} \mid \det(A) \neq 0\}$; the general linear group of degree d over $\mathbb{R}$.	36
M_d	$:= \{M \in \mathrm{GL}_d \mid M = D \cdot P$ for $D \in \mathrm{GL}_d$ diagonal and $P \in \mathrm{P}_d\}$; the group of (real) monomial matrices of degree d.	36
$\mathcal{C}_d$	$\equiv C(\mathbb{I}; \mathbb{R}^d) := \{x : \mathbb{I} \to \mathbb{R}^d \mid \mathbb{I} \ni t \mapsto x(t) =: x_t$ continuous$\}$; the space of continuous paths from $\mathbb{I}$ (compact interval) into $\mathbb{R}^d$; we use $\mathbb{I} = [0,1]$ unless mentioned otherwise.	38

π^I_J	the canonical projection from $E_I := \{(u_i)_{i\in I} \mid u_i \in E_i$ for all $i \in I\}$ onto E_J ($(E_i \mid i \in I)$ some indexed family of sets, $J \subseteq I$); that is $\pi^I_J((u_i)_{i\in I}) = (u_i)_{i\in J}$ where the tuple-indexation follows the order of I and J, resp. The superscript I will be omitted if the domain of π^I_J is clear. E.g.: $\pi_{\{1,3,5\}}(x_1, x_2, \cdots, x_6) = (x_1, x_3, x_5)$, and $\pi_i := \pi_{\{i\}}$.	38
IC	A stochastic process in $\mathbb{R}^d$ is called *IC* if it has mutually independent component processes.	38
$\Delta_m(\mathbb{I})$	$:= \{(t_1, \ldots, t_m) \in \mathbb{I}^{\times m} \mid t_1 < \ldots < t_m\}$ ($m \in \mathbb{N}$); the (relatively) open m-simplex on $\mathbb{I}^{\times m}$; we denote $\Delta_m \equiv \Delta_m([0,1])$.	39
wrt./ s.t./ wlog	'with respect to'/ 'such that'/ 'without loss of generality'	40
$\mathrm{int}(A)$	the topological interior of a set $A \subseteq \mathbb{R}^d$ (wrt. the Euclidean topology).	42
J_φ or D_φ	$:= \left(\frac{\partial}{\partial x_j}\varphi_i\right)_{ij}$; the Jacobian of $\varphi \equiv (\varphi_1, \cdots, \varphi_d)^\intercal \in C^1(G;\mathbb{R}^d)$.	42
$C^{k,k}(D)$	$:= \left\{h : D \to \mathbb{R}^d \;\middle|\; h \in \mathrm{Diff}^k(G) \text{ for some open } G \supseteq D\right\}$, $D \subseteq \mathbb{R}^d$; the set of all C^k-invertible transformations on D.	46
$f_1 \times f_2$	$: U_1 \times U_2 \to V_1 \times V_2,\ (u_1, u_2) \mapsto (f_1(u_1), f_2(u_2))$; the Cartesian product of two maps $f_i : U_i \to V_i$ ($i = 1, 2$).	47
$\mathrm{Diff}^k(G)$	$:= \{h : G \to \mathbb{R}^d \mid h : G \twoheadrightarrow h(G) \text{ is a } C^k\text{-diffeomorphism}\}$; the set of all functions $h \in C^k(G;\mathbb{R}^d)$ which are one-to-one with $h^{-1} \in C^k(h(G);\mathbb{R}^d)$, for $G \subseteq \mathbb{R}^d$ some open subset.	48
$\mathrm{diag}_{i\in[d]}[a_i]$	$:= (a_i \cdot \delta_{ij})_{ij}$; the diagonal matrix with main diagonal $(a_1, \ldots, a_d)$.	52
$\nabla^\times$	$:= \left\{(\lambda_\nu) \in \mathbb{R}^d \;\middle|\; \exists\, i, j \in [d],\, i \neq j : \lambda_i = \lambda_j\right\}$; the set of all vectors in $\mathbb{R}^d$ whose coordinates are not pairwise distinct.	53
$d(a, B)$	$\equiv \mathrm{dist}_d(a, B) := \inf\{d(a, b) \mid b \in B\}$, for (M, d) a given metric space; the distance between a point $a \in M$ and a non-empty set Y of M.	94
$\mathfrak{C}_m$	the set of all cross-shuffles $\mathfrak{C} \equiv \bigsqcup_{k=2}^d \mathcal{W}_k$ (Proposition 4.2.2) of fixed word-length m.	95
$\mathcal{M}_1$	$:= \mathcal{M}_1(\mathcal{C}_d)$; set of Borel probability measures on $(\mathcal{C}_d, \lVert\cdot\rVert_\infty)$.	135
D_Y	see (3.7); the spatial support of a signal $Y \in \mathcal{M}_1$.	138
$C^{0,0}(D)$	$:= \{\tilde{f} : D \to \mathbb{R}^d \mid \tilde{f} \text{ homeom. onto } \tilde{f}(D)\}$; set of all homeomorphisms on D.	138
$\mathrm{ev}_X(g)$	$\equiv \mathrm{ev}(X, g) := g(X)$ for $g : D_X \to \mathbb{R}^d$; the evaluation map at the process X.	139
$C^{1,1}$	$\equiv C^{1,1}(\mathbb{R}^d) := C^{1,1}(\mathbb{R}^d;\mathbb{R}^d)$; group of C^1-diffeomorphisms on $\mathbb{R}^d$.	143
$\mathfrak{D}$ resp. $\dot{\mathfrak{D}}$	see (5.29)/(5.54); space of $\mathfrak{sig}$-coordinatisable resp. premetrizable signals.	149/159
$\bar{\mathbb{R}}$	$:= \mathbb{R} \cup \{-\infty, \infty\}$; the affine closure of the real numbers.	152
$[d]^\star_{\leq m}$	$:= \bigcup_{0\leq j\leq m} [d]^{\times j}$; the set of all multi-indices of order up to m ($m \in \mathbb{N}$).	157
$\mathrm{ddiag}_{i\in[d]}[a_i]$	$:= (a_i \cdot \delta_{ij})_{ij}$; the diagonal matrix with main diagonal $(a_1, \ldots, a_d)$; for $\theta = (\theta_{ij}) \in \mathbb{R}^{d\times d}$ we write $\mathrm{ddiag}(\theta) = \mathrm{ddiag}_{i\in[d]}[\theta_{ii}]$.	161

$\langle u, v \rangle_2$	$:= u_1 v_1 + \ldots + u_d v_d$; the dot product of two vectors $u = (u_i), v = (v_i)$ in $\mathbb{R}^d$.	171
Sym_d^+	$:= \{C \in \mathbb{R}^{d\times d} \mid C^\intercal = C,\ u^\intercal C u > 0$ for all $u \in \mathbb{R}^d \setminus \{0\}\}$; the set of all symmetric positive definite matrices in $\mathbb{R}^{d\times d}$.	173
$\mathrm{P}_d^\pm$	$:= \{(\epsilon \cdot \delta_{\sigma(i),j})_{i,j\in[d]} \in \mathrm{GL}_d \mid \epsilon \in \{0,1\},\ \sigma \in S_d\}$; signed permutation matrices.	208
wlog/wrt./s.t.	"without loss of generality"/ "with respect to" / "such that".	
i.o.w.	"in other words"	

Chapter 1

Introduction

A common problem in science and engineering is that an observed quantity, X, is determined by an unobserved source, S, which one is interested in. Denoting by f the deterministic relationship between X and S, one thus arrives at the equation

$$X = f(S) \tag{1.1}$$

where X is known but both the relation f and the source S are unknown.

The premise that the data X is determined by its source S reflects in the assumption that f is a deterministic function, while the premise that S can be completely inferred from X — i.e. that no information be lost in the process of going from S to X — is reflected in the assumption that the function f is one-to-one; for simplicity, it is typically also assumed that f is onto. Any function f of this kind will be referred to as a mixing transformation.

The central challenge, known as the problem of *Blind Source Separation (BSS)*, then becomes to infer — or 'identify' — the hidden source S from the given data X:

Under which assumptions is it possible to recover the source data S in (1.1) *if only its mixture X is observed? To what extent can such a recovery be achieved and how can it be performed in practice?* (1.2)

It is clear that without additional assumptions, the above problem of inference (1.2) is severely underdetermined: If X and equation (1.1) is the only information available but both f and S are unknown, then we may generally find infinitely many possible 'explanations' $(\tilde{S}, \tilde{f})$ for X which all satisfy (1.1) but are not otherwise meaningfully related to the true explanation (S, f) underlying the data. In many cases, however, this 'indeterminacy of S given X with f unknown' can be controlled by imposing certain statistical conditions on the source S.

The following simple example illustrates this situation.

Example 1.0.1. Suppose that you are on a video-call and want to follow the simultaneous speeches of two speakers S^1 and S^2, modelled as real-valued time series each. As the propagation of sound adheres to the superposition principle, the acoustic signals X^1 and X^2 that reach your left and right ear, respectively, may be modelled as linear mixtures $X^i = a_{i1}S^1 + a_{i2}S^2$ of the individual speech signals S^1 and S^2. Denoting $X \equiv (X^1, X^2)^\intercal$ and $S \equiv (S^1, S^2)^\intercal$ and $A \equiv (a_{ij}) \in \mathbb{R}^{2\times 2}$, the relation between the audio data X and its underlying sources S can hence be expressed by the model equation $X = A \cdot S$, which for A invertible is a special case of (1.1) for the linear map $f := A$. The above problem (1.2) then becomes to recover the constituent speeches S^1 and S^2 from their observed mixtures X^1, X^2 alone, given that the relationship between X and S is linear.

Without further assumptions, the true explanation (S, A) of the observable X cannot be distinguished from any of its 'alternative explanations' $\{(\tilde{S}, \tilde{A}) \equiv (B \cdot S, AB^{-1}) \mid B \in \mathbb{R}^{2\times 2}$ invertible$\}$. But if the speech signals S^1 and S^2 were assumed to be uncorrelated, say, then the above family of best-approximations of (S, A) reduced to $\{(\tilde{S}, \tilde{A}) \equiv (B\Lambda \cdot S, A\Lambda^{-1}B^\intercal) \mid \Lambda \in \mathbb{R}^{2\times 2}$ (invertible) diagonal, $B \in \mathbb{R}^{2\times 2}$ orthogonal$\}$;[1] hence if they are uncorrelated, S^1 and S^2 may be recovered from X uniquely up to scale and a rotation. ◆

This simple observation can be significantly improved by way of the classical Darmois-Skitovich theorem [130, 43, 145] which implies that for f linear, the original source S may be identified from X even up to scaling and a permutation of its components if S is modelled as a random vector whose coordinates S^i are not only uncorrelated but statistically independent. This mathematical insight, elaborated in P. Comon's seminal framework [37], quickly became the theoretical foundation of *Independent Component Analysis (ICA)*, a popular statistical method that has since seen far-reaching theoretical investigations and extensions, e.g. [9, 139], and has been successfully implemented in numerous widely-applied algorithms, e.g. [14, 27, 69, 79]; see for instance [51, 80, 111] as well as the monographs [38, 81] for an overview.

Comon's contribution is arguably the most conceptionally influential answer to the above inference task (1.2) to date that was both practically relevant and mathematically rigorous. However, Comon's approach applies to linear relationships (1.1) between X and S only, because among nonlinear mixing functions on $\mathbb{R}^d$ there are many 'non-trivial' transformations that preserve the mutual statistical independence of their input vectors [84].

[1] Indeed: The assumption of uncorrelatedness complements the original model equation (1.1) by the additional (statistical) source condition $\mathrm{Cov}(\tilde{S}, \tilde{S}) = \mathrm{Cov}(S, S) = \mathrm{I}_2$, which implies that $B^\intercal B = \mathrm{Cov}(\tilde{S}, \tilde{S}) = \mathrm{I}_2$ (where the components of $\tilde{S}$ are assumed to be scaled to unit variance).

This is a substantial limitation not only from a theoretical perspective but also in applications, where real-world data is often assumed to depend nonlinearly on certain nonredundant (independent) explanatory source signals and the instantaneous invertible nonlinear model (1.1) is deemed an adequate description of this dependence. See for instance [5, 40, 48, 70, 86, 94, 115] and the references therein for a few according example applications of nonlinear BSS ranging from the analysis of star clusters in interstellar gas clouds and biomedical tissue monitoring during surgery over electroencephalography and molecular simulation to statistical process monitoring, vibration analysis and stock market prediction.

Overcoming the traditional confinement to linearity has thus been a long-standing scientific endeavour, and the past twenty-six years have seen various attempts of establishing alternative identifiability approaches to recover multivariate data from their nonlinear transformations. Prominent ideas in this direction include the optimisation of mutual information over outputs of (adversarial) neural networks, e.g. [4, 22, 73, 93, 148], or the idea of 'linearising' the generative relation (1.1) by mapping the observable X into a high-dimensional feature space where it is then subjected to a linear ICA-algorithm [68].

More recently, the works of Hyvärinen et al. [82, 83, 85] achieved significant progress regarding the recovery of nonlinearly mixed sources with temporal structure (e.g. time series, instead of random vectors in $\mathbb{R}^d$) by first augmenting the observed mixture of these sources with an auxiliary variable such as time [82] or its history [83], and then training logistic regression to discriminate ('contrast') between the thus-augmented observable and some additional 'variation' of the data. This variation is obtained by augmenting the observable with a randomized auxiliary variable of the same type as before, thus linking the asymptotical recovery of the source $S = f^{-1}(X)$ to a trainable optimisation problem, namely the convergence of a universal function approximator (e.g. a neural network) learning a classification task. These identifiability results were recently extended and embedded into the context of variational autoencoders in [93].

Another basic question of substantial importance to both the theory and statistical practice of Blind Source Separation is the issue of statistical robustness, that is if and to what extent the solutions to (1.2) persist when their underlying assumptions on f and S are violated. This question is of significant interest for the blind inversion of linear mixtures already, not in the least due to the overwhelming dominance and utility of linear separation methods in most contemporary BSS applications. Yet despite explicit calls for research on this topic [80], existing investigations are few and confined to rather specific violation scenarios, numerical partialities or very restrictive assumptions overall [1, 11, 12, 24, 29, 30, 89, 92, 125, 138, 139, 144]; a comprehensive robustness analysis on the effects of general violations of the supporting identifiability assumptions for BSS or ICA, let alone for time-dependent sources, does not seem to have been presented yet.

In this thesis, we revisit the inference problem (1.2) for stochastic processes $X = (X_t)$ and $S = (S_t)$ with recent tools from stochastic analysis. Departing from the classical framework of Comon [37], the focus of this work is on the above-mentioned aspects of nonlinearity and robustness. The following section summarizes our main contributions to the literature.

1.1 Summary of Main Contributions

This thesis has two main parts. In the first part, consisting of Chapters 3 and 4, we propose a new approach to the problem of nonlinear blind source separation (1.2) for multidimensional time-dependent signals that leverages modern tools from stochastic analysis: For an unknown discrete- or continuous-time signal $S = (S_t)$ in $\mathbb{R}^d$ and an unknown function $f : \mathbb{R}^d \to \mathbb{R}^d$, a new statistical method to recover S from its pointwise transformation $X = f(S) \equiv (f(S_t))$ via 'signature cumulants' is presented.

In essence, we provide a new algorithm[2] that performs the inversion, or 'retransformation,'

$$X \longmapsto S \tag{1.3}$$

of the generative relation (1.1) in the case that f and S are not explicitly known and f is sufficiently differentiable and (by necessity) invertible[3] but otherwise *arbitrarily nonlinear*.

Finding ways to achieve this 'blind inversion' (1.3) has been of long-standing scientific interest, and efforts in this direction gave rise to an established area of specialised statistical research that has been very active for nearly three decades now. Apart from only a small number of exceptions, however, related works were predominantly confined to the very limiting assumption that the hidden relation f be a linear map on $\mathbb{R}^d$ – the few existing approaches towards the blind inversion of nonlinear causal relations were either heuristic or required f to belong to very narrowly defined function classes only, and it was not until the recent breakthroughs of Hyvärinen et al. that the first mathematically justified ideas for the blind inversion of general nonlinear relationships between X and S have emerged. Our work is a contribution to the dawning research on nonlinear blind inversion.

[2] That is, an explicitly computable map – or estimator, in the statistical sense – that takes in [a realisation of] the mixture X and returns an 'optimal' approximation of [the corresponding realisation of] S as an output.

[3] Invertible at least on the smallest subset of $\mathbb{R}^d$ which is actually reached by S, but see Definition 3.2.4 and (3.8).

The material of Chapters 3 and 4 is based on the paper [120], and the content of Chapter 5 is to appear as [119].

The second part of the thesis consists of Chapter 5 and presents the first comprehensive robustness framework for the blind recovery of linearly mixed time-dependent signals. Here, we provide a general and flexible topological notion of statistical robustness for the inverse problem (1.2), and demonstrate the theoretical and practical utility of our robustness concept on a generic BSS estimator that performs the inversion (1.3) under ICA-typical identifiability assumptions. The first part and the second part of this thesis are independent.

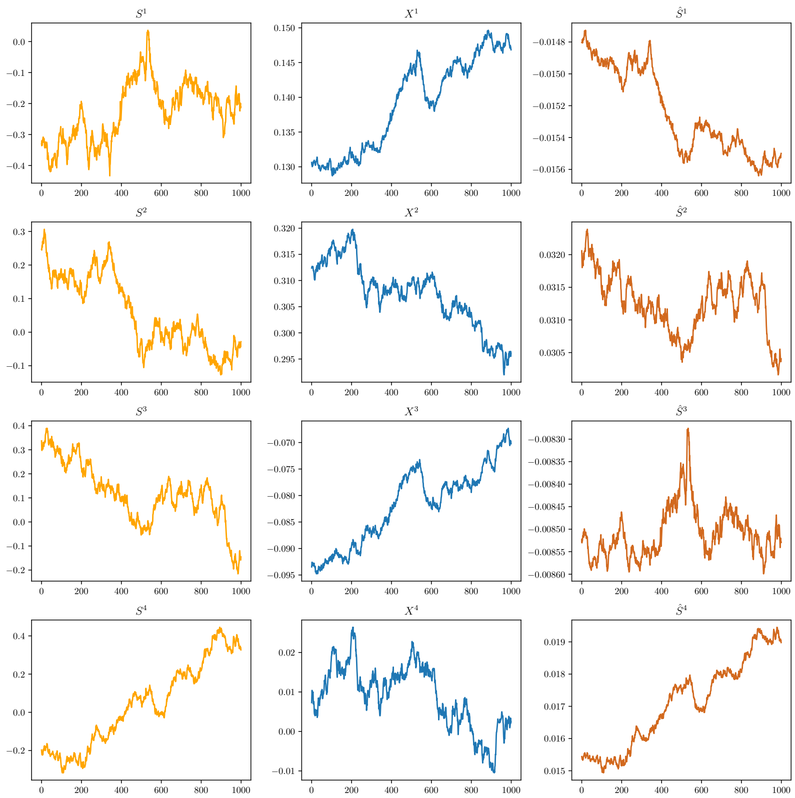

Figure 1.1: *A source S with four components $S^1, \ldots, S^4$ (orange) is mapped under some nonlinear transformation f, resulting in the observed mixture X (blue). We present a new method to recover the original source S from its mixture X up to a minimal deviation: Given X, this approach returns an estimate $\hat{S} = (\hat{S}^1, \cdots, \hat{S}^4)$ (brown) that approximates S up to the original order of its channels and a monotone scaling. In this example, $\hat{S}^1$ estimates the original component S^2, and $\hat{S}^2$ estimates S^3, and $\hat{S}^3$ estimates S^1, and $\hat{S}^4$ estimates S^4. (See Example 1.1.1 for details.)*

Below is a summarising discussion of our main contributions, listed in order of appearance.

1.1.1 Identifiability of Nonlinear Mixtures (Theorems 3.3.3, 3.3.7)

To achieve a meaningful recovery (1.3) of the source S from X, we need to compensate for the blindness regarding f and S by imposing some additional assumptions on the latter. The most established such assumption, and arguably the most relevant one in practice, is that the component signals of S are statistically independent; we adopt this assumption throughout. Many of the conceptional issues that arise in the nonlinear blind reconstruction of an independent-component source S from X can then be anticipated from the classical i.e. linear case $f \in \mathbb{R}^{d\times d}$ already. Similar to the classical case, cf. Theorem 3.1.1,

- the blindness[4] underlying (1.3) makes an exact recovery of S impossible, but statistical prior information on the source allows to identify S from X up to a minimal ambiguity, namely up to a permutation and monotone scaling of the source's original component signals;

- the estimates $\hat{S}$ of the original source S that are minimally ambiguous in the above sense preserve the initial condition of intercomponental independence (IC), but under some natural assumptions on S the converse is also true: those retransformations of X which are IC must be minimally ambiguous to S.

These insights into the blind inversion (1.3), which are rigorously discussed in Section 3.3, are the mathematical heart of our approach. Especially the equivalence stated in the last point, which is made precise in Theorems 3.3.3 and 3.3.7, is a central new finding:

Under some mild statistical conditions on the source S, we can show that the assumed IC property of the source is strong enough to trivialise[5] the action of any spatial diffeomorphism which preserves this property; in other words: their property of having minimal intercomponental statistical dependence distinguishes the minimally ambiguous estimates $\hat{S}$ of S from any other invertible nonlinear transformations of X.

This makes 'minimisation of intercomponent-dependence' an illuminating optimisation principle for the initially blind search for S, which immediately translates into the following strategy for the desired nonlinear blind inversion (1.3):

$$\text{as an estimate } \hat{S} \text{ for } S, \quad \text{choose} \quad \hat{S} = \theta_\star(X), \; \theta_\star \text{ invertible, s.t. } \theta_\star(X) \text{ is IC}\,; \tag{1.4}$$

i.e., the right retransformations of X are those that minimise intercomponental dependence.

[4] That is, the fact that the inverse problem (1.3) is inherently underdetermined since the constituents f and S of the RHS in (1.1) are both unknown.

[5] Here, 'trivialise' means reduce to the composition of a permutation and a componentwise monotone scaling.

As mentioned, the sources S for which this strategy works are those that 'carry their IC property well enough' for this property to characterise them, up to minimal ambiguity, among their (invertible) nonlinear transformations. But not every source is of this kind, as becomes particularly clear from considering two 'unrecoverable' statistical extreme cases: If the source S is deterministic,[6] then the IC property is void and a meaningful blind inversion (1.3) of the source's mixtures is generally impossible. If S is constant in time, i.e. $S = (Z)_{t\in[0,1]}$ for some random vector Z in $\mathbb{R}^d$, then the IC property on S cannot manifest cross-componentally over different time-points and is then generally too weak to support the strategy (1.4) for nonlinear mixtures, see [84] and Example 3.1.4.

These unidentifiable source types can be seen as degenerate extremes that are naturally interpolated by the mathematical model class of continuous-time stochastic processes, and said interpolation can be controlled at the level of the second-order finite-dimensional distributions (fdds) of such processes, see Section 3.2. In fact, we formulate three regularity assumptions on the family of fdds of a source S which enable the IC-based identifiability (1.4) of the source by ensuring that it is sufficiently far away from the above degeneracies (Section 3.3). More specifically, our non-degeneracy assumptions on the source require that sufficiently many of its fdds admit a probability density which is sufficiently complex in that it satisfies one of the following conditions:

(a) the density avoids local factorisations and is not of a certain 'pathological' Gaussian-like shape, as is made precise in Definition 3.3.2;

(b) the density has locally non-vanishing mixed log-derivatives that lie outside certain nullsets, as specified in Definition 3.3.5.

While the non-factorizability and non-vanishing-log-derivative conditions ensure that the source is 'stable enough' to make its IC property unfold[7] into its component signals in such a way that the ('residual') action inflicted upon S by the composition of the mixing transformation f with an IC-enforcing retransformation [as in (1.4)] does not collapse when considered jointly at different points in time, the exclusion of Gaussian-like shapes or algebraically degenerate density configurations ensures that this residual action on S is 'expressive' enough (as per implying a non-degenerate eigenspectrum of a Jacobian). All of this is made precise in Section 3.3.1 and the proofs of Theorems 3.3.3 and 3.3.7.

[6] That is, if S attains exactly one sample path with probability one.

[7] Instead of holding it merely within its fixed-time marginals, as in the generally unidentifiable case of IC random vectors in $\mathbb{R}^d$.

The source conditions (a) and (b) again generalise classical theory in a natural way (cf. Section 3.1 and the remarks on p. 43 and Remark 3.5.4), and in Section 3.5 we illustrate their broad applicability by compiling a set of widely used signal classes to which these conditions apply.

Thus far, our work has established the dependence-minimising approach (1.4) as a successful mathematical strategy to achieve the nonlinear blind source separation task (1.3), see Theorem 3.3.3 and Theorem 3.3.7: We identified natural probabilistic conditions (a) and (b) on the source which guarantee that its IC property manifests strongly enough to characterise that source among any invertible (re)transformations of X up to some inevitable ambiguity[8].

1.1.2 Nonlinear Blind Inversion via Optimisation (Theorem 4.2.3)

Building on the above, Chapter 4 proposes a way to turn this theoretical inversion strategy into a ready-to-use statistical method that can be easily implemented in practice. What we need to do for this is provide the observer of the mixture X with three things, namely

- a set Θ of invertible candidate demixing transformations on $\mathbb{R}^d$ which is 'large enough' to include approximations of the original inverse f^{-1} up to permutation and scale;[9]
- a 'pair of goggles' ϕ that allows the observer to gauge the degree of intercomponental statistical dependence of any given (re)transformation of X: the weaker the statistical dependence between the component signals of a process Y, the smaller shall be $\phi(Y) \in \mathbb{R}_{\geq 0}$; the desired inversion (1.3) is then performed [via (1.4)] by choosing those transformations $\theta(X)$, $\theta \in \Theta$, of X for which the value $\phi(\theta(X))$ is minimal;
- an automatable optimisation procedure that combines Θ and ϕ and returns

$$\theta_\star \in \operatorname*{arg\,min}_{\theta \in \Theta} \phi(\theta(X)) \quad \text{and then} \quad \hat{S} = \theta_\star(X) \tag{1.5}$$

  as the desired [minimally ambiguous] estimate of S, in accordance with (1.4).

The above is formalised in Theorem 4.2.3. A natural choice in practice is to implement Θ as the realisation space of an invertible artificial neural network (NN) with d input nodes, cf. Remark 4.2.4 (ii) and Section 4.8.3. Adding ϕ as a loss function to the NN, the optimisation (1.5) can then be performed efficiently via backpropagation; for details see Sections 4, 4.4 and 4.8.

[8] That is, as we recall, up to a permutation and monotone scaling of the source's original component signals.

[9] See the hypothesis on Θ that is formulated in Theorem 4.2.3.

Intuitively speaking, in the course of the optimisation (1.5) the observer gradually performs the desired inversion (1.3) directly by comparing different transformations of the data and choosing as most akin to the true inverse those that minimize the ϕ-quantified statistical dependence of X. For nonlinear f the theoretical justification of this procedure is new, while the underlying idea of source separation via quantified dependence minimisation is a well-established concept for the recovery of linearly mixed random vectors in $\mathbb{R}^d$, see e.g. Corollary 3.1.3.

Inspired by another classical concept, cf. (3.3) on page 36, in Section 4 we propose as dependence quantification ϕ a 'cross-cumulant'-based energy functional of the form

$$\phi(Y) = \sum_{m=2}^{\infty} \sum_{\boldsymbol{q}_m} \bar{\kappa}_{\boldsymbol{q}_m}(Y)^2 \tag{1.6}$$

where $\bar{\kappa}_{\boldsymbol{q}}(Y)$ denotes 'the (standardised) signature cumulant at index $\boldsymbol{q}$' of a stochastic process Y in $\mathbb{R}^d$, see Definition 4.1.2 and Notation 4.2.1, and the inner sums run over all 'cross-shuffles' of word-length m (see (4.21) on page 94). The entirety of all signature cumulants $(\bar{\kappa}_{\boldsymbol{q}}(Y))$, which can be thought of as a carefully chosen 'coordinate vector' for the distribution of the multidimensional stochastic process Y, provides a hierarchical and parsimous description of the statistical dependence relations within Y, which may occur simultaneously between coordinates and over different points in time. The functional (1.6) summarises the aspects of this description that are most central for us, namely 'how much' of this dependence there is between the multiple component signals of Y. Since the above ϕ vanishes over exactly those processes that are IC (Proposition 4.2.2), the functional (1.6) is well suited to operationalise the inversion strategy (1.4) via the optimisation scheme (1.5), as described in Theorem 4.2.3; further aspects are discussed in Sections 4.4 and 4.8.

1.1.3 Consistency (Theorem 4.4.13)

Up to this point, we discussed the method (1.5) in a setting where the whole distribution of X is assumed to be known. This idealisation is of course difficult to uphold in practice, where mixtures are typically not available as continuous-time stochastic processes and only discrete-time sample trajectories of X, i.e. finite sequences of data points in $\mathbb{R}^d$, are observed.

The statistical guarantees of Theorem 4.4.13 ensure that our method remains applicable under these practical constraints. More specifically, a statistical consistency analysis of the procedure (1.5) requires to simultaneously deal with

– time-discretization: if S, and hence X, are continuous-time signals, then 'full' sample observations of the underlying model X (i.e. continuous paths in $\mathbb{R}^d$) are not available in real-world applications, where only discrete-time data can be used;[10]

– finite samples: typically, one of two situations arise in applications. One is that n presumably independent [discrete-time] sample trajectories of the observable are recorded, e.g. medical recordings of n patients. The other situation is that only one [discrete-time] sample trajectory of X is given and ergodicity or mixing assumptions are invoked to make inference about the underlying distribution; for example, this situation is common in finance and economics.

We show that under general conditions, which for example are satisfied by many classical SDE models, our method (1.5) is (strongly) consistent in a sense that addresses both of these points: As the grid of observational time-points gets finer and the length of the observed time series increases, our method (1.5) produces a signal $\hat{S}$ that gets closer to the unobserved source signal S, even when the model for S is formulated in continuous time; see Theorem 4.4.13 for the precise statement. Additionally, Theorem 4.4.13 shows that our method is robust under approximations of the contrast function ϕ (for computational purposes, the series (1.6) of signature cumulants needs to be truncated in practice). The key ingredients to establish this result are tools from stochastic analysis, natural assumptions on the topology of function approximators (e.g., deep neural networks), and statistical approaches to the optimality of extremum estimators. Practitioners may find the displayed algorithm in Section 4.6 a useful summary.

Our exposition is complemented by a number of numerical examples (Section 4.8) which further illustrate the practical utility of our method by applying it to a series of nonlinear blind inversion problems (1.3) with multidimensional source signals in discrete and continuous time. As a concrete illustration of our blind inversion method (1.5) and its underlying procedure, let us draw on one of these examples here (see Section 4.8.3 for details).

Example 1.1.1. Imagine a professional context where you are interested in a set of 'hidden' quantities $S^1, \ldots, S^d$ that are related to some observable data $X^1, \ldots, X^d$ via an unknown invertible function $f : \mathbb{R}^d \to \mathbb{R}^d$. Assume further that these quantities are time-dependent, so that $S^i = (S^i_t)$ is a real-valued time series (in discrete or continuous time) and $(X^1_t, \cdots, X^d_t) = f(S^1_t, \cdots, S^d_t)$, and that you may regard $S^1, \ldots, S^d$ as mutually statistically independent. Suppose for example that $d = 4$ and your application context is the

[10] In spirit, this is similar to the well-developed statistical question of parameter estimation for stochastic differential equations where also only time-discretized sample trajectories are observed.

vibration analysis of wind turbines for fault detection. The quantities of interest S^i could then, e.g., be vibration responses excited by cracked gears or other engine faults in the turbine, which are mixed together during their transmission by an unknown, generally non-linear [48] mixing process f determined by the gearbox configuration; the resulting mixtures X^i are observable vibrations recorded by multi-channel sensors at the outside of the gear-box.[11] Statistically, your recorded data is a time-discretised sample $\mathfrak{x} \equiv (\mathfrak{x}^1_j, \cdots, \mathfrak{x}^4_j)^\intercal \in \mathbb{R}^{4\times\mathbb{N}}$ drawn from the stochastic process $X = (X^1, \cdots, X^4)$; it might locally look like the blue signals shown in the middle of Figure 1.1. In your search for the hidden vibrations $\mathfrak{s}$ that 'caused' your data, with $\mathfrak{s} \equiv (\mathfrak{s}^1_j, \cdots, \mathfrak{s}^4_j)^\intercal \in \mathbb{R}^{4\times\mathbb{N}}$ seen as a discretised sample observation of the process $S = (S^1, \cdots, S^4)$, your ignorance with regards to the actual relation f between $\mathfrak{s}$ and $\mathfrak{x}$ requires you to perform a blind inversion (1.3). You know that the closest ("minimally ambiguous") estimate $\hat{\mathfrak{s}}_\star$ of $\mathfrak{s}$ that you could then obtain is one that coincides with the original $\mathfrak{s}$ up to some permutation and scale, that is where

$$\hat{\mathfrak{s}}_\star = \Big(\alpha_1(\mathfrak{s}^{\tau(1)}_j), \cdots, \alpha_4(\mathfrak{s}^{\tau(4)}_j)\Big) \tag{1.7}$$

for τ some permutation of $\{1, \ldots, 4\}$ and $\alpha_i : \mathbb{R} \to \mathbb{R}$ strictly monotone. Based on the (very likely correct) assumption that the model $S = (S^1, \cdots, S^4)$ of $\mathfrak{s}$ satisfies one of the non-degeneracy assumptions (a) or (b), your hope is to arrive at such an optimal estimate (1.7) of $\mathfrak{s}$ by subjecting your data to the proposed inversion scheme (1.5). For this, you need to specify a sufficiently large (see above) set of retransformations Θ [of $\mathfrak{x}$] with a suitable approximation topology, e.g. a neural network, and cap the series (1.6) at some finite order m_0.[12] Next, compute from $\mathfrak{x}$ and each $\theta \in \Theta$ a consistent estimate $\hat{\phi}(\theta)$ of the (capped) contrasts $\phi(\theta(X))$. This may be done as summarized in the algorithm of Sect. 4.6. By our method (1.5), an empirical estimate of the best approximation (1.7) of $\mathfrak{s}$ is then obtained by finding a minimizer

$$\hat{\theta}_\star \in \arg\min_{\theta\in\Theta} \hat{\phi}(\theta) \quad \text{and setting} \quad \hat{\mathfrak{s}} := \hat{\theta}_\star(\mathfrak{x}). \tag{1.8}$$

This inversion scheme comes with (strong) consistency guarantees that ensure the convergence $\hat{\mathfrak{s}} \to \hat{\mathfrak{s}}_\star$ in the limit of infinite data, see Theorem 4.4.13. Applying this method to the sensory observations $\mathfrak{x}$ [← Fig. 1.1, blue column] of our example returns a source estimate $\hat{\mathfrak{s}} \in \mathbb{R}^{4\times\mathbb{N}}$ as shown in the rightmost (brown) displays of Fig. 1.1. Comparing this with the true vibration signals $\mathfrak{s}^1, \ldots, \mathfrak{s}^4$ [← Figure 1.1, orange column] underlying $\mathfrak{x}$, we find that

[11] This particular application context is motivated by and adapted from [48, 102] and references therein.

[12] For this example, detailed in Section 4.8.3, we used the network $\Theta := \Theta_2$ specified in (4.106) and capped the contrast (1.6) at $m_0 = 5$ (in general, m_0 may be chosen larger the more complicated the nonlinearity f is assumed to be). For more information, including on the mixing f and the optimisation (1.8), see also Rem. 4.8.4 and [140].

indeed: to a good approximation, the estimate $\hat{\mathfrak{s}}$ coincides with $\mathfrak{s}$ up to permutation and scale, as desired. ◆

We emphasize that the above methodology in its entirety, including any of our definitions or theorems, applies to both continuous-time and discrete-time signals alike[13]. The latter type includes signals that are "genuinely discrete", i.e. generated from a discrete-time process, and signals that are of continuous origin but "discretely observed", i.e. obtained from sampling a continuous-time process at a discrete set of time points. These cases are treated in detail in Sections 3.7, 4.4 and 4.5.0.1, which are referenced accordingly throughout the text.[14]

In total, Chapters 3 and 4 combine to a general and flexible new statistical method for the nonlinear blind source separation of multidimensional time-dependent signals.

1.1.4 Robust Blind Source Separation (Definition 5.3.10, Theorem 5.6.14)

In the second part of this thesis (Chapter 5) we return to the classical version of the BSS problem (1.2) where the idealised relation between source and observable is assumed linear, and study how the accuracy of the blind inversion (1.3) will then be affected when typical identifiability assumptions on f and S are violated. The question of how sensitively a BSS-solution reacts to violations of its underlying statistical assumptions is naturally of particular significance in applications, where identifiability assumptions are often seen as mere idealisations and their infringement is the norm, as the following example illustrates.

Example 1.1.2 (A More Realistic Cocktail Party)**.** Suppose that you monitor a crowded room in which there are $d \in \mathbb{N}$ people, say $(S^1, \ldots, S^d) =: S$, each engaged in conversation. Suppose further that each of their speeches are potentially relevant and informative to you, and that it is thus your objective to get a (consensual) recording of what each speaker S^i has said. To this end, you have installed d voice recorders at different places in the room.

However, instead of talking directly into your microphones, the speakers are so engaged that their utterances overlap acoustically so that, instead of a single clear voice per recorder, each of your microphones records a mixture $X^i = f_i(S^1, \ldots, S^d)$ of the d different voices. Will you still be able to recover the speaker's individual contributions from your recordings?

[13] For the case of discrete-time signals, everything basically applies as in the continuous case up to very minor modifications necessitated by the change from (path-)connected to discrete realisations of the underlying signals.

[14] For overview: Section 3.7.1 explicates our identifiability theory (Chapter 3) for the exact inversion of genuinely discrete mixtures, while the (asymptotic) recovery of signals from (samples of) their discretely observed nonlinear mixtures is developed as part of Section 4.4, see Theorem 4.4.13 in particular, and in Section 4.5.0.1.

In the first instance at least, it seems plausible (by linearity of sound propagation) to model each of the recordings as an unknown linear superposition of the speakers' voices, i.e. to assume that $X \equiv (X^1, \cdots, X^d) = f(S)$ for some $f = (f_1, \cdots, f_d) : \mathbb{R}^d \to \mathbb{R}^d$ linear. In addition, although possibly a bit harder to justify in this context, it surely is of considerable mathematical convenience to further assume that the sought voices are[15] statistically independent. From these angles, it is then compelling to try and recover the desired individual speech signals according to (5.1) by applying an ICA algorithm of one's choice to the recorded mixtures X.

While intuitively this is certainly a plausible idea, a few limiting considerations regarding the validity of this approach and its underlying structural assumptions quickly arise:

- due to the interactive nature of human conversation, it is almost certain that some or all of the speech signals – say those from speakers involved in a joint discussion – will not actually be statistically independent of each other, at least not entirely;
- owing to the spatial distance between the microphones and various phenomena affecting the transmission and propagation of sound such as acoustic dispersion or attenuation, the actual relation between the recorded mixtures and the speaker's individual voices is often only approximately linear but factually nonlinear, at least mildly so;
- the voice signals themselves are typically not iid in time but naturally exhibit a great deal of intertemporal statistical complexity, and both the physical propagation of these signals and their recording are rarely perfectly unperturbed but typically subject to various forms of noise and technical corruptions, such as ambient sounds or asynchronous recording at non-homogeneous sampling rates between different microphones.

Given these complications that make the true mixing constellation deviate from the simplifying idealisation of linearity and independence, will an ICA of the recorded mixtures still result in a meaningful reconstruction of the individual vocal expressions that, if probably not exact, is still 'good enough' to at least identify the main message that each speaker has communicated? Or will the violations of the underlying model assumptions affect the ICA performance so severely that the returned signals are distorted beyond intelligibility or the anticipated signal separation breaks down entirely? Is the transition between these outcomes continuous, and can the accuracy of the returned signals be informatively quantified against the 'amount of distortion' that causes the true scenario to deviate from its idealised identifiability assumptions? ◆

[15] . . . sampled from some sound-describing stochastic processes $S^i = (S^i_t)_{t\geq 0}$ in $\mathbb{R}$, which in turn are . . .

The points listed in the example each violate the consistency assumptions of most popular ICA methods, but a comprehensive theory to analyse the impact of such violations has not yet been proposed. In Chapter 5 we aim to fill this gap with a flexible topological framework built to quantify how the accuracy of a blind inversion method is affected by general infringements of its prior consistency assumptions. Drawing on general considerations of Huber [75], we propose a formal notion of robustness for BSS and ICA based on the idea of having the dependence (1.3) between (estimated) source and observable be pointwise-continuous at those causes of X (again described as a stochastic process) for which the underlying identifiability assumptions of the estimator are exactly met (Definition 5.3.10).

Our further approach then consists of an informative quantification of this continuity: Using a graded family of path-space functionals pertaining to the signature map from rough path theory (Section 2.2), we derive a coarse premetric topology on the space of laws of stochastic processes that we then extend to an explicit ICA-tailored topology on the space of all relevant hidden causes of a BSS-observable (Section 5.5). This topology is convenient to handle and weak enough to serve as a suitable reference topology for our robustness concept (Section 5.4) and thus combines to a flexible and model-free device to capture general violations of the statistical identifiability assumptions that lie at the heart of a given ICA method.

The theoretical considerations of Chapter 5 are brought to fruition in Section 5.6, where we introduce a generic BSS-estimator that achieves the blind inversion (1.3) under a conventional ICA-assumption (Theorem 5.6.10) and further enjoys general and informative robustness guarantees in the form of entirely explicit moduli of continuity with regards to the above ICA-topology (Theorem 5.6.14). The practical utility of our robustness approach is demonstrated with a number of statistical applications (Section 5.7) which include (non-asymptotic) robustness bounds for the practical infringements listed in Example 1.1.2.

Both parts of this thesis use as one of their central tools the signature map from rough paths theory. To account for its relevance to this thesis, the following chapter introduces the signature and provides a panoptic summary of those of its features that will be most important to us.

Chapter 2

Non-Local Analysis of Ordered Data

Several aspects of this thesis are built with recourse to the *signature representation* of a path in $\mathbb{R}^d$, a central perspective on time-ordered data that originates from rough paths theory [54, 104, 105]. Essentially, the signature offers an informative discretisation of a path in terms of an algebraically organised selection of certain non-local (in time) features of the path that are sufficient to characterise that path up to an inessential equivalence relation.[1] Following a definition of what (random) 'paths' actually are (Section 2.1), this chapter gives an essentially self-contained introduction to their non-local signature representation.

The route we take is novel and short: provided that a given path $x : I \to \mathbb{R}^d$ is continuous enough, say of bounded variation, then elementary functional analysis allows us to naturally identify this path with a bounded finite-rank operator on the space of continuous functions over I, see Section 2.2. This identification, which is a trade-off between dimension and linearity, has the advantage that bounded linear functionals admit powerful compressions via restriction to smaller subspaces of their domain (cf. Hahn-Banach; such a restriction is a lossless compression of the functional if the restricted functional admits a unique linear extension back to the full domain). The signature pertains to this compression principle:

Here, the subspace to which the operator x will be restricted is the span $\mathcal{A}_x$ of all iterated integral functions from the path x (Definition 2.2.3) and the signature of x, denoted $\mathfrak{sig}(x)$, is then precisely the transformation matrix (cf. linear algebra) of the restriction $x|_{\mathcal{A}_x}$ wrt. the 'eigenbasis' of iterated integral functions of x, see Section 2.2.2.[2] Now is when the

[1] Intuitively, we may liken the signature coefficients of a path to the interior angles of a triangle: Both characterise a global object (path; triangle) up to some 'inessential' equivalence relation ('tree-like equivalence' [66]; similarity) by exhausting the 'global geometry' of that object through an opportune selection of some of its intrinsic features that are characteristic of this object ('universal eigenstructure' of a path when seen as an operator on the continuous functions (see Section 2.2.2); the interior angles of a triangle).

[2] Since the assignment $x \mapsto x|_{\mathcal{A}_x}$ is functional, we may – under a slight abuse of terminology – also refer to the signature $\mathfrak{sig}(x)$ as the transformation matrix of the *full* operator x (instead of just its $\mathcal{A}_x$-restriction).

specific structure of the iterated integrals comes into play: the importance of these integral functions is that they are not algebraically independent but in fact highly interrelated (they form an algebra), and this additional level of internal structure on $\mathcal{A}_x$ implies (quite intricately [66]) that the transformation matrix $\mathfrak{sig}(x)$ characterises the underlying operator $x|_{\mathcal{A}_x}$ uniquely[3] and that the restriction[4] $x \mapsto x|_{\mathcal{A}_x}$ is unique up to tree-like equivalence. This combines to the initial statement that the signature characterises a path up to tree-like equivalence [66]. This proposition justifies to interpret the sequence of numbers $\mathfrak{sig}(x)$ as a coordinate vector that identifies its path x uniquely (up to tree-like equivalence).

By duality, this signature coordinatisation of paths naturally extends to an analogous representation for the laws of stochastic processes, resulting in what is known as the expected signature of a process (Section 2.2.3). As the signature coordinates of paths and their laws inherit from $\mathcal{A}_x$ some internal algebraic structure, it is opportune to account for this fact by treating the associated representations not as mere sequences in $\mathbb{R}$ but as elements of a suitably constructed coordinate space in which this structure can be better reflected and exploited. The space of formal power series turns out to be an appropriate choice (Section 2.3.1), and its endowment with natural structure-preserving bijections (Section 2.3.2) and a suitable topology (Section 2.3.3) will prove useful in later chapters.

2.1 Paths and Stochastic Processes

For a set V and a linearly-ordered set I, we call a *data stream* in V over I any function

$$x \;:\; I \to V, \quad t \mapsto x_t, \tag{2.1}$$

writing $x \equiv (x_t)_{t\in I} \equiv (x_t)$ for short. In this thesis, we assume V to be a Euclidean vector space and x to have distinguished 'features', by which we mean the coefficients of (x_t) with respect to some fixed basis in V. We further require any two data points x_s and x_t to be 'close by' if their indices $s, t \in I$ are close wrt. the order topology on I.

These assumptions are specific for the following central data type.

Definition 2.1.1 (Path, Components)**.** A data stream $x \equiv (x_t)_{t\in I}$ in V over I is called a *path* in V over I if $V \equiv (V, \langle\cdot,\cdot\rangle)$ is a Euclidean vector space, I is a compact interval and the map (2.1) is continuous. Given V and I, the family of all paths in V over I is denoted

$$\mathcal{C}(I;V) := \{x \,|\, x \text{ is a path in } V \text{ over } I\}, \quad \text{and} \quad \mathcal{C}_d := \mathcal{C}([0,1];\mathbb{R}^d) \quad (d := \dim(V))$$

[3] . . . and hence much more strongly than a generic transformation matrix defined on an (algebraically independent) basis would; recall that those 'generic' matrices characterise their underlying endomorphisms only up to similarity.

[4] . . . defined on the domain of all continuous bounded-variation paths from I to $\mathbb{R}^d$ . . .

is the canonical representative of $C([0,1];V)$ once a fixed basis $(e_i)_{i\in[d]}$ of V is chosen. In this case, we denote by $x^i \equiv (x^i_t)_{t\in I} := (\langle x_t, e_i\rangle)_{t\in I}$ the *i-th component of* x, where $i \in [d]$. We write $x \equiv (x^i)_{i\in[d]} \equiv (x^i)$ for a path x with components $x^1, \ldots, x^d$.

Clearly, if $x \equiv (x^i) \in \mathcal{C}_d$ then $x^i \in \mathcal{C}_1$ for each $i \in [d]$. Recall further that the uniform norm $\|x\|_\infty := \sup_{t\in[0,1]} |(x^1_t, \ldots, x^d_t)|$ defines a complete topology τ_∞ on $\mathcal{C}_d$ with associated Borel σ-algebra $\mathcal{B}(\mathcal{C}_d) := \sigma(\tau_\infty)$; under this topology, the Borel space $\mathcal{C}_d$ is Cartesian.

Lemma 2.1.2. *The Borel spaces* $\mathcal{C}_d$ *and* $(\mathcal{C}_1\times\cdots\times\mathcal{C}_1, \mathcal{B}(\mathcal{C}_1)^{\otimes d})$ *are canonically isomorphic.*

Proof. The space $\mathcal{C}_d$ inherits the Cartesian structure of $\mathbb{R}^d = \mathbb{R}\times\cdots\times\mathbb{R}$ via the canonical Borel isometry $\psi : \mathcal{C}_1\times\cdots\times\mathcal{C}_1 \to \mathcal{C}_d$ given by

$$\mathcal{C}_1^{\times d} \ \ni \ ((x^1_t)_{t\in\mathbb{I}}, \cdots, (x^d_t)_{t\in\mathbb{I}}) \quad \longmapsto \quad \left[\mathbb{I} \ni t \ \mapsto \ \sum_{i=1}^{d} x^i_t \cdot e_i\right] \ \in \ \mathcal{C}_d \tag{2.2}$$

for $(e_i)_{i\in[d]}$ the standard basis[5] of $\mathbb{R}^d$. Indeed: It is clear that ψ is a linear isometry with respect to the Banach norms $\|(x^1,\cdots,x^d)\|_\alpha := \max_{i\in[d]}\|x^i\|_\infty$ and $\|x\|_\beta := \max_{i\in[d]}\|\pi_i \circ x\|_\infty$ on $\mathcal{C}_1^{\times d}$ and $\mathcal{C}_d$, respectively (where $\pi_i := \langle\cdot, e_i\rangle$). As the norms $\|\cdot\|_\beta$ and $\|\cdot\|_\infty$ are equivalent and hence induce the same Borel-structure on $\mathcal{C}_d$, the Banach isometry (2.2) extends to an isomorphism of Borel spaces

$$(\mathcal{C}_d, \|\cdot\|_\infty) \ \cong \ (\mathcal{C}_1\times\cdots\times\mathcal{C}_1, \|\cdot\|_\alpha)$$

as claimed. But since the factors $\mathcal{C}_1$ are separable metric spaces, so is $\mathcal{C}_1^{\times d}$. As $\|\cdot\|_\alpha$ induces the topology of componentwise convergence, this implies that $\mathcal{B}(\mathcal{C}_1^{\times d}) = \mathcal{B}(\mathcal{C}_1)^{\otimes d}$ (see e.g. [16, Appendix M (p. 244)]), which proves the claim. □

Remark 2.1.3. We call a data stream (2.1) *discrete* if I is countable. Since a discrete data stream on V can be continuously injected into $\mathcal{C}_d$ via piecewise-linear interpolation (certainly if $|I| < \infty$; see Section 4.7.3), it will be no loss of generality for us to in the following focus on paths over the standard interval $[0,1]$ only.

The parametrisation (2.1) of a generic path $x \in \mathcal{C}_d$ is generally still 'too irregular' to allow for a meaningful global analysis of x. However, by imposing a control on the modulus

$$(s,t) \quad \mapsto \quad |x_t - x_s| \tag{2.3}$$

of a path x, we can specify regularity conditions on (2.1) that allow for a certain 'Hahn-Banach-type compression' of x which is known as the signature of x.

[5] Or, for the sake of generality, any fixed basis of V.

More specifically, what we want is a control on the modulus (2.3) which ensures that the parametrization (2.1) of a path $x = (x^i)$ is 'tight enough' for this path to induce a unique linear operator (specifically, the vectorial Stieltjes-Lebesgue measure $\mu_x = (\mathrm{d}x^i)$ identified as a finite-rank operator on $\mathcal{C}_1$; the latter identification is due to the Riesz-Markov-Kakutani representation theorem, which we recall to establish a one-to-one correspondence between Borel measures and linear functionals on locally compact spaces) that can then be 'faithfully compressed' to its restriction on a certain 'eigenspace' of x; see Section 2.2.

A sufficient condition on (2.3) for such an 'eigenspace compression' of a path is as follows.

Definition 2.1.4 (p-Variation)**.** Let $p \geq 1$. A data stream x in $(V, |\cdot|)$ over I is said to be of *finite p-variation* if I is a compact interval and

$$\|x\|_{p\text{-var}} := \left[\sup_{\mathcal{D}} \sum_{(t_\nu)\in\mathcal{D}} |x_{t_\nu} - x_{t_{\nu-1}}|^p\right]^{1/p} \; < \; \infty \tag{2.4}$$

where the supremum is over the set $\mathcal{D}$ of all (finite) dissections (t_ν) of I; set $\mathcal{BV}_p := \{x \in \mathcal{C}_d \mid \|x\|_{p\text{-var}} < \infty\}$ for the space of all such paths. We call *of bounded variation* a path with finite 1-variation, and denote by $\mathcal{BV} := \mathcal{BV}_1$ the space of all such bounded-variation paths.

Remark 2.1.5. For each $p \geq 1$, the space $\mathcal{BV}_p$ can be endowed with the *p-variation topology*, τ_p, which is the topology that the p-variation seminorm (2.4) defines on $\mathcal{BV}_p$. Moreover, as $p \mapsto \|x\|_{p\text{-var}}$ is decreasing we have that $\mathcal{BV}_p \subset \mathcal{BV}_q$ for any $q > p$, whence it makes sense to also consider $\mathcal{BV}_{1|p} := \mathrm{cl}_{\mathcal{BV}_p}(\mathcal{BV}; \tau_p)$, the topological closure of $\mathcal{BV}$ in $\mathcal{BV}_p$ with respect to τ_p; see e.g. [104, Sect. 1.2] for details. ◆

As is typical in probability and statistics, we may regard a given path $x \in \mathcal{C}_d$ as a *realisation* or *sample* of some $\mathcal{C}_d$-valued random variable X, that is consider the lift $x = X(\omega)$ for some $\omega \in \Omega$, where $(\Omega, \mathscr{F}, \mathbb{P})$ is some fixed probability space supporting $X : \Omega \to \mathcal{C}_d$ ('*random path*'). This 'data model' for (2.1) is naturally formalised by the following central notion.

Definition 2.1.6 (Stochastic Process)**.** A *continuous stochastic process in* $\mathbb{R}^d$ is a map

$$X : \Omega \to \mathcal{C}_d \quad \text{such that} \quad \omega \mapsto X(\omega) \equiv (X_t(\omega))_{t\in I} \quad \text{is} \quad (\mathscr{F}, \mathcal{B}(\mathcal{C}_d))\text{-measurable,} \tag{2.5}$$

where $\mathcal{B}(\mathcal{C}_d) = \sigma(\pi_t \mid t \in I)$ denotes the Borel σ-algebra on the Banach space $(\mathcal{C}_d, \|\cdot\|_\infty)$. Writing $X_t(\omega) \equiv (X_t^1(\omega), \cdots, X_t^d(\omega))^\intercal \in \mathbb{R}^d$ for each $\omega \in \Omega$, the scalar processes $X^i = (X_t^i)_{t\in I}$ $(i \in [d])$ are called the *coordinate processes* or the *components* of $X \equiv (X^1, \cdots, X^d)$, and the pushforward measure $\mathbb{P}_X := \mathbb{P} \circ X^{-1} : \mathcal{B}(\mathcal{C}_d) \to [0, 1]$ is called *the law of X*.

Remark 2.1.7. (i) From a more local perspective, Definition 3.2.1 is equivalent to the description of a continuous stochastic process X as an I-indexed family $X = (X_t)_{t\in I}$ of random vectors[6] X_t in $\mathbb{R}^d$ such that the map $X(\omega) \;:\; I \ni t \mapsto X_t(\omega) \in \mathbb{R}^d$ is continuous for each $\omega \in \Omega$, see e.g. [134, Sect. II.27].

(ii) In stricter terminology, a stochastic process $X \equiv (X_t(\omega))_{t\in I} : I \times \Omega \to \mathbb{R}^d$ as defined by (2.5) is called *jointly measurable.* If the underlying probability space $(\Omega, \mathscr{F}, \mathbb{P})$ is filtered then it can carry measurability notions stronger than (3.4) such as progressive measurability or predictability, see e.g. [157, Proposition 2.23]. For our purposes, the weak notion of measurability given by (2.5) will suffice.

Unless mentioned otherwise, any given stochastic processes is assumed continuous.

2.2 The Signature Representation of Paths and Their Laws

We now formally introduce the signature representation of a path as outlined on page 2.

2.2.1 Stieltjes-Embeddings of Paths and Iterated Integrals

Let $x : [0,1] \to \mathbb{R}^d$ be a path of bounded variation, and recall that paths of this regularity can be embedded into the space of signed measures via their Stieltjes measures.

Lemma 2.2.1 (Bounded-Variation Paths $\hookrightarrow$ Signed Measures)**.** *For each path $x \in \mathcal{C}_1$ of bounded variation, there exists a unique signed measure $\mu_x : \mathcal{B}([0,1]) \to \mathbb{R}$, the* Stieltjes measure *of x, such that $\mu_x = \mu_x^+ - \mu_x^-$ for measures $\mu_x^\pm : \mathcal{B}([0,1]) \to \mathbb{R}_+$ given by the identities $\mu_x^+((s,t]) = \frac{1}{2}\left[\|x|_{(s,t]}\|_{1\text{-var}} + x_{s,t}\right]$ and $\mu_x^-((s,t]) = \frac{1}{2}\left[\|x|_{(s,t]}\|_{1\text{-var}} - x_{s,t}\right]$ for each $0 \le s \le t \le 1$, where $x_{s,t} := x_t - x_s$. In particular, for each $\varphi \in \mathcal{C}_1$ we have that*

$$\int_0^1 \varphi_t \, \mu_x(\mathrm{d}t) \;=\; \int_0^1 \varphi_t \, \mathrm{d}x_t, \tag{2.6}$$

where the RHS is the classical Riemann-Stieltjes integral of φ with respect to x.

Although this one-to-one (cf. Rem. 2.2.2) correspondence between BV paths and (signed) measures is very well-known, let us account for its importance by including a brief proof.

Proof. Fix any $x \in \mathcal{C}_1$ with $\|x\|_{1\text{-var}} < \infty$, and abbreviate $v_t(x) := \|x \cdot \mathbb{1}_{[0,t]}\|_{1\text{-var}}$, $t \in [0,1]$. Note that $v_t(x) \le v_1(x) < \infty$ and $\|x|_{(s,t]}\|_{1\text{-var}} = v_t(x) - v_s(x)$ for any $0 \le s \le t \le 1$, in direct consequence of (2.4). Given the elementary measure-theoretical fact that the family $\mathcal{R} := \{(s,t] \mid 0 \le s \le t \le 1\} \subseteq 2^{[0,1]}$ of left-open subintervals of $[0,1]$ is a set-theoretic semiring generating the Borel-σ-algebra on $[0,1]$, i.e. that $\sigma(\mathcal{R}) = \mathcal{B}([0,1])$), Carathédory's

[6] For us every random vector in $\mathbb{R}^d$ is Borel, i.e. an $(\mathscr{F}, \mathcal{B}(\mathbb{R}^d))$-measurable function.

classical extension theorem guarantees that the set-functions $\mu_x^\pm : \mathcal{R} \to \mathbb{R}_+$ defined via σ-additive extension of the assignments $(s,t] \mapsto \frac{1}{2}\left[\|x|_{(s,t]}\|_{1\text{-var}} \pm x_{s,t}\right]$, respectively, can be uniquely extended to measures $\mu_x^\pm$ on $\mathcal{B}([0,1])$. Setting $\mu_x := \mu_x^+ - \mu_x^-$ then yields a signed measure on $\mathcal{B}([0,1])$, and for any other measures $\tilde{\mu}^\pm$ on $\mathcal{B}([0,1])$ with $\mu_x = \tilde{\mu}^+ - \tilde{\mu}^-$ it holds that $\tilde{\mu}^\pm = \mu_x^\pm$ as the Jordan decomposition of μ_x is unique (e.g. [64, Sect. 29 (p. 122)]).

To prove (2.6), let $\varphi : [0,1] \to \mathbb{R}$ be a continuous function and take any sequence $(\mathcal{Z}^n := \{0 = t_0^{(n)} < t_1^{(n)} < \ldots < t_{m_n}^{(n)} = 1\} \mid n \in \mathbb{N})$ of partitions of $[0,1]$ with vanishing mesh size $|\mathcal{Z}^n| := \max\{|t_j^{(n)} - t_{j-1}^{(n)} \mid j = 1, \ldots, m_n\} \to 0$ as $n \to \infty$. Then the RHS of (2.6) exists and we have (cf. [154, Satz 6.2])

$$\int_0^1 \varphi_t \, \mathrm{d}x_t = \lim_{n\to\infty} \sum_{j=1}^{m_n} \varphi_{t_{j-1}^{(n)}}\big(x_{t_j^{(n)}} - x_{t_{j-1}^{(n)}}\big) = \lim_{n\to\infty} \int_{[0,1]} \varphi_s^{(n)} \, \mu_x(\mathrm{d}s) \tag{2.7}$$

with $\varphi_s^{(n)} := \sum_{j=1}^{m_n} \varphi_{t_{j-1}^{(n)}} \cdot \mathbb{1}_{(t_{j-1}^{(n)}, t_j^{(n)}]}(s) + \varphi_0 \cdot \mathbb{1}_{\{0\}}(s)$.[7] As $\lim_{n\to\infty} |\mathcal{Z}^n| = 0$ and φ is right-continuous, $\lim_{n\to\infty} \varphi^{(n)} = \varphi$ holds μ_x-almost everywhere; moreover, $(\varphi^{(n)})$ is L^1-dominated due to $\varphi^{(n)} \leq \bar{\varphi} := M_\varphi \cdot \mathbb{1}_{[0,1]}$ with $M_\varphi := \max_{s\in[0,1]} |\varphi_s|$ for all $n \in \mathbb{N}$ and $\int_{[0,1]} \bar{\varphi} \, \mathrm{d}|\mu_x| \leq M_\varphi(\mu_x^+((0,1]) + \mu_x^-((0,1])) < \infty$, where $|\mu_x| = \mu_x^+ + \mu_x^-$ denotes the total variation of μ_x. The identity (2.6) hence follows from (2.7) and dominated convergence. □

Remark 2.2.2. (i) Denoting $\mathcal{M}_\mathbb{R}(\mathbb{I}) := \{\mu : \mathcal{B}([0,1]) \to \mathbb{R} \mid \mu \text{ is } \sigma\text{-additive}\}$ for the space of (finite) signed Borel-measures on the interval $\mathbb{I} := [0,1]$, notice that it is justified to call the assignment $\mathcal{C}_1 \cap \mathrm{BV} \ni x \mapsto \mu_x \in \mathcal{M}_\mathbb{R}(\mathbb{I})$ of Lemma 2.2.1 an 'embedding' because it is almost injective in the sense that: if $\mu_x = \mu_y$, then $x_{0,t}$ $\big(= \mu_x([0,t])\big) = y_{0,t}$ for each $0 \leq t \leq 1$ and hence $x = y + c$ for some $c \in \mathbb{R}$.

(ii) While any BV-path in $\mathcal{C}_1$ induces a Lebesgue-Stieltjes measure in $\mathcal{M}_\mathbb{R}(\mathbb{I})$, the converse association is also true in that for every $\mu \in \mathcal{M}_\mathbb{R}(\mathbb{I})$ there exists a unique (up to constant intercept) bounded-variation path $x \in \mathcal{C}_1$ such that $\mu = \mu_x$; see e.g. [53, Thm. 1.16] together with [64, Sect. 29 (p. 122)].

(iii) A path $x = (x^1, \cdots, x^d)$ is in $\mathcal{BV}$ iff $x^1, \ldots, x^d \in \mathcal{C}_1 \cap \mathrm{BV}$, and we have the embedding

$$\mathcal{BV} \ni x \equiv (x^i) \hookrightarrow \mu_x \equiv (\mu_{x^1}, \cdots, \mu_{x^d}) \in \mathcal{M}_\mathbb{R}(\mathbb{I})^{\times d}. \tag{2.8}$$

If α, β are continuous bounded-variation paths from $[0,1]$ to $\mathbb{R}$ then so is their *Stieltjes product*

$$\alpha \rtimes \beta := \left(\int_0^t \alpha_s \, \mathrm{d}\beta_s\right)_{0\leq t\leq 1}. \tag{2.9}$$

[7] Notice that $\mu_x\big((t_{j-1}^{(n)}, t_j^{(n)}]\big) = \mu_x\big((0, t_j^{(n)}]\big) - \mu_x\big((0, t_{j-1}^{(n)}]\big) = \big(\mu_x^+\big((0, t_j^{(n)}]\big) - \mu_x^+\big((0, t_{j-1}^{(n)}]\big)\big) + \big(\mu_x^-\big((0, t_{j-1}^{(n)}]\big) - \mu_x^-\big((0, t_j^{(n)}]\big)\big) = \frac{1}{2}\big[2x_{t_j^{(n)}} - 2x_{t_{j-1}^{(n)}}\big] = x_{t_j^{(n)}} - x_{t_{j-1}^{(n)}}$.

Let us now introduce a countable family of $\mathcal{C}_1$-nonlinearities (the 'iterated integrals' of x) that is 'characteristic of x' in the sense that their dual pairings (Riesz-Markov) with the Stieltjes measure $\mu_x \equiv (\mu_{x^1}, \ldots, \mu_{x^d}) \in \mathcal{M}_{\mathbb{R}}(\mathbb{I})^{\times d}$ characterise the path x up to a constant intercept at best and up to a minor ambiguity ('tree-like equivalence' [66]) at worst.

Definition 2.2.3 (Interpaths)**.** For a path $x = (x^i) \in \mathcal{C}_d$ of bounded variation, we define the *interpaths* $(\tilde{x}^w \mid w \in [d]^\star) \subset \mathcal{C}_1 \cap \mathrm{BV}$ *of* x recursively via

$$\tilde{x}^\emptyset := 1 \qquad \text{and} \qquad \tilde{x}^{wi} := \tilde{x}^w \rtimes (x^i - x_0) \tag{2.10}$$

for each letter $i \in [d]$ and every multiindex $w \in [d]^\star$. We call *eigenspan of* x the subspace

$$\mathcal{A}_x := \mathrm{span}_{\mathbb{R}}\{\tilde{x}^w \,|\, w \in [d]^\star\} \quad \text{of} \quad \mathcal{C}_1. \tag{2.11}$$

The following is well-known, see e.g. [66, Lemma 5.15] (or Proposition 5.9.3 (ii)).

Lemma 2.2.4. *The eigenspan $\mathcal{A}_x$ of a path $x \in \mathcal{BV}$ is a subalgebra of $\mathcal{C}_1$.*

Notice that the product (2.9) defines a noncommutative and associative binary operation on the path space $\mathcal{C}_1^{\mathrm{bv}} := \mathcal{C}_1 \cap \mathrm{BV}$, turning this space into an associative algebra. This allows us to identify the algebra of functions (2.11) with the ring of non-commutative polynomials in the formal (path) variables $x^1, \ldots, x^d$, cf. also Definition 2.3.1.

Remark 2.2.5. Written out, the interpaths (2.10) of $x = (x^1, \cdots, x^d)$ read

$$\tilde{x}^i_t = x^i_t - x^i_0, \quad \tilde{x}^{ij}_t = \int_0^t x^i_{0,s}\,\mathrm{d}x^j_s, \quad \tilde{x}^{ijk}_t = \int_0^t \int_0^s \int_0^r \mathrm{d}x^i_q \mathrm{d}x^j_r \mathrm{d}x^k_s, \quad \ldots$$

for each $i, j, k \in [d]$ and $0 \leq t \leq 1$, and are thus precisely the iterated integrals from [31].

2.2.2 The Signature – A Non-Local Characterisation of Paths

The unique[8] representation (2.8) of a continuous bounded-variation path $x : [0,1] \to \mathbb{R}^d$ as a tuple μ_x of signed Borel measures on $[0,1]$ allows us to view the infinite-dimensional object x from the convenient angle of functional analysis, namely as the finite-rank operator

$$x : \mathcal{C}_1 \to \mathbb{R}^d, \quad \varphi \mapsto x(\varphi) \equiv \int_0^1 \varphi\,\mathrm{d}\mu_x := \left(\int_0^1 \varphi\,\mathrm{d}\mu_{x^1}, \ldots, \int_0^1 \varphi\,\mathrm{d}\mu_{x^d}\right) =: (x^i(\varphi)). \tag{2.12}$$

This is a potent 'change-of-variables' for x, namely from its time-local default parametrisation (2.1) to the 'time-global' parametrisation (2.12) in which x now appears as a bounded linear map over the much richer domain $\mathcal{C}_1$ instead of a generally nonlinear function over $[0,1]$. This trade-off (2.12) between dimension and linearity brings an advantage:

While the time-domain $[0,1]$ in the (2.1)-parametrization of x has no distinguished subsets that allow for a particularly informative compression[9] of x in general, the vastly

---

[8] ... up to a constant intercept ...

[9] For instance, while the discretization $\mathcal{C}_d \ni x \mapsto (x_t)_{t \in \mathbb{I} \cap \mathbb{Q}} \in \mathbb{R}^{d \times \mathbb{N}}$ is an injection of $\mathcal{C}_d$ into a lower-dimensional domain (thus a compression in our general understanding of the term) it is usually not very insightful.

larger size of $\mathcal{C}_1$ in combination with the linearity of x in (2.12) opens up new possibilities to achieve such an 'informatively compressed' representation of x, namely in the spirit of the extension theory of Hahn-Banach: can we find a 'convenient' subspace $\mathcal{A}$ of $\mathcal{C}_1$ for which

(a) the restriction $x|_{\mathcal{A}}$ admits a unique linear extension to $L(\mathcal{C}_1, \mathbb{R}^d) \cong (\mathcal{C}_1')^{\times d}$, and

(b) $x|_{\mathcal{A}}$ is 'well-structured' and allows for an efficient and insightful analysis of x ?

Since the 'path-operator' x in (2.12) is bounded-linear[10] and hence in $L(\mathcal{C}_1, \mathbb{R}^d)$, point (a) is equivalent to the requirement that $(\mathcal{BV}/\mathbb{R}^d \cong)\, L(\mathcal{C}_1, \mathbb{R}^d) \ni x \mapsto x|_{\mathcal{A}}$ be an injective assignment and, thus, a lossless compression. Point (b) asks for the compressed representation $x|_{\mathcal{A}}$ of x to be more efficient to analyse than the full path and to provide insights that its default time-local representation (2.1) does not readily reveal.

By choosing $\mathcal{A}$ as the span $\mathcal{A}_x$ of its iterated integrals, the signature of a path x offers such a compressed representation of (2.12) for which the above, and much more, is satisfied.

Definition 2.2.6 (Signature)**.** For $x \equiv (x^i) \in \mathcal{BV}$ with interpaths $(\tilde{x}^w)_{w\in[d]^\star}$, the sequence[11]

$$\mathfrak{sig}(x) \equiv (\mathfrak{sig}_w(x) \mid w \in [d]^\star), \quad \text{with} \quad \mathfrak{sig}_\emptyset(x) := 1 \;\text{ and }\; \mathfrak{sig}_{wi}(x) := x^i(\tilde{x}^w) \tag{2.13}$$

for each letter $i \in [d]$ and every multiindex $w \in [d]^\star$, is called *the signature of* x. The associated map $\mathfrak{sig} : \mathcal{BV} \to \mathbb{R}^{[d]^\star}$, $x \mapsto \mathfrak{sig}(x)$, is called *the signature transform.*

Remark 2.2.7. From the definition of $\mathcal{A}_x$ and the linearity of (2.12), we clearly have

$$\mathfrak{sig}(x) \;\cong\; (x(\varphi) \mid \varphi \in \mathcal{A}_x) \qquad \text{(see (2.24))}, \tag{2.14}$$

and by the definition (2.10) of the interpaths $(\tilde{x}^w)$ of x (cf. Rem. 2.2.5) we also have

$$\mathfrak{sig}_{i_1\cdots i_m}(x) = \int_{0\leq t_1\leq\cdots\leq t_m\leq 1} \mathrm{d}x_{t_1}^{i_1} \cdots \mathrm{d}x_{t_m}^{i_m} \quad \text{for any } \; i_1, \ldots, i_m \in [d],\, m \in \mathbb{N}, \tag{2.15}$$

whence it is clear that Definition 2.2.6 coincides with [104, Definition 2.6], cf. also Sect. 2.3.3.

The signature owes its significance to the finer structural properties of its underlying algebra $\mathcal{A}_x$, as these give that [66]

[10] Indeed, since for each $\varphi \in \mathcal{C}_1$ we have $x(\varphi) = \sum_{i=1}^d \langle \mu_{x^i}, \varphi\rangle \cdot e_i$ with $|\langle \mu_{x^i}, \varphi\rangle| \leq \|\mu_{x^i}\|_{\mathrm{TV}}\|\varphi\|_\infty$ and $\|\mu_{x^i}\|_{\mathrm{TV}} = \|x^i\|_{1\text{-var}}$, it follows that $\|x\| \equiv \sup_{\|\varphi\|_\infty=1} |x(\varphi)| \leq d\|x\|_{1\text{-var}} < \infty$. (Here, $\langle \mu, \varphi\rangle := \int_0^1 \varphi \,\mathrm{d}\mu$ and $\|\cdot\|_{\mathrm{TV}}$ is the total variation norm on $\mathcal{M}_{\mathbb{R}}(\mathbb{I})$.)

[11] Note that (2.13) may be referred to as a sequence upon lexicographical ordering of the multiindices w.

(a′) the map $(L(\mathcal{C}_1, \mathbb{R}^d) \cong)\, \mathcal{BV}/\mathbb{R}^d \ni x \mapsto x|_{\mathcal{A}_x}$ is injective up to tree-like equivalence;[12]

(b′) the restriction $x|_{\mathcal{A}_x}$ is uniquely determined by $\mathfrak{sig}(x) \in \mathbb{R}^{[d]^\star}$

(iow., the difference between the tuple $(x(\varphi) \mid \varphi \in \mathcal{A}_x)$ and the map $x|_{\mathcal{A}_x}$ is not essential). The above combines to the faithfulness of the signature transform, see [66, Theorem 4]:

(c) For any $x, y \in \mathcal{BV}$, if $\mathfrak{sig}(x) = \mathfrak{sig}(y)$ then $x \doteq y$ up to tree-like equivalence.

Point (c) suggest to interpret $\mathfrak{sig}(x)$ as a (basis-free[13]) non-local coordinate vector of x.

Remark 2.2.8 (Rough Paths)**.** Definition 2.2.6 can also be generalised to paths beyond $\mathcal{BV}$. For paths rougher than bounded p-variation for $p \geq 2$, the essential 'iteration algebra' (2.9) underlying the representation (2.13), which for bounded variation paths holds by necessity, needs to be imposed up to a certain finite 'germ level' (the rougher the path, the higher up this imposition needs to go) from which the full signature representation can then be uniquely extended. This (mainly technical) extension is at the centre of the theory of rough paths, and we refer the interested reader to the monographs [54, 104, 105] for instructive accounts on this topic. ◆

This signature representation of a path enjoys a number of useful algebraic properties and invariances, see for instance [54, 104, 105] or Section 2.3.3 and Section 5.9.2.

2.2.3 The Expected Signature – Coordinates for Stochastic Processes

The signature transform compresses a bounded-variation path in $\mathcal{C}_d$ to an element in $\mathbb{R}^{[d]^\star}$ (Section 2.2.2). We will now extend this compression to the laws of random paths in $\mathbb{R}^d$. This extension can be motivated by an analogous relation between the classical moments for laws in $\mathcal{B}(\mathbb{R}^d)$ and the 'moment coordinates' of vectors in $\mathbb{R}^d$, see Remark 2.2.10 below.

Definition 2.2.9 (Expected Signature)**.** For $X = (X_t^1, \cdots, X_t^d)_{t\in[0,1]}$ a stochastic process in $\mathbb{R}^d$ with sample-paths of bounded variation, the collection of real numbers (if it exists) $\mathfrak{S}(X) := (\sigma_{\boldsymbol{i}}(X))_{\boldsymbol{i}\in[d]^\star}$ defined by the expected iterated Stieltjes integrals

$$\sigma_{\boldsymbol{i}}(X) := \mathbb{E}\Big[\int_{0\leq t_1\leq t_2\leq\cdots\leq t_m\leq 1} \mathrm{d}X_{t_1}^{i_1}\mathrm{d}X_{t_2}^{i_2}\cdots\mathrm{d}X_{t_m}^{i_m}\Big] \quad \text{for} \quad \boldsymbol{i} = (i_1, \ldots, i_m) \tag{2.16}$$

and with $\sigma_\emptyset(X) := 1$, is called *the expected signature of* X.

[12] Note that for the stronger assertion $(a)|_{\mathcal{A}=\mathcal{A}_x}$ to hold as stated, it is sufficient that the algebra $\mathcal{A}_x$ separates points (which holds, e.g., when one of the path components is strictly monotone). Indeed: If $\mathcal{A}_x$ separates points, then Lemma 2.2.4 implies via the Stone-Weierstrass theorem that $\mathcal{A}_x$ lies dense $\mathcal{C}_1$. As any bounded linear operator with norm-complete range admits a unique linear extension to the closure of its domain, $(a)|_{\mathcal{A}=\mathcal{A}_x}$ follows. (If the component x^{i_0} of $x = (x^i)$ is strictly monotone, then the interpath $\tilde{x}^{i_0} \equiv x^{i_0} - x_0^{i_0} \in \mathcal{A}_x$, and thus $\mathcal{A}_x$, separates points.)

[13] More accurately, a coordinate vector wrt. the canonical 'eigenbasis' of x that is given by (2.10).

By the next remark, the expected signature (2.16) of a random path in $\mathcal{C}_d$ can be regarded as a natural 'noncommutative' extension to $\mathcal{B}(\mathcal{C}_d)$ of the classical moment-based coordinates (2.17) on $\mathcal{B}(\mathbb{R}^d)$; this justifies calling $\mathfrak{S}(X)$ a 'coordinate vector' for (the law of) X.

Remark 2.2.10 (Motivation)**.** Many results in statistics are based on the well-known fact that the distribution of a random vector $Z = (Z^1, \cdots, Z^d)$ in $\mathbb{R}^d$ can be characterised by a set of coordinates with respect to a basis of nonlinear functionals on $\mathbb{R}^d$. More specifically, any such vector Z can be assigned its 'moment coordinates' $(\mathfrak{m}_{\boldsymbol{i}}(Z))_{\boldsymbol{i}\in[d]^\star} \subset \bar{\mathbb{R}}$ defined by

$$\mathfrak{m}_{i_1\cdots i_m}(Z) \;:=\; \mathbb{E}\left[Z^{i_1}\cdots Z^{i_m}\right] \;=\; \int_{\mathbb{R}^d} x_{i_1}\cdots x_{i_m}\,\mathbb{P}_Z(\mathrm{d}x). \tag{2.17}$$

As the linear span of the monomials $\{x_{\boldsymbol{i}} \equiv x_{i_1}\cdots x_{i_m} \mid \boldsymbol{i} \equiv (i_1, \ldots, i_m) \in [d]^\star\}$ is uniformly dense in the spaces of continuous functions over compact subsets of $\mathbb{R}^d$, the coordinatisation (2.17) is faithful in the sense that, under certain conditions [96], the (coefficients of) the moment vector $\mathfrak{m}(Z) \equiv (m_{\boldsymbol{i}}(Z))_{\boldsymbol{i}\in[d]^\star}$ determine the distribution of Z uniquely.

Now, if instead of a random vector in $\mathbb{R}^d$ one seeks to find a convenient coordinatisation for the distribution of a stochastic process X in $\mathbb{R}^d$, i.e. a random path in $\mathcal{C}_d$, then one can – perhaps surprisingly – resort to a natural generalisation of (2.17), which will be the expected signature (2.16) of X: Analogous to how the monomials $\{x_{\boldsymbol{i}} \mid \boldsymbol{i} \in [d]^\star\}$ are a basis[14] of nonlinear functionals on $\mathbb{R}^d$ that provides coordinates $(\mathfrak{m}_{i_1\cdots i_m} \mid (2.17)) \subset \bar{\mathbb{R}}$ for a random vector in $\mathbb{R}^d$, the signature coordinates $\mathfrak{sig}(x) \equiv (\mathfrak{sig}_{\boldsymbol{i}}(x) \mid \text{as in (2.13)},\ \boldsymbol{i} \in [d]^\star)$ of (regular enough, say of bounded variation) paths $x \in \mathcal{C}_d$ induce nonlinear functionals

$$\mathfrak{sig}_{i_1\cdots i_m} \;:\quad x \quad\longmapsto\quad \mathfrak{sig}_{i_1\cdots i_m}(x) = \int_{0\leq t_1\leq t_2\leq\ldots\leq t_m\leq 1} \mathrm{d}x^{i_1}_{t_1}\mathrm{d}x^{i_2}_{t_2}\cdots\mathrm{d}x^{i_m}_{t_m} \tag{2.18}$$

on applicable subspaces of $\mathcal{C}_d$ which define a dual basis[15] for the vector $x \in \mathcal{C}_d$ ('noncommutative moments'). Thus, these moments yield 'dual coordinates' $(\sigma_{\boldsymbol{i}})_{\boldsymbol{i}\in[d]^\star} \subset \overline{\mathbb{R}}$ for (the law of) a random path X in $\mathcal{C}_d$ via

$$\begin{aligned} \sigma_{i_1\cdots i_m}(X) \;&:=\; \int_{\mathcal{C}_d} \mathfrak{sig}_{i_1\cdots i_m}(x)\,\mathbb{P}_X(\mathrm{d}x) \qquad\qquad (i_1, \ldots, i_m \in [d],\ m \geq 0) \\ &=\; \mathbb{E}\Big[\int_{0\leq t_1\leq t_2\leq\cdots\leq t_m\leq 1} \mathrm{d}X^{i_1}_{t_1}\mathrm{d}X^{i_2}_{t_2}\cdots\mathrm{d}X^{i_m}_{t_m}\Big]. \end{aligned} \tag{2.19}$$

Similarly still to the monomial dual basis $\{x_{i_1}\cdots x_{i_m}\}$ on $\mathbb{R}^d$, the linear span of the above functionals $\{\mathfrak{sig}_{i_1\cdots i_m} \mid (2.18)\}$ is closed under pointwise multiplication and hence forms an algebra over the space of applicable paths in $\mathcal{C}_d$, from which one obtains that their linear span

[14] Cf. the trivial fact the monomials $x_1 = \langle\cdot\,, e_1\rangle, \ldots, x_d = \langle\cdot\,, e_d\rangle$ determine each vector in $\mathbb{R}^d$ uniquely.

[15] In the sense that the coefficients $(\mathfrak{sig}_{i_1\cdots i_m}(x) \mid (2.18)) \equiv \mathfrak{sig}(x)$ determine the path x (essentially) uniquely, as detailed in Section 2.2.2.

is uniformly dense in the space of continuous functions over (certain) compact subsets of $\mathcal{C}_d$ (Stone-Weierstrass), see e.g. [104, Thm. 2.15] (and Lemma 2.3.5 (iv)). Consequently,[16] one can infer in analogy to (2.17) that the dual coefficients (2.19) define a complete (i.e. 'characteristic') set of coordinates for the distribution of a stochastic process $X = (X_t^1, \cdots, X_t^d)_{t \in \mathbb{I}}$ in $\mathbb{R}^d$ that has compact support (and sample paths of bounded variation). ◆

Remark 2.2.11. (i) The above-mentioned assumption of compact support for $\mathbb{P}_X$ is of course much too restrictive on a non-locally compact space like $\mathcal{C}_d$, but under additional decay conditions [33] or by using a normalization [34, Theorem 5.6] it can be shown that the coefficients $(\sigma_{i_1 \cdots i_k}(X) \mid (2.19))$ indeed characterize the distribution of X uniquely even if the compactness assumption is dropped. Extending the definition of (2.13) to paths less regular ('rougher') than of bounded variation is at the centre of the Theory of Rough Paths ([54, 104, 105]).

(ii) The first application of the coordinates (2.19) in statistics was given in [123] for SDE parameter estimation, with more recent applications including the development of non-commutative cumulants [19] and Hurst parameter estimation [47].

2.3 A Coordinate Space for the Laws of Stochastic Processes

To make the sequences $\mathfrak{sig}(x)$ and $\mathfrak{S}(X)$ more convenient to work with, we will embed them into an appropriate space where their information can be analysed more efficiently.

Note for this that in contrast to the elements of an actual vector space basis, the family of interpaths (2.10) is not algebraically independent in the sense that there are nontrivial polynomial relations among the elements $\tilde{x}^w$ of the algebra $\mathcal{C}_1$, see Lemma 2.2.4. This complexity of (2.11) is accounted for by, instead of treating it as a mere vector in $\mathbb{R}^{[d]^\star}$, embedding the sequences $\mathfrak{sig}(x)$ and $\mathfrak{S}(X)$ as elements of (the closure of) the following algebra.

2.3.1 The Space of Formal Power Series

Let $[d]^*$ be the *free monoid* on the alphabet $[d] = \{1, \ldots, d\}$, denoting by $*$ the monoid operation (word concatenation), and identify each multiindex $(i_1, \cdots, i_m) \in [d]^*$ with the *word* $\mathtt{i}_1 \cdots \mathtt{i}_m$ $(\equiv \mathtt{i}_1 * \cdots * \mathtt{i}_m) \in [d]^*$ it defines, with ϵ $(\leftrightarrow \emptyset)$ the empty word.

[16] Recall that by Riesz representation theorem, a (signed) Borel measure on a compact metric space K acts as a continuous linear functional over the space $C(K)$ of continuous functions on K and is hence uniquely determined by its (dual) functional action on a dense subset of $C(K)$.

Definition 2.3.1 (Coordinate Space)**.** For $[d]^*$ as above, define $\mathbb{R}[d]^*$ as the *free associative $\mathbb{R}$-algebra on* $[d]^*$, that is

$$\mathbb{R}[d] := \bigoplus_{w\in[d]^*} \mathbb{R}w \tag{2.20}$$

with multiplication the $\mathbb{R}$-bilinear extension of $*$, for $\oplus$ the external direct sum and $\mathbb{R}w$ the free $\mathbb{R}$-module over $\{w\}$. Denote further by $\mathbb{R}[[d]]$ the $[d]$-adic completion[17] of $\mathbb{R}[d]$, that is

$$\mathbb{R}[[d]] = \{\boldsymbol{t} : [d]^* \to \mathbb{R} \mid \boldsymbol{t} \text{ is a map}\} \equiv \left\{\textstyle\sum_{w\in[d]^*} \boldsymbol{t}(w)\cdot w \;\middle|\; \boldsymbol{t} \in \mathbb{R}[[d]]\right\}; \tag{2.21}$$

i.e., $\mathbb{R}[[d]]$ is the *ring of formal power series* in the variables $\mathtt{1}, \ldots, \mathtt{d}$ over $\mathbb{R}$.

A few basic remarks about the spaces $\mathbb{R}[d]$ and its closure $\mathbb{R}[[d]]$ are in order.

Remark 2.3.2. (i) Clearly, every infinite $[d]^\star$-indexed sequence $a \equiv (a_w)_{w\in[d]^\star}$ with real entries, such as $\mathfrak{sig}(x)$ or $\mathfrak{S}(X)$, can be embedded into $\mathbb{R}[[d]]$ via

$$\boldsymbol{a} \;=\; \textstyle\sum_{w\in[d]^*} \boldsymbol{t}_a(w)\cdot w \;\in \mathbb{R}[[d]] \quad \text{with} \quad \boldsymbol{t}_a(\mathtt{i}_1\cdots\mathtt{i}_m) := a_{i_1\cdots i_m}.$$

(ii) Grouping the (free) summands in (2.20) and (2.21) by their wordlength, we have

$$\mathbb{R}[d] \;=\; \bigoplus_{m=0}^{\infty} V_m \quad \text{and} \quad \mathbb{R}[[d]] = \prod_{m=0}^{\infty} V_m \quad \text{for} \quad V_m := \bigoplus_{w\in[d]^*:\,|w|=m} \mathbb{R}w, \tag{2.22}$$

where the *length* $|w| \in \mathbb{N}_0$ of a word $w \in [d]^*$ is defined as the number of its letters. Since the words in $[d]^*$ are a basis of $\mathbb{R}[d]$, any element $\boldsymbol{a} = (a_w)$ of $\mathbb{R}[d]$ can be uniquely written in the form

$$\boldsymbol{a} \;=\; \textstyle\sum_{w\in[d]^*} a_w\cdot w \;=\; \displaystyle\sum_{m=0}^{\infty}\sum_{i_1,\cdots,i_m\in[d]} a_{i_1\cdots i_k}\cdot \mathtt{i}_1\cdots\mathtt{i}_k \;=:\; \sum_{m=0}^{\infty} \boldsymbol{a}_m, \tag{2.23}$$

where all but finitely many of the coefficients $a_w \in \mathbb{R}$ are zero; the elements of $\mathbb{R}[[d]]$ are of the same form except that for these we may have $a_w \neq 0$ for infinitely many words $w \in [d]^*$. This defines projections $\pi_m : \mathbb{R}[[d]] \to V_m$, $\pi_m(\boldsymbol{a}) := \boldsymbol{a}_m \equiv \sum_{|w|=m} a_w\cdot w$ for each $m \geq 0$; we call $(\boldsymbol{a}_m)_{m\geq 0}$ the *graded representation* of an element $\boldsymbol{a} \in \mathbb{R}[[d]]$ and refer to $\boldsymbol{a}_m$ as its *homogeneous component* of degree m.

(iii) The space $\mathbb{R}[[d]]$ is the topological (wrt. the $[d]$-adic topology, cf. Section 2.3.4) dual of $\mathbb{R}[d]$, in symbols: $\mathbb{R}[[d]] = \mathbb{R}[d]'$, and the space $\mathbb{R}[d]$ is the algebraic dual of $\mathbb{R}[[d]]$, in symbols: $\mathbb{R}[d] = \mathbb{R}[[d]]^\star$. The dual pairing between these spaces is given by

$$\langle\cdot,\cdot\rangle \;:\; \mathbb{R}[[d]]\times\mathbb{R}[d] \;\ni\; (\boldsymbol{a},\boldsymbol{b}) \;\mapsto\; \langle\boldsymbol{a},\boldsymbol{b}\rangle = \textstyle\sum_{w\in[d]^*}\langle\boldsymbol{a},w\rangle\cdot\langle\boldsymbol{b},w\rangle \;\in\; \mathbb{R},$$

with $\langle\cdot,\cdot\rangle : \mathbb{R}[[d]]^{\times 2} \to \bar{\mathbb{R}}$ the (infinite) bilinear extension of $\langle u, v\rangle := \delta_{uv}$, $u, v \in [d]^*$.

[17] See for instance [132, p. 17].

Remark 2.3.3. The coordinate space $\mathbb{R}[d]^*$ is not just an $\mathbb{R}$-vector space but a twofold *bialgebra* (in fact: a *Hopf algebra* with antipode given by the signed order-reversion of words), namely wrt. the two multiplications given by (a) the concatenation product $*$ (the $\mathbb{R}$-bilinear extension of the word-concatenation on $[d]^*$), and (b) the shuffle product $\sqcup\!\sqcup$ from (4.6), see [132, pp. 29 and 31] for details. The bi-algebra structures associated to these two products are an algebraic reflection of the duality between (2.17) and (2.18), cf. also (4.5) and Proposition 4.2.2. ◆

By Remark 2.3.2 (i), we may now organise the information contained in the signature $\mathfrak{sig}(x) = (\mathfrak{sig}_w(x) \mid w \in [d]^\star)$ of a path $x = (x^i) \in \mathcal{BV}$ in form of the formal power series

$$\mathfrak{sig}(x) = \sum\nolimits_{w\in[d]^*} \mathfrak{sig}_w(x) \cdot w = \epsilon + \sum_{i=1}^{d} \Big[\sum\nolimits_{w\in[d]^*} \mathfrak{sig}_{wi}(x) \cdot w\Big] \cdot i \;=\; \epsilon + \sum_{i=1}^{d} \mathfrak{sig}^{(i)}(x) \cdot i$$

where we call $\mathfrak{sig}^{(i)}(x) := \sum_{w\in[d]^*} \mathfrak{sig}_{wi}(x) w \in \mathbb{R}[[d]]$ the *i-th partial signature of x*. By (2.13),

$$x^i(\varphi) \;=\; \left\langle \mathfrak{sig}^{(i)}(x), \boldsymbol{\ell}_\varphi \right\rangle \qquad \text{for every} \;\; \varphi \in \mathcal{A}_x, \tag{2.24}$$

where $\boldsymbol{\ell}_\varphi := \sum_{w\in[d]^*} c_w \cdot w \in \mathbb{R}[d]$ for any fixed tuple[18] $(c_w) \subset \mathbb{R}$ with $\varphi = \sum_{w\in[d]^*} c_w \cdot \tilde{x}^w$. It will likewise be mathematically convenient to also embed the expected signature, i.e. the coordinate sequence $\mathfrak{S}(X) = (\sigma_w(X))$ with entries (2.16), as a formal power series in $\mathtt{1}, \ldots, \mathtt{d}$,

$$\mathfrak{S}(X) \;=\; \sum\nolimits_{w\in[d]^*} \sigma_w(X) \cdot w \;\in\; \mathbb{R}[[d]]. \tag{2.25}$$

It turns out that the expected signature (2.25) is somewhat 'bloated' in that it exhibits a certain level of internal algebraic redundancy. In fact, as we will see next, it turns out that the coordinate vector (2.25) is 'close to an exponential', which allows for the information it contains to be efficiently compressed by a logarithmic 'change of coordinates'.

2.3.2 Group-Like Elements, Exponentials, and the Logarithm

Following (2.25), the expected signature $\mathfrak{S}(X)$ of a process X in $\mathbb{R}^d$ can be seen as a coordinate vector of X wrt. the monomial standard basis $\mathfrak{B} := \{\mathtt{i}_1 \cdots \mathtt{i}_k \mid i_1, \ldots, i_k \in [d],\; k \geq 0\}$ of $\mathbb{R}[d]$. The vector $\mathfrak{S}(X)$ itself, however, is contained in a *nonlinear* subspace of $\mathbb{R}[[d]]$; more specifically, it is in fact 'close to an exponential' in the following sense:

Algebraically, $\mathfrak{S}(X)$ lies in the convex hull of the set

$$\mathcal{R}_d \;:=\; \mathfrak{sig}(\mathcal{BV}) \;=\; \{\, \sum\nolimits_{w\in[d]^*} \mathfrak{sig}_w(x) \cdot w \mid x \in \mathcal{BV}\} \;\subseteq^{19}\; \mathcal{G}_d \tag{2.26}$$

[18] Recall that while not necessarily unique, such a tuple exists by definition (2.11) of the eigenspan.

where $\mathcal{G}_d$ is the subgroup[20] of $(\mathbb{R}_1[[d]], *)$ of *group-like elements*, or *characters*, defined by

$$\mathcal{G}_d := \Big\{ \boldsymbol{t} \in \mathbb{R}_1[[d]] \;\Big|\; \forall\, \boldsymbol{p}, \boldsymbol{q} \in \mathbb{R}[d] \,:\, \langle \boldsymbol{t}, \boldsymbol{p} \sqcup\!\sqcup \boldsymbol{q} \rangle = \langle \boldsymbol{t}, \boldsymbol{p} \rangle \cdot \langle \boldsymbol{t}, \boldsymbol{q} \rangle \Big\}. \tag{2.27}$$

The above definition features the *shuffle product* $\sqcup\!\sqcup : \mathbb{R}[[d]]^{\times 2} \to \mathbb{R}[[d]]$ which is given by infinite bilinear extension (from $([d]^*)^{\times 2}$ to $\mathbb{R}[[d]]^{\times 2}$) of the assignment

$$\sqcup\!\sqcup \;:\quad (\mathtt{i}_1 \cdots \mathtt{i}_k,\, \mathtt{i}_{k+1} \cdots \mathtt{i}_{k+\ell}) \;\longmapsto\; \textstyle\sum_{\tau \in S_{k,\ell}} \mathtt{i}_{\tau(1)} \cdots \mathtt{i}_{\tau(k+\ell)}, \tag{2.28}$$

where the above sum is taken over the family[21] of (k, ℓ)-shuffle permutations

$$S_{k,\ell} := \{\tau \in S_{k+\ell} \mid \tau(1) < \cdots < \tau(k) \;\text{ and }\; \tau(k+1) < \cdots < \tau(k+\ell)\}.$$

As exemplified by (2.27), the (commutative) product (2.28) captures and accounts for the noncommutativity that stems from (products of) the iterated integrals (2.10). In fact, $\mathcal{G}_d$ is a Lie group with Lie algebra[22] (e.g. [132, Thm. 3.2]; cf. (2.23) for notation)

$$\mathcal{L}_d := \Big\{ \boldsymbol{t} \equiv \textstyle\sum_{m \geq 0} \boldsymbol{t}_m \in \mathbb{R}_0[[d]] \;\Big|\; \boldsymbol{t}_m \in \mathcal{L}_d^{[]},\; \forall\, n \geq 0 \Big\},$$

where $\mathcal{L}_d^{[]}$ denotes the *free Lie algebra on* $\mathit{1}, \ldots, \mathit{d}$, i.e. the smallest $\mathbb{R}$-submodule of $\mathbb{R}[d]$ that is closed under the Lie bracket $\mathbb{R}[d]^{\times 2} \ni (\boldsymbol{p}, \boldsymbol{q}) \mapsto [\boldsymbol{p}, \boldsymbol{q}] := \boldsymbol{p} * \boldsymbol{q} - \boldsymbol{q} * \boldsymbol{p}$;[23] the space $\mathbb{R}[d]$ is then precisely the envelopping algebra of $\mathcal{L}_d^{[]}$. This relationship implies the identity

$$\mathcal{G}_d \;=\; \exp(\mathcal{L}_d) \qquad \text{with} \qquad \exp(\boldsymbol{t}) \equiv e^{\boldsymbol{t}} := \sum_{m \geq 0} \frac{1}{m!} \boldsymbol{t}^{*m} \tag{2.29}$$

for $\boldsymbol{t} \in \mathbb{R}_0[[d]]$, where the series converges in the $[d]$-adic topology on $\mathbb{R}[[d]]$ (cf. Sect. 2.3.4).

Now from the inclusion $\mathfrak{S}(X) \subset \mathcal{G}_d$ and (2.29), it is reasonable to expect a 'more parsimonious' (than $\mathfrak{S}(X)$) coordinatisation of a process X wrt. $\mathfrak{B}$ to be given by the $\mathfrak{B}$-coordinates of the *faithful linearisation* $\Phi(\mathfrak{S}(X))$ of (2.25) which is effected by $\Phi(\boldsymbol{t}) := \log(\boldsymbol{t})$ with

$$\log(\boldsymbol{t}) := \sum_{m \geq 1} \frac{(-1)^{m-1}}{m} (\boldsymbol{t} - \epsilon)^{*m} \tag{2.30}$$

for $\boldsymbol{t} \in \mathbb{R}_1[[d]]$, since we have $\Phi \circ \exp = \mathrm{id}_{\mathbb{R}_0[[d]]}$ (cf. Remark 2.3.4). This linearised coordinate description $\Phi(\mathfrak{S}(X))$ of $\mathbb{P}_X$ underlies the concept of signature cumulants, see Section 4.1.

[19] For this inclusion, see e.g. [132, Thm. 3.2 ((i) $\leftrightarrow$ (iii)) and Cor. 3.5].

[20] See e.g. [132, Cor. 3.5], with us denoting $\mathbb{R}_\alpha[[d]] := \big\{ \boldsymbol{t} \in \mathbb{R}[[d]] \;\big|\; \langle \boldsymbol{t}, \epsilon \rangle = \alpha \big\}$ for $\alpha = 0, 1$.

[21] Note that in general, $S_{k,\ell}$ *is not* a subgroup of $S_{k+\ell}$, as for instance both $\sigma := (3\ 2)$ and $\tau := (2\ 4\ 3)$ are elements of $S_{2,2}$, but their composition $\rho := \tau \circ \sigma = (3\ 4)$ is not (due to $\rho(3) = 4 > 3 = \rho(4)$).

[22] The elements of $\mathcal{L}_d$ are called *Lie series*, while the elements of $\mathcal{L}_d^{[]}$ are called *Lie polynomials*.

[23] For a definition of 'free Lie algebra' see e.g. [132, Sect. 0.2 (p. 4)], and note that the given minimality-characterisation of $\mathcal{L}_d^{[]}$ holds by [132, Thm. 0.5 and Thm. 0.4].

2.3.3 A Locally Convex Topology

As of yet, the space $\mathbb{R}[[d]]$ has been regarded simply as an 'algebraic container' for the coordinate sequences (2.13) and (2.25). In order to make the information provided by these vectors amenable to some further mathematical analysis, including considerations of convergence, it will be convenient to endow their coordinate space $\mathbb{R}[[d]]$ with a suitable topology.

Recalling the product structure (2.22) of $V^\infty := \mathbb{R}[[d]]$, convenient topologies on this space can be defined as follows: We identify V^∞ with the tensor algebra $\prod_{m\geq 0}^\infty V_m$ (canonically isomorphic[24] to (2.22)) where $V_0 = \mathbb{R}$ and $V_m = V_1^{\otimes m}$ for $V_1 \equiv (\mathbb{R}^d, |\cdot|_2)$, and endow this space with its natural tensor algebra structure with 1.[25] Denote by $\|\cdot\|_m$ the Euclidean (i.e., $|\cdot|_2$-induced) tensor norm on V_m, and (as in Rem. 2.3.2 (ii)) write $\pi_m : V^\infty \to V_m$ ($\hookrightarrow V$), $\pi_m((v_j)_{j\geq 0}) = v_m$, for the canonical projection of V^∞ onto its m^{th} factor. Let further $V_{[m]} := \prod_{\nu=0}^m V_\nu$ be the truncated tensor algebra, and $\pi_{[m]} := \sum_{\nu=0}^m \pi_\nu$ the truncation map.

Remark 2.3.4 (Truncation)**.** Notice that $V_{[m]}$ comes equipped with a natural algebra structure, namely the one realised as the quotient of V^∞ by the ideal $\prod_{\nu>m} V_\nu$; the map $\pi_{[m]}$ is then the canonical quotient epimorphism. Note in particular that

$$\pi_{[m]}(\log(\boldsymbol{t})) = \sum_{k=1}^m \frac{(-1)^{k-1}}{k}\left[\pi_{[m]}(\boldsymbol{t}-\epsilon)\right]^{\underline{*}k} =: \log_{[m]}(\boldsymbol{t})\,, \tag{2.31}$$

defining a (bijective) polynomial map $\log_{[m]} : V_{(1)} \to V_{[m]}$ ($\hookrightarrow V^\infty$; the space $V_{[m]}$ is embedded as a (closed) linear subspace of V^∞ but not as a subalgebra). In the above, $\underline{*}$ denotes the multiplication in the algebra $V_{[m]}$, i.e.: $\pi_{[m]}(\boldsymbol{t}_1)\,\underline{*}\,\pi_{[m]}(\boldsymbol{t}_2) \overset{\text{def}}{=} \pi_{[m]}(\boldsymbol{t}_1 * \boldsymbol{t}_2)$, for any $\boldsymbol{t}_1, \boldsymbol{t}_2 \in V^\infty$.

Our *topological coordinate space* for (random) paths and their laws is then the space

$$V := \left\{\boldsymbol{t} \in V^\infty \;\middle|\; |||\boldsymbol{t}|||_\lambda := \textstyle\sum_{m\geq 0}\|\pi_m(\boldsymbol{t})\|_m \cdot \lambda^m \;<\; \infty,\ \forall\, \lambda > 0\right\} \tag{2.32}$$

equipped with the locally convex topology induced by the (fundamental) family of norms $(|||\cdot|||_\lambda \mid \lambda > 0)$. (Note that the subspace topology on $V_m \subset V$ coincides with the (Euclidean) topology on $(V_m, \|\cdot\|_m)$.) Compare with [33, Section 2], where the locally m-convex algebra (2.32) was first established for the analysis of signatures and their expectation.

The fact that $\mathfrak{sig}(x) \equiv (\mathfrak{sig}_m(x))_{m\geq 0}$ (recall (2.23)) decays factorially in the degree of its homogeneous components $\mathfrak{sig}_m(x)$, namely (see e.g. [104, Theorem 3.7] for a reference)

$$\left\|\mathfrak{sig}_m(x)\right\|_m \;\leq\; \frac{\|x\|_{1\text{-var}}^m}{m!\cdot c} \qquad \text{for each}\;\; m \geq 0 \tag{2.33}$$

[24] Via $[d]^* \ni \mathtt{i}_1\cdots\mathtt{i}_m \leftrightarrow e_{i_1}\otimes\cdots\otimes e_{i_m} \in V_m$ and $\epsilon \leftrightarrow 1 \in \mathbb{R}$, with $(e_i)_{i\in[d]}$ the standard basis of $\mathbb{R}^d$.

[25] See e.g. [104, Sect. 2.2.1, Rem. 1.24 f.] for details.

with $c := 1 + \sum_{k=3}^{\infty}(\frac{2}{k-2})^2$, implies that $\mathfrak{sig}(x), \mathfrak{S}(Y) \in V$, see also Lemma 2.3.5 below.

For convenience, we will also employ the dilation maps

$$\delta_\lambda \;:\; V \to V, \quad (v_m)_{m\geq 0} \mapsto (\lambda^m \cdot v_m)_{m\geq 0}, \qquad (\lambda > 0)$$

as well as the subspaces $V_{(c)} := \{\boldsymbol{t} \in V \mid \pi_0(\boldsymbol{t}) = c\}$ and $V_{(c)}^{\infty}$ (defined analogously), and recall that the space $\mathcal{BV} = \mathcal{C}_d \cap \mathrm{BV}$ of continuous $\mathbb{R}^d$-valued paths of bounded variation can be endowed with the *p-variation topology* (any $p \geq 1$) defined via the seminorm (2.4).

The next lemma collects basic facts on V, (2.13) and (2.30) that are useful for Section 4.4.

Lemma 2.3.5. *Let V and $\mathfrak{sig}$ and $\log$ be as above, and $\rho > 1$. Then the following holds:*

(i) *the space V is a separable and metrizable Hausdorff space;*

(ii) *the signature transform $x \mapsto \mathfrak{sig}(x)$ defines a map $\mathfrak{sig} : \mathcal{BV} \to V$ which for any $1 \leq p < 2$ is continuous wrt. the p-variation topology on $\mathcal{BV}$;*

(iii) *the signature is invariant under order-preserving time-domain reparametrisations of its arguments, i.e. $\mathfrak{sig}(x) = \mathfrak{sig}(x_\varphi)$ for $x_\varphi \equiv (x_{\varphi(t)})_{t\in\mathbb{J}}$ with $\varphi \in C(\mathbb{J};\mathbb{I})$ strictly monotone;*

(iv) *for each $\varphi \in C(\mathcal{K})$ with $\mathcal{K}/\mathcal{T} \subset \mathcal{BV}$ compact, there is a sequence of index-polynomials $(\boldsymbol{\ell}_j)_{j\in\mathbb{N}}$ in $V^\circ := \bigoplus_{m\geq 0} V_m \subset V^\infty$ such that $\varphi = \lim_{j\to\infty}\langle \mathfrak{sig}(\cdot), \boldsymbol{\ell}_j\rangle$ wrt. $\|\cdot\|_\infty$;*

(v) *the capped logarithm $\log_{[m]} : V_{(1)} \to V$ from (2.31) satisfies $\log_{[m]} = \log_{[m]} \circ \pi_{[m]}$ and is continuous for each $m \geq 0$, and $\log$ from (2.30) maps subsets of $\{\boldsymbol{t} \in V_{(1)} \mid |||\boldsymbol{t} - 1|||_\lambda \leq 1\}$ to subsets of $\{\boldsymbol{\ell} \in V^\infty \mid |||\boldsymbol{\ell}|||_\rho \leq \sum_{m\geq 0}(2\rho/\lambda)^m\}$ for any $\lambda > 2\rho$;*

(vi) *for each $m \geq 0$, it holds that on $|||\cdot|||_\rho$-bounded subsets the projections $\pi_{[m]} : V^\infty \to V^\infty$ converge uniformly wrt. $|||\cdot|||_1$ to the identity operator on V^∞;*

(vii) *for each $\lambda > 0$, we have that $\delta_\lambda \circ \log = \log \circ \delta_\lambda$ and $\delta_\lambda[\mathfrak{sig}(x)] = \mathfrak{sig}(\lambda \cdot x)$, any $x \in \mathcal{BV}$.*

Proof. Statements (i), (ii) and (iii) are well-known, see e.g. [33, Cor. 2.4 and Cor. 5.5] and [66, Thm. 4] (or simply apply the change of variables theorem for direct verification). The approximation property (iv), which is sometimes referred to as the *universality* of the signature, is an immediate consequence of the Stone-Weierstrass theorem and the fact that the set $\{\langle \mathfrak{sig}(\cdot), \boldsymbol{\ell}\rangle \mid \boldsymbol{\ell} \in V^\circ\}$ is a subalgebra of $C(\mathcal{K})$ which contains the constants and separates points (the latter due to [66, Thm. 4] which implies that $\mathfrak{sig} : \mathcal{BV}/\mathcal{T} \to V$ is injective), see e.g. [34, Thm. 5.6. (2)].

(v) : As is immediate from (2.31), the map $\log_{[m]}$ is a polynomial and hence V-valued and continuous, the latter by the fact that (both $\pi_{[m]}$ and) the multiplication $*$ ($\cong$ tensor multiplication $\otimes$; cf. footnote 24) on V is continuous (e.g. [33, Section 3]). The commutativity of $\log_{[m]}$ and $\pi_{[m]}$ is clear again from (2.31). As to the boundedness assertion, let $\lambda > 2\rho$ and denote $B_\lambda := \{\boldsymbol{t} \in V_{(1)} \mid |||\boldsymbol{t} - 1|||_\lambda \leq 1\}$. Then in particular $\sup_{\boldsymbol{t}\in B_\lambda} \|\pi_m(\boldsymbol{t})\|_m \ \leq \ \lambda^{-m}$ for every $m \geq 0$, whence for each $\boldsymbol{t} \in B_\lambda$ and $\boldsymbol{\ell} := \log(\boldsymbol{t})$ we have that, for all $m \geq k \geq 1$,

$$\left\|\pi_m\left[(\boldsymbol{t}-1)^{*k}\right]\right\|_m \ \leq \ \sum_{\substack{m_1+\ldots+m_k=m \\ m_\nu \geq 1}} \|\pi_{m_1}(\boldsymbol{t}) * \cdots * \pi_{m_k}(\boldsymbol{t})\|_m \ \leq \ \binom{m-1}{k-1} \cdot \lambda^{-m}$$

(as the tensor norms $\|\cdot\|_m$ are each submultiplicative), and hence find from (2.30) that

$$\|\pi_m(\boldsymbol{\ell})\|_m \ \leq \ \sum_{k=1}^{m} \frac{1}{k}\binom{m-1}{k-1}\lambda^{-m} \ \leq \ 2^m\lambda^{-m} \qquad \text{for each} \ \ m \geq 1,$$

implying that $\sup_{\ell\in\log(B_\lambda)} |||\boldsymbol{\ell}|||_\rho \leq \sum_{m\geq 1}(2\rho/\lambda)^m < \infty$, as desired.

(vi) : Let $B \subset V^\infty$ be bounded wrt. $|||\cdot|||_\rho$, i.e. suppose that $\beta_\rho := \sup_{\boldsymbol{t}\in B} |||\boldsymbol{t}|||_\rho < \infty$. Then there will be some $0 < q < 1$ together with an index $m_0 \geq 1$ such that

$$\sup_{\boldsymbol{t}\in B} \|\pi_m(\boldsymbol{t})\|_m \ \leq \ q^m \qquad \text{for each} \quad m \geq m_0. \tag{2.34}$$

Indeed: Assuming otherwise that the above does not hold, we for any given $q \in (0,1)$ obtain the existence of a sequence $(\boldsymbol{t}^{(n)})_{n\in\mathbb{N}} \subset B$ with the property that

$$\|\pi_{m_n}(\boldsymbol{t}^{(n)})\|_{m_n} \ > \ q^{m_n} \qquad \text{for each} \quad n \in \mathbb{N}$$

for some strictly increasing sequence $(m_n)_{n\in\mathbb{N}} \subset \mathbb{N}$. But choosing $q > \rho^{-1}$ then implies that

$$\beta_\rho \ \geq \ \sup_{n\in\mathbb{N}} |||\boldsymbol{t}^{(n)}|||_\rho \ \geq \ \sup_{n\in\mathbb{N}} (\rho \cdot q)^{m_n} \ = \ \infty$$

in contradiction to the $|||\cdot|||_\rho$-boundedness of B. Thus (2.34) holds, and with it (by convergence of the geometric series) the claimed uniform $|||\cdot|||_1$-convergence of $(\pi_{[m]})_{m\in\mathbb{N}}$.

(vii) : This is clear by inspection of (2.30) and (2.18), respectively. □

2.3.4 Complementary Remarks

2.3.4.1 On the Series exp and log

The $[d]$-adic topology on $\mathbb{R}[[d]]$ is induced by any of the *ultrametric* distances

$$d_\theta(\boldsymbol{a}, \boldsymbol{b}) := \theta^{\varrho(\boldsymbol{a}-\boldsymbol{b})} \ \text{ for } \ \theta \in (0,1) \text{ and } \varrho(\boldsymbol{c}) := \min\{|w| \mid w \in [d]^* : \langle \boldsymbol{c}, w\rangle \neq 0\} \tag{2.35}$$

and $\varrho(\boldsymbol{0}) := +\infty$ (with $\boldsymbol{0} := 0 \cdot \epsilon$). This topology admits the neighbourhood system

$$\mathcal{U}_{\boldsymbol{a}} \equiv \left(U_L(\boldsymbol{a}) \,\middle|\, L \subset [d]^* \text{ finite}\right), \quad U_L(\boldsymbol{a}) := \{\boldsymbol{b} \in \mathbb{R}[[d]] \mid \langle \boldsymbol{b}, w\rangle = \langle \boldsymbol{a}, w\rangle, \ \forall\, w \in L\},$$

and it is the coarsest topology wrt. which the linear functionals $\mathbb{R}[[d]] \ni \boldsymbol{a} \mapsto \langle \boldsymbol{a}, w\rangle \in \mathbb{R}$ are continuous for each $w \in [d]^*$. See [132, p. 17] for details. As shown next, it is wrt. this topology that the series (2.29) and (2.30) converge.

Let $(\boldsymbol{a}_i)_{i\in\mathbb{N}}$ be a sequence in $\mathbb{R}[[d]]$. The fact that the distance function in (2.35) satisfies the ultrametric inequality directly implies that the series $\sum_{i\geq 0} \boldsymbol{a}_i$ converges (in $\mathbb{R}[[d]]$) if and only if $\boldsymbol{a}_i \to 0$ as $i \to \infty$. Given the above neighbourhood system, this in particular implies that, for any $\boldsymbol{a} \in \mathbb{R}[[d]]$, the series $\sum_{i\geq 0} c_i \cdot \boldsymbol{a}^{*i}$ (any $(c_i)_{i\geq 0} \subset \mathbb{R}$, and $\boldsymbol{a}^{*i} := \boldsymbol{a} * \cdots * \boldsymbol{a}$ (i-many times)) converges if $\varrho(\boldsymbol{a}) \geq 1$ (as then $\varrho(c_i \cdot \boldsymbol{a}^{*i}) \geq i$).

Thus for any $\boldsymbol{a} \in \mathbb{R}_1[[d]]$ (i.e. whose scalar intercept is one) we have convergence in $\mathbb{R}[[d]]$ for all of the following power series: $\sum_{k\geq 0} \boldsymbol{a}^{*k}$ as well as, recall (2.29) and (2.30),

$$\exp(\boldsymbol{a} - \mathbf{1}) \equiv e^{\boldsymbol{a}-\mathbf{1}} = \sum_{m\geq 0} \frac{1}{m!}(\boldsymbol{a} - \mathbf{1})^{*m} \quad \text{and} \quad \log(\boldsymbol{a}) = \sum_{m\geq 1} \frac{(-1)^{m-1}}{m}(\boldsymbol{a} - \mathbf{1})^{*m} \tag{2.36}$$

with $\mathbf{1} := 1 \cdot \epsilon$, where $\exp(\log(\boldsymbol{a})) = \boldsymbol{a}$ and $\log(e^{\boldsymbol{a}-\mathbf{1}}) = \boldsymbol{a} - \mathbf{1}$ by direct calculation.

In particular, an element $\boldsymbol{a} \in \mathbb{R}[[d]]$ is invertible iff $a_0 := \langle \boldsymbol{a}, \epsilon\rangle \neq 0$, and then

$$\boldsymbol{a}^{-1} \left[= a_0^{-1} \cdot \left(\mathbf{1} - \frac{\boldsymbol{a}'}{a_0}\right)^{-1} = a_0^{-1} \cdot \sum_{k\geq 0} \left(\frac{\boldsymbol{a}'}{a_0}\right)^{*k} \right] = \frac{1}{a_0} \sum_{k\geq 0} \left(\mathbf{1} - \frac{\boldsymbol{a}}{a_0}\right)^{*k}$$

for $\boldsymbol{a}' := a_0 \cdot \mathbf{1} - \boldsymbol{a}$. Also, if $\boldsymbol{s} := \sum_{i\geq 0} \boldsymbol{a}_i$ converges then, by the continuity of $\langle \cdot, w\rangle$, $\langle \boldsymbol{s}, w\rangle = \sum_{i\geq 0} \langle \boldsymbol{a}_i, w\rangle \in \mathbb{R}$ for each $w \in [d]^*$. This together with (2.36) yields that $\exp : \mathbb{R}_0[[d]] \to \mathbb{R}_1[[d]]$ and $\log : \mathbb{R}_1[[d]] \to \mathbb{R}_0[[d]]$ are bijections with $\exp^{-1} = \log$.

2.3.4.2 The Geometry of $\mathcal{R}_d$

While the log transform (2.30) yields an efficient compression of the somewhat redundant 'exponential' coordinatisation (2.25) of $\mathbb{P}_X$, we do not currently know if $\kappa(X) \in \mathcal{L}_d$ in general. Certainly this would hold if the signature space $\mathcal{R}_d$ was convex, which seems to be an open question; clearly, at least, the space $\mathcal{R}_d$ is not generally an affine space. Indeed: If the space $\mathcal{R}_d \subset \mathbb{R}[[d]]$ was affine, then (due to $\epsilon \in \mathcal{R}_d$) its translation $\mathcal{R}_d' := \mathcal{R}_d - \epsilon$ would be linear and so would be its m-projections $R_m := \pi_m(\mathcal{R}_d')$ for each $m \geq 0$ (cf. Remark 2.3.2 (ii)). While this holds for $m = 0, 1$, the case $m = 2$ already fails in dimension two: in this case, $R_2 = \{q(z) \equiv \sum_{i,j=1,2} \mathfrak{sig}_{ij}(z) \cdot ij \mid z \equiv (x, y)^\intercal \in \mathcal{BV}_2\}$ with $q_{11}(z) \equiv \mathfrak{sig}_{11}(z) = x_{0,1}^2 \geq 0$ for each $z \in \mathcal{BV}_2$; hence for any $\tilde{z} \in \mathcal{BV}_2$ with $q_{11}(\tilde{z}) > 0$ we have $\lambda \cdot q(\tilde{z}) \notin R_2$ for all $\lambda < 0$, showing that R_2 cannot be a vector space. In fact:[26] since $q_{22}(z) \equiv \mathfrak{sig}_{22}(z) = y_{0,1}^2$ and $q_{12}(z) \equiv \mathfrak{sig}_{12}(z) = \int_0^1 x_{0,t} \,\mathrm{d}y_t$ and $q_{21}(z) \equiv \mathfrak{sig}_{21}(z) = \int_0^1 y_{0,t} \,\mathrm{d}x_t$, we find that $(q_{12} + q_{21})^2 = q_{11}q_{22}$ (via partial integration) with no interdependencies among q_{12}, q_{11}, q_{22},

[26] We thank Terry Lyons for contributing this observation.

and hence that $R_2 \equiv \{(q_{11}(z), q_{22}(z), q_{12}(z), q_{21}(z)) \mid z \in \mathcal{BV}_2\} \subseteq \{(X, Y, A_1, A_2) \in \mathbb{R}^4 \mid (A_1 + A_2)^2 - XY = 0\}$, implying that R_2 is a non-affine three-dimensional *algebraic surface* locally parametrized by the first three components (q_{11}, q_{22}, q_{12}) of its constituent vectors.

Chapter 3

Nonlinear Independent Component Analysis

Motivated by recent breakthroughs of Hyvärinen and Morioka [82, 83], we revisit the inference problem (1.2) for stochastic processes[1] $X = (X_t)$ and $S = (S_t)$ with recent tools from stochastic analysis. We lay some theoretical foundations for the blind inversion (1.3) and provide general identifiability results that generalise Comon's classical independence-based inversion criterion from linear mixtures of random vectors to nonlinear mixtures of discrete- and continuous-time stochastic processes, see Theorems 3.1.1, 3.3.3, and 3.3.7. On a theoretical level, working with infinite-dimensional (i.e. path-valued) random variables poses new challenges that we will address by using rough path theory. From an applied perspective, many models are naturally formulated in continuous time rather than in discrete time (e.g. in biology, physics, medicine or finance), which our approach accounts for by naturally covering both discrete-time and continuous-time models alike, including Stochastic Differential Equations (SDEs) in particular.

This chapter is organised as follows. Our formal exposition towards the recovery of nonlinearly mixed independent sources begins by recalling the main results of [37] as conceptional points of reference (Section 3.1). The core of our identifiability theory is developed in the subsequent three sections: advocating for the incorporation of time as an integral dimension of our source model (Section 3.2), we show how sources admitting a non-degenerate 'temporal structure' harbour sufficient mathematical richness to encode any nonlinear action performed upon them as a sort of 'intrinsic statistical fingerprint', based on which the constituent relation (1.1) may then be inverted up to a minimal deviation by maximizing an independence criterion (Section 3.3). Our approach covers sources of various types of statistical regularity, including popular time series models, various Gaussian processes and

[1] Throughout, "stochastic process" means "continuous stochastic proces" unless mentioned otherwise.

Geometric Brownian Motion (Section 3.5). The chapter concludes with an outlook on non-linear identifiability via control on marginal distributions (Section 3.6) and an explication of how, as promised in the introduction, all results and methods in this chapter are directly applicable to the separation of discrete-time signals as well (Section 3.7).

3.1 Comon's Framework of Linear Independent Component Analysis

Our approach to the problem of nonlinear Blind Source Separation (1.2) for stochastic processes can be regarded as a natural extension of Comon's identifiability framework [37]. This section briefly recalls the main results of this classical framework as conceptional points of reference.

Theorem 3.1.1 (Comon [37, Theorem 11]). *Let $S = (S^1, \cdots, S^d)^\intercal$ be a random vector in $\mathbb{R}^d$ with mutually independent, non-deterministic components $S^1, \ldots, S^d$ of which at most one is Gaussian. Let further $X = C \cdot S$ for an orthogonal matrix $C \in \mathbb{R}^{d\times d}$. Then, for any orthogonal matrix $\theta \in \mathbb{R}^{d\times d}$, we have the following characterisation:*

$$\begin{aligned} (\tilde{S}^1, \cdots, \tilde{S}^d) := \theta \cdot X \;=\; \Lambda P \cdot S \quad &\textit{for some} \quad (\Lambda, P) \in \Delta_d \times \mathrm{P}_d \\ \textit{if and only if} \quad \tilde{S}^1, \ldots, \tilde{S}^d \quad &\textit{are mutually independent.} \end{aligned} \tag{3.1}$$

The significance of Theorem 3.1.1 is that it characterises — up to some minimal deviation, namely their scaling and re-ordering — the independent sources $S^1, \ldots, S^d$ underlying an observable linear mixture $X = A \cdot (S^1, \ldots, S^d)^\intercal$ as precisely those transformations $\theta_\star \cdot X =: (X^1_{\theta_\star}, \ldots, X^d_{\theta_\star})$ of the data whose components $X^i_{\theta_\star}$ are mutually independent.

Remark 3.1.2. (i) The orthogonality constraint of Theorem 3.1.1 imposes no loss of generality with regards to general linear mixtures since any invertible linear relation $X = A \cdot S$, $A \in \mathrm{GL}_d$, between X and S can be reduced to an orthogonal one by performing a principal component analysis on X.

(ii) The proof of Theorem 3.1.1 is based on the remarkable probabilistic fact that any two linear combinations of a family of statistically independent random variables can themselves be statistically independent only if each random variable of this family which has a non-zero coefficient in both of the linear combinations is Gaussian. (A result which is known as the Darmois-Skitovich theorem, see [43, 145].) This accounts for the theorem's somewhat curious 'non-Gaussianity' condition.

(iii) On a historical note, we thank Samuel Cohen for making us aware that the above works are in fact all predated by the earlier identifiability considerations [130] of Reiersøl.

Theorem 3.1.1 enables the recovery of S from X by way of solving an optimisation problem.

Corollary 3.1.3 ([37]). *Let X and S be as in Theorem 3.1.1. Then for any function*[2] $\phi : \mathcal{M}_1(\mathbb{R}^d) \to \mathbb{R}_+$ *such that* $\phi(\mu) = 0$ *iff* $\mu = \mu^1 \otimes \cdots \otimes \mu^d$, *it holds that*[3]

$$\left[\arg\min_{\theta \in \Theta} \phi(\theta \cdot X)\right] \cdot X \;\subseteq\; \mathrm{M}_d \cdot S \tag{3.2}$$

where $\mathrm{M}_d := \{\Lambda \cdot P \mid (\Lambda, P) \in \Delta_d \times \mathrm{P}_d\}$ *is the subgroup of monomial matrices and* $\Theta \subset \mathrm{GL}_d$ *is the subgroup of orthogonal matrices.*

In other words: For f linear and $S = (S^1, \cdots, S^d)^\intercal$ a random vector with mutually independent, non-Gaussian components, the constituent relationship (1.1) between the observable X and its source S can be inverted (up to a minimal deviation) by optimizing some independence criterion ϕ over a set of candidate transformations Θ applied to X.

Partially driven by their applicability (3.2) to ICA, a variety of such criteria ϕ, referred to in [37] as contrast functions, have been developed.

The 'original' independence criterion ϕ_c proposed in [37] quantifies the statistical dependence between the components Y^i of a random vector $Y = (Y^1, \cdots, Y^d)$ in $\mathbb{R}^d$ via the sum of the squares of all standardized cross-cumulants $\kappa^Y_{i_1 \cdots i_j}$ of Y up to r^{th}-order (see [37, Sect. 3.2] and cf. (4.5)), i.e. via the quantity

$$\phi_c(Y) \;:=\; \sum_{j=2}^{r} \sum_{i_1, \ldots, i_j}^{\times} (\kappa^Y_{i_1 \cdots i_j})^2 \qquad (r \geq 2) \tag{3.3}$$

where the inner sum runs over the indices $i_1, \ldots, i_j \in [d]$ corresponding to (4.5).
Initially proposed in [37], the statistic (3.3) originates from a truncated Edgeworth-expansion of mutual information in terms of the standardized cumulants of its argument.

A variety of alternatives to (3.3) soon followed, including kernel-based independence measures [9, 60], a variety of (quasi-) maximum-likelihood objectives, e.g. [13, 112, 126], as well as mutual information and approximations thereof, e.g. [26, 37, 77, 78].

While successfully achieving the separability of linear mixtures, Theorem 3.1.1 has its limitations: Being based on somewhat of a probabilistic curiosity (Rem. 3.1.2 (ii)), it might not be surprising that the characterisation (3.1) cannot be generalised to guarantee the

[2] Here and in the following, $\mathcal{M}_1(V) := \{\mu : \mathcal{B}(V) \to [0,1] \mid \mu \text{ is a (Borel) probability measure}\}$ denotes the space of probability measures over the Borel σ-algebra $\mathcal{B}(V) := \sigma(\mathcal{T})$ of a topological space $(V, \mathcal{T})$.

[3] We write $\mu^i := \mu \circ \pi_i^{-1}$ for the i^{th} marginal of a (Borel) measure μ on $\mathbb{R}^d$. We further abuse notation by writing $\phi(Z) := \phi(\mathbb{P}_Z)$ for any random vector $Z : (\Omega, \mathscr{F}, \mathbb{P}) \to (\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$.

recovery of independent scalar sources from substantially more general nonlinear mixtures of them [84]. Roughly speaking, the reason for this is that for a single random vector in $\mathbb{R}^d$, the statistical property of componental independence is too weak to characterise the nonlinear mixing transformations preserving this property as 'trivial' in a sense made precise by Definition 3.3.1 below. The following example illustrates this.

Example 3.1.4 (Comon's Criterion (3.1) Does Not Apply to Nonlinearly Mixed Vectors in $\mathbb{R}^d$)**.** Let S^1 and S^2 be independent with S^1 Rayleigh-distributed[4] of scale 1 and S^1 uniformly distributed over $(-\pi, \pi)$, and consider the nonlinear mixing transformation f given by $f(u, v) := (u\cos(v), u\sin(v))$ (transformation from polar to Cartesian coordinates). Then even though their functional relation f to S^1, S^2 is 'non-trivial' (i.e. f is not monomial in the sense of Definition 3.3.1), the mixed variables X^1 and X^2 defined by $(X^1, X^2) := f(S^1, S^2)$ are [normally distributed and] statistically independent.[5] ◆

Example 3.1.4 is based on the 'Box-Muller transform', a well-known subroutine from computational statistics. A systematic procedure of constructing 'unidentifiable' nonlinear mixtures of IC random vectors in $\mathbb{R}^d$ is given in [84].

3.2 Modelling Sources as Stochastic Processes

A central direction along which the blind recovery of the source S from its nonlinear mixture X can be controlled is the amount of statistical structure that S carries: If the source S is deterministic, then no additional information is given and a meaningful recovery of S from X is generally impossible, cf. Example 1.0.1. If, on the other hand, the source S were to be described merely as a random vector in $\mathbb{R}^d$, then a recovery of S from X is possible but in general only if X is a linear function of S, cf. [37, 84] and Example 3.1.4. A key insight from [83] is to go for the middle ground (see Remark 3.2.3): if we demand the source S to have a 'non-degenerate temporal structure' and exploit this in a suitable manner, then the recovery of S from even its nonlinear mixtures is possible. To formalize such temporal statistical dependencies requires us to model the source S as a stochastic process. To this end, we use this section to briefly recall foundational notions from stochastic analysis (Section 3.2.1) and provide some basic notions and lemmas (Section 3.2.2) that we will use for our subsequent identifiability results in Section 3.3.

[4] Recall that a real random variable Z is said to be *Rayleigh-distributed* of scale $\lambda > 0$ if Z has a probability density of the form $p_Z(u; \lambda) = u\lambda^{-2}e^{-\frac{u^2}{2\lambda^2}}\mathbb{1}_{[0,\infty)}(u)$.

[5] Note that since the density p_S of (S^1, S^2) reads $p_S(s_1, s_2) = \frac{1}{2\pi}s_2 e^{-s_2^2/2}$, the (joint) density $p_X = (p_S \circ f) \cdot |\det J_f|^{-1}$ of (X^1, X^2) factorizes, implying the independence of X^1 and X^2 as claimed.

3.2.1 Stochastic Processes Interpolate Statistical Extremes

Here and throughout, let $\mathbb{I}$ be a compact interval, $d \in \mathbb{N}$ be some fixed integer, write $\mathcal{C}_d \equiv C(\mathbb{I};\mathbb{R}^d) := \{x : \mathbb{I} \to \mathbb{R}^d \mid \text{the map } \mathbb{I} \ni t \mapsto x(t) =: x_t \text{ is continuous}\}$ for the space of continuous paths in $\mathbb{R}^d$, and let $(\Omega, \mathscr{F}, \mathbb{P})$ denote a fixed probability space.

Definition 3.2.1 (Source Model)**.** We call a *continuous stochastic process in* $\mathbb{R}^d$ any map

$$S : \Omega \to \mathcal{C}_d \quad \text{s.t.} \quad \omega \mapsto S(\omega) \equiv (S_t(\omega))_{t\in\mathbb{I}} \quad \text{is} \quad (\mathscr{F}, \mathcal{B}(\mathcal{C}_d))\text{-measurable}, \tag{3.4}$$

where $\mathcal{B}(\mathcal{C}_d) = \sigma(\pi_t \mid t \in \mathbb{I})$ denotes the Borel σ-algebra on the Banach space $(\mathcal{C}_d, \|\cdot\|_\infty)$. Writing $S_t(\omega) \equiv (S_t^1(\omega), \cdots, S_t^d(\omega))^\intercal \in \mathbb{R}^d$ for each $\omega \in \Omega$, the scalar processes $S^i \equiv (S_t^i)_{t\in\mathbb{I}}$ ($i \in [d]$) are called the *component processes* or the *components* of $S \equiv (S^1, \cdots, S^d)$. We say that a stochastic process $S = (S^1, \cdots, S^d)$ *has independent components*, or that S *is IC*, if its distribution $\mathbb{P}_S := \mathbb{P} \circ S^{-1}$ satisfies the factor-identity[6]

$$\mathbb{P}_{(S^1,\cdots,S^d)} \;=\; \mathbb{P}_{S^1} \otimes \cdots \otimes \mathbb{P}_{S^d}. \tag{3.5}$$

Remark 3.2.2. From a more local perspective, Definition 3.2.1 is equivalent to the description of a continuous stochastic process S as an $\mathbb{I}$-indexed family $S = (S_t)_{t\in\mathbb{I}}$ of random vectors[7] S_t in $\mathbb{R}^d$ such that the map $S(\omega) : \mathbb{I} \ni t \mapsto S_t(\omega) \in \mathbb{R}^d$ is continuous for each $\omega \in \Omega$; e.g. [134, Sect. II.27]. Consequently (cf. also Section 2.1.2), the independence condition (3.5) is equivalent to

$$(S^1_{t_1^{(1)}}, \cdots, S^1_{t_{k_1}^{(1)}}), (S^2_{t_1^{(2)}}, \cdots, S^2_{t_{k_2}^{(2)}}), \cdots, (S^d_{t_1^{(d)}}, \cdots, S^d_{t_{k_d}^{(d)}}) \quad \text{mutually } \mathbb{P}\text{-independent}$$

for any finite selection of time-points $t_1^{(1)}, \ldots, t_{k_1}^{(1)}, \ldots, t_1^{(d)}, \ldots, t_{k_d}^{(d)} \in \mathbb{I}$, $k_1, \ldots, k_d \in \mathbb{N}_0$.

Stochastic processes can be given a prominent role in the BSS-context, namely as natural interpolants between deterministic signals and random vectors. While the first type of signal is the unidentifiable default model for the source in (1.1), the latter is the predominant source model in classical ICA-approaches. More specifically, the following is easy to see.

Proposition 3.2.3 (Stochastic Processes Interpolate Between Extremal Source Models)**.** *Let $S = (S_t)_{t\in\mathbb{I}}$ be a continuous stochastic process in $\mathbb{R}^d$ such that either*

(a) *S_s and S_t are independent for each $s, t \in \tilde{\mathbb{I}}$ with $s \neq t$, or*

(b) *$S_s = S_t$ almost surely for each $s, t \in \tilde{\mathbb{I}}$ with $s \neq t$,*

[6] Strictly speaking, (3.5) reads $\mathbb{P}_{(S^1,\cdots,S^d)} = \big(\mathbb{P}_{S^1} \otimes \cdots \otimes \mathbb{P}_{S^d}\big) \circ \psi^{-1}$, an identity of measures on $\mathcal{B}(\mathcal{C}_d)$, where $\psi : \mathcal{C}_1^{\times d} \to \mathcal{C}_d$ is a canonical isometry defining the Cartesian identification $\mathcal{C}_d \cong \mathcal{C}_1^{\times d}$ (Lemma 2.1.2).

[7] For us every random vector in $\mathbb{R}^d$ is Borel, i.e. $(\mathscr{F}, \mathcal{B}(\mathbb{R}^d))$-measurable.

for some $\tilde{\mathbb{I}} \subset \mathbb{I}$ *dense. Then* S *is either a single path in* $\mathcal{C}_d$ *almost surely (i.e.* S *is deterministic;* 'statistically trivial'*) namely iff (a) holds, or the sample-paths of* S *are constant almost surely (i.e.* S *is a random vector;* 'temporally trivial'*) namely iff (b) holds.*

Proof. Let S be defined on $(\Omega, \mathscr{F}, \mathbb{P})$ and fix $t_0 \in \mathbb{I}$. As the realisations of S are continuous, we have $S_{t_0}(\omega) = \lim_{t \to t_0} S_t(\omega)$ for each $\omega \in \Omega$, whence there is a sequence $\mathfrak{t} \equiv (t_j)_{j \in \mathbb{N}} \subset \mathbb{I}$ with $\lim_{j \to \infty} S_{t_j} = S_{t_0}$ pointwise on Ω. By assumption we can choose $\mathfrak{t} \subset \tilde{\mathbb{I}}$, i.e. such that the elements of $(S_{t_j})_{j \in \mathbb{N}}$ are (a) pairwise independent, or (b) equal almost surely.

Suppose first that (a) holds, and let $(\tilde{S}_j)_{j \in \mathbb{N}}$ be a sequence of independent random variables s.t. $\tilde{S}_j \overset{\mathrm{d}}{=} S_{t_j}$ for each $j \in \mathbb{N}$. Clearly then $(\tilde{S}_j, \tilde{S}_k) \overset{\mathrm{d}}{=} (S_{t_j}, S_{t_k})$ for each $j, k \in \mathbb{N}$, whence both (S_{t_j}) and $(\tilde{S}_j)$ converge to S_{t_0} in probability (recall that $S_{t_j} \to S_{t_0}$ in probability since $S_{t_j} \to S_{t_0}$ a.s.). Consequently, there is a subsequence $(j_\nu)_\nu \subset \mathbb{N}$ such that $\lim_{\nu \to \infty} \tilde{S}_{j_\nu} = S_{t_0}$ almost surely, whence the limit S_{t_0} is constant a.s. by Kolmogorov's zero-one law. In other words, there exists a $\mathbb{P}$-full set $\tilde{\Omega}_{t_0} (\in \mathscr{F}$, assuming $\mathscr{F}$ is $\mathbb{P}$-complete) together with a constant $c_{t_0} \in \mathbb{R}^d$ such that $S_{t_0}(\omega) = c_{t_0}$ for each $\omega \in \tilde{\Omega}_{t_0}$. As $t_0 \in \mathbb{I}$ was arbitrary, we hence find that

$$\forall\, s \in \hat{\mathbb{I}} := \mathbb{I} \cap \mathbb{Q} \;:\; \exists\, (c_s, \tilde{\Omega}_s) \in \mathbb{R}^d \times \mathscr{F} \;:\; S_s(\omega) = c_s \;\text{ for each } \omega \in \tilde{\Omega}_s,$$

which implies that $(S_t(\omega))_{t \in \hat{\mathbb{I}}} = \tilde{\gamma}$ for all $\omega \in \tilde{\Omega} := \bigcap_{r \in \hat{\mathbb{I}}} \tilde{\Omega}_r \in \mathscr{F}$, where $\tilde{\gamma} \equiv (\tilde{\gamma}_s)_{s \in \hat{\mathbb{I}}} := (c_s)_{s \in \hat{\mathbb{I}}}$. But since $\hat{\mathbb{I}} \subset \mathbb{I}$ dense and since the sample paths of S are continuous by assumption, we must have $S(\omega) = \gamma$ for all $\omega \in \tilde{\Omega}$, where $\gamma \in C(\mathbb{I}; \mathbb{R}^d)$ denotes the (unique) continuous extension of the path $\tilde{\gamma}$ from $\hat{\mathbb{I}}$ to $\mathbb{I}$. As $\tilde{\Omega}$ has full measure, the claim follows.

Combining the sample-continuity of S with the fact that the countable intersection of $\mathbb{P}$-full sets has probability one, the claimed consequent of (b) follows immediately. ☐

Proposition 3.2.3 asserts that both deterministic signals (a) as well as random vectors (b) can be seen as degenerate stochastic processes, and that for a given stochastic process $S = (S_t)_{t \in \mathbb{I}}$ this degeneracy manifests on the level of its 2^{nd}-order finite-dimensional distributions, i.e. on

$$\text{the distributions of} \qquad \{(S_s, S_t) \mid (s,t) \in \Delta_2(\mathbb{I})\} \tag{3.6}$$

where the index set $\Delta_2(\mathbb{I}) := \{(s,t) \in \mathbb{I}^{\times 2} \mid s < t\}$ is the (relatively) open 2-simplex on $\mathbb{I} \times \mathbb{I}$. In the following, we refer to (3.6) as the *temporal structure* of a stochastic process $S = (S_t)_{t \in \mathbb{I}}$.

The following is essential: As mentioned above and illustrated in the next section, if the temporal structure of the IC source S in (1.1) is 'degenerate' in the sense of Proposition

3.2.3 (a), (b), then S is unidentifiable from X unless f is of a very specific form, e.g. linear (cf. Theorem 3.1.1). Conversely, we will argue that if the source S has a temporal structure which is 'non-degenerate' (in some specified sense) and satisfies some additional regularity assumptions, then $S = (S^1, \cdots, S^d)$ will be identifiable from even its nonlinear mixtures up to a permutation and monotone scaling of its components S^i (Theorems 3.3.3, 3.3.7, 4.2.3).

3.2.2 Stochastic Processes as Sources: Basic Notions and Assumptions

Recall that the BSS problem (1.2) concerns the recovery of the source S from its image X under some mixing transformation f on $\mathbb{R}^d$. It is thus clear that given X, the map f can be analysed only on that part of its domain that is actually reached by S during the time X is observed. With this in mind, we introduce the 'spatial support' of a stochastic process as the smallest closed subset of $\mathbb{R}^d$ which contains (the trace of) $\mathbb{P}$-almost each sample path of the process.[8]

Definition 3.2.4 (Spatial Support)**.** For $Y = (Y_t)_{t\in\mathbb{I}}$ a (continuous) stochastic process in $\mathbb{R}^d$, the *spatial support* of Y is defined as the set

$$D_Y = \overline{\bigcup_{t\in\mathbb{I}} \operatorname{supp}(Y_t)} \tag{3.7}$$

with $\operatorname{supp}(Y_t) \equiv \operatorname{supp}(\mathbb{P}_{Y_t}) =: D_{Y_t}$ denoting the support of the distribution of Y_t, and where the closure is taken wrt. the Euclidean topology on $\mathbb{R}^d$.

(Readers uncomfortable with (3.7) may for simplicity assume that $D_S = \mathbb{R}^d$ throughout.)

The following elementary properties of the set (3.7) will be useful to us.

Lemma 3.2.5. *Let $Y = (Y_t)_{t\in\mathbb{I}}$ be a stochastic process in $\mathbb{R}^d$ which is continuous with spatial support D_Y. Then the following holds:*

(i) *if $f : D_Y \to \mathbb{R}^d$ is a homeomorphism onto $f(D_Y)$, then $D_{f(Y)} = \overline{f(D_Y)}$;*

(ii) *the traces $\operatorname{tr}(Y(\omega)) := \{Y_t(\omega) \mid t \in \mathbb{I}\}$ are contained in D_Y for $\mathbb{P}$-almost each $\omega \in \Omega$;*

(iii) *for each open subset U of D_Y there is some $t^\star \in \mathbb{I}$ with $\mathbb{P}(Y_{t^\star} \in U) > 0$;*

(iv) *if each random vector Y_t, $t \in \mathbb{I}$, admits a continuous Lebesgue density on $\mathbb{R}^d$, then D_Y is the closure of its interior;*

[8] Analogous to how the support $D_Z := \operatorname{supp}(Z)$ of a random vector Z in $\mathbb{R}^d$ is the smallest closed subset of $\mathbb{R}^d$ within which Z is contained with probability one.

(v) *if each random vector Y_t, $t \in \mathbb{I}$, admits a continuous Lebesgue density υ_t such that $\upsilon^x : \mathbb{I} \ni t \mapsto \upsilon_t(x)$ is continuous for each $x \in D_Y$, we for $\dot{D}_t := \{\upsilon_t > 0\}$ have that the set*

$$\bigcup_{(s,t) \in \Delta_2(\mathbb{I})} \dot{D}_s \cap \dot{D}_t \quad \textit{is} \quad \textit{dense} \quad \textit{in} \quad D_Y.$$

Proof. (i): Recalling the support $\mathrm{supp}(\mu)$ of a Borel-measure $\mu : \mathcal{B}(\mathbb{R}^d) \to [0,1]$ to be defined as the smallest closed subset $C \subseteq \mathbb{R}^d$ having total mass $\mu(C) = 1$, it is easy to see that $\mathrm{supp}(\mu \circ f^{-1}) = \overline{f(\mathrm{supp}(\mu))}$ for any continuous f whose domain includes C_μ. (Indeed: Denoting $C_\nu := \mathrm{supp}(\nu)$ for brevity, we find $\mu(f^{-1}(C_{\mu \circ f^{-1}})) = 1$ and hence $C_\mu \subseteq f^{-1}(C_{\mu \circ f^{-1}})$ and thus $\overline{f(C_\mu)} \subseteq C_{\mu \circ f^{-1}}$ (since both f and the closure operation preserve set inclusion); conversely, for $\tilde{C} := \overline{f(C_\mu)}$ we have $f^{-1}(\tilde{C}) \supseteq C_\mu$ and hence $\mu \circ f^{-1}(\tilde{C}) = 1$ and thus $C_{\mu \circ f^{-1}} \subseteq \tilde{C}$ as desired.) This implies that $D_{f(Y_t)} = \mathrm{supp}(\mathbb{P}_{Y_t} \circ f^{-1}) = \overline{f(D_{Y_t})}$ for each $t \in \mathbb{I}$, from which the assertion $D_{f(Y)} = \overline{f(D_Y)}$ is readily obtained via (3.7) and the continuity of f. (Indeed: On the one hand, $f(D_Y) \subseteq \overline{f(\bigcup_{t \in \mathbb{I}} D_{Y_t})} \subseteq \overline{\bigcup_{t \in \mathbb{I}} D_{f(Y_t)}} = D_{f(Y)}$ where the first inclusion holds by the fact that the continuous image of a closure is contained in the closure of the continuous image and the second inclusion holds by the above-established identity of fixed-time supports; on the other hand, the latter identity also implies that $D_{f(Y)} = \overline{\bigcup_{t \in \mathbb{I}} \overline{f(D_{Y_t})}} = \overline{\bigcup_{t \in \mathbb{I}} f(D_{Y_t})} = \overline{f(\bigcup_{t \in \mathbb{I}} D_{Y_t})} \subseteq \overline{f(D_Y)}$, where once more the last inclusion holds since the map f and the closure operator both preserve set inclusions.)

(ii): By definition of D_{Y_t}, the event $\tilde{\Omega}_t := Y_t^{-1}(D_{Y_t})$ has probability one for each $t \in \mathbb{I}$, and hence so does the countable intersection $\tilde{\Omega} := \bigcap_{t \in \mathbb{I} \cap \mathbb{Q}} \tilde{\Omega}_t$. Now taking any $\tilde{\omega} \in \tilde{\Omega}$, we find by construction that the sample path $\gamma^{\tilde{\omega}} = Y(\tilde{\omega})$ satisfies $\mathrm{im}(\gamma^{\tilde{\omega}}|_{\tilde{\mathbb{I}}}) \subset D_Y$ for $\tilde{\mathbb{I}} := \mathbb{I} \cap \mathbb{Q}$. But since $\tilde{\mathbb{I}}$ is dense and the sample path $\gamma^{\tilde{\omega}}$ is continuous, we obtain $\mathrm{tr}(Y(\tilde{\omega})) = \mathrm{im}(\gamma^{\tilde{\omega}}) \subseteq D_Y$ by the fact that D_Y is closed.

(iii): We proceed by contradiction: If $\mathbb{P}(Y_t \in U) = 0$ for each $t \in \mathbb{I}$, then $\mathbb{P}(Y_t \in U^c) = 1$ and thus (as U^c is closed) $\mathrm{supp}(\mathbb{P}_{Y_t}) \subseteq U^c$ for each $t \in \mathbb{I}$, yielding $D_Y \equiv \overline{\bigcup_{t \in \mathbb{I}} \mathrm{supp}(\mathbb{P}_{Y_t})} \subseteq U^c$ in contradiction to $U \subset D_Y$.

(iv): Recall that in our notation, $D_Y = \overline{\bigcup_{t \in \mathbb{I}} D_{Y_t}}^{|\cdot|}$ for $D_{Y_t} := \mathrm{supp}(Y_t)$. Since Y_t admits a continuous Lebesgue density $\chi_t \in C(\mathbb{R}^d)$ (vanishing identically outside of D_{Y_t}) by assumption, each of the sets $D_{Y_t} \overset{\mathrm{def}}{=} \overline{\chi_t^{-1}(\{0\}^c)}^{|\cdot|}$, $t \in \mathbb{I}$, is the closure of an open set. Denote $D'_Y := \bigcup_{t \in \mathbb{I}} D_{Y_t}$. Then for each $u \in D_Y$ there is a sequence $(u_n)_n \subset D'_Y$ with $\lim_{n \to \infty} u_n = u$, and for each $n \in \mathbb{N}$ there is a sequence $(u_{n,m})_m \subset \mathrm{int}(D_{Y_{t_n}}) \subseteq \mathrm{int}(D_Y)$ (some $t_n \in \mathbb{I}$) with $\lim_{m \to \infty} u_{n,m} = u_n$ by the fact that each element of $\langle D_{Y_t} \rangle_{t \in \mathbb{I}}$ is the closure of its interior. With this, it is easy to see that there is a subsequence $(m_n)_n \subset \mathbb{N}$

such that $\lim_{n\to\infty} u_{n,m_n} = u$, proving $u \in \overline{\mathrm{int}(D_Y)}$ as claimed. The remaining inclusion $\overline{\mathrm{int}(D_Y)} \subseteq D_Y$ is clear as $\mathrm{int}(D_Y) \subset D_Y$ and D_Y is closed.

(v): Suppose that the set $\tilde{D} := \bigcup_{(s,t)\in\Delta_2(\mathbb{I})} \dot{D}_s \cap \dot{D}_t$ is not dense in D_S. Then by (iv) the set $\tilde{D}$ is not dense in the interior of D_S, whence there exists $x_0 \in D_S$ and $\varepsilon > 0$ such that $B_\varepsilon(x_0) \subset D_S \setminus \tilde{D}$. Hence by (iii), there must then be some $t^\star \in \mathbb{I}$ such that $\mathcal{O} := \dot{D}_{t^\star} \cap B_\varepsilon(x_0) \neq \emptyset$. Now since $\mathcal{O} \subset \tilde{D}^c$, we for each $x \in \mathcal{O} \subseteq \dot{D}_{t^\star}$ have that $x \in \dot{D}_{t^\star} \cap (\dot{D}_s^c)$ for all $s \neq t^\star$, the latter implying that $x \in \dot{D}_{t^\star} \setminus \left(\bigcup_{s\neq t^\star} \dot{D}_s\right)$ and hence $\upsilon_{t^\star}(x) > 0$ and $\upsilon_s(x) = 0$ for all $s \neq t\star$, contradicting the continuity of $s \mapsto \upsilon_s(x)$. □

Given the above, we can describe the transformation f that maps S to X via[9] (1.1) as

$$\text{a homeomorphism}^{10} \quad f : D_S \to D_X, \tag{3.8}$$

with the action of f outside of D_S and D_X being irrelevant (and inaccessible) to us.

We now introduce smoothness conditions on the density which we require later on.

Definition 3.2.6. A random vector Z in $\mathbb{R}^n$ will be called C^k*-distributed*, $k \in \mathbb{N}_0$, if its distribution admits a Lebesgue density $\varsigma \in C^k(G)$ for $G := \mathrm{int}(\mathrm{supp}(\varsigma))$; if ς is C^k on some open neighbourhood of $x_0 \in \mathbb{R}^n$, then Z will be called C^k*-distributed around* x_0.

Remark 3.2.7. We recall that for $\vartheta : \mathbb{R}^n \to \mathbb{R}^n$ a C^ℓ-diffeomorphism, $\ell \geq 1$, the classical transformation formula for densities asserts that the image $\tilde{Z} := \vartheta(Z)$ of a C^k-distributed random vector Z with density ς is itself $C^{k\wedge(\ell-1)}$-distributed with density $\tilde{\varsigma}$ given by

$$\tilde{\varsigma} \;=\; (\varsigma \circ \vartheta^{-1}) \cdot |\det J_{\vartheta^{-1}}|. \tag{3.9}$$

The action of the mixing map (3.8) on the source can be profitably captured by imposing the temporal structure (3.6) of S to meet the following analytical regularity condition:

In the following, a stochastic process $Y = (Y_t)_{t\in\mathbb{I}}$ in $\mathbb{R}^d$ will be called

C^k*-regular at* $(s,t) \in \Delta_2(\mathbb{I})$ if the random vector (Y_s, Y_t) is C^k-distributed;

the process Y will be called C^k*-regular at* $((s,t), y_0) \in \Delta_2(\mathbb{I}) \times \mathbb{R}^{2d}$ if the random vector (Y_s, Y_t) is C^k-distributed around $y_0 \in \mathbb{R}^{2d}$ and its density at y_0 is positive.

[9] Recall that $X = f(S)$ means: $X_t = f(S_t)$ for each $t \in \mathbb{I}$.

[10] Note that while the assumption of invertibility of f is canonical, the additionally imposed bi-continuity of the mixing transformation f is a technical condition to ensure that the sample-continuity of the considered processes is preserved under any of the operations that follow.

Remark 3.2.8. Note that if Y is C^k-regular at (s,t), then the boundary of the support of the joint density of (Y_s, Y_t) is a Lebesgue nullset. (A direct consequence of Sard's theorem.)

Remark 3.2.9. Given a C^k-distributed random vector $Z = (Z^1, \cdots, Z^n)$ in $\mathbb{R}^n$ with density ς and some fixed subset $I \subseteq [n]$, say $I = \{i_1, \ldots, i_k\}$ with $k := |I|$, denote by $Z_I := \pi_I(Z) = (Z^{i_1}, \cdots, Z^{i_k})$ the random vector in $\mathbb{R}^k$ which is given by the projection of Z to its I-indexed subcoordinates. Then Z_I is C^k-distributed with Lebesgue density ς_I given by

$$\varsigma_I \;=\; \int_{\mathbb{R}^{n-k}} \varsigma \; \mathrm{d}x_1 \cdots \widehat{\mathrm{d}x_{i_1}} \cdots \widehat{\mathrm{d}x_{i_k}} \cdots \mathrm{d}x_n. \tag{3.10}$$

As an immediate consequence, we have the inclusion

$$\operatorname{supp}(\varsigma) \;\subseteq\; \operatorname{supp}(\varsigma_{[k]}) \times \operatorname{supp}(\varsigma_{[n]\setminus[k]}) \qquad \text{for each} \quad k \in [n]. \tag{3.11}$$

Indeed, setting $C_k := \operatorname{supp}(\varsigma_{[k]})$ and $C'_k := \operatorname{supp}(\varsigma_{[n]\setminus[k]})$, we note that $\mathbb{P}(Z \in C_k \times C'_k) = \mathbb{P}(Z_{[k]} \in C_k,\, Z_{[n]\setminus[k]} \in C'_k) \geq \mathbb{P}(Z_{[k]} \in C_k) + \mathbb{P}(Z_{[n]\setminus[k]} \in C'_k) - 1 = 1$ and hence $C_k \times C'_k \supseteq \operatorname{supp}(\mathbb{P}_Z) = \operatorname{supp}(\varsigma)$, as claimed.

The theory of ICA knows two prominent 'exceptional cases' for which the recovery of an IC random vector S in $\mathbb{R}^d$ from even its linear mixtures X cannot be guaranteed without further assumptions, namely the cases in which

(i) more than one of the components of S is Gaussian (cf. Theorem 3.1.1), or

(ii) the source S is 'statistically trivial' in the sense of Proposition 3.2.3 (a).

As it turns out, a generalised version of these pathologies carries over to the first and more 'static' of our separation principles (Theorem 3.3.3), owing to the fact that certain analytical forms of the joint distributions constituting (3.6) will be 'too simple' to guarantee nonlinear identifiability even for sources whose temporal structure (3.6) is not otherwise degenerate.

Generalising (i) and (ii) from 'spatial' to 'inter-temporal statistics', these exceptional types of joint distributions[11] will be named 'pseudo-Gaussian' and 'separable', respectively:

Definition 3.2.10 (Non-Gaussian, (Regularly) Non-Separable)**.** A function $\varsigma : G \to \mathbb{R}$, $G \subseteq \mathbb{R}^2$ open, is called *pseudo-Gaussian* if there are functions $\varsigma_1, \varsigma_2, \varsigma_3 : \mathbb{R} \to \mathbb{R}$ for which

$$\varsigma(x,y) \;=\; \varsigma_1(x) \cdot \varsigma_2(y) \cdot \exp(\pm\, \varsigma_3(x) \cdot \varsigma_3(y))$$

[11] Distributional pathologies similar to Definition 3.2.10 have been first described in [83]. More specifically, the above notions of (strict) non-separability and pseudo-Gaussianity generalise the notions [83, Def. 1 and Def. 2], respectively, see Section 3.3.4.

holds on all of G; the function ς will be called *separable* if the above holds for $\varsigma_3 \equiv 0$. The function $\varsigma : G \to \mathbb{R}$ will be called *strictly non-Gaussian* if it is such that

$$\varsigma|_{\mathcal{O}} \quad \text{is not pseudo-Gaussian,} \quad \text{for every open subset } \mathcal{O} \text{ of } G\,; \tag{3.12}$$

the property of ς being *strictly non-separable* is declared mutatis mutandis. Furthermore, the function $\varsigma : G \to \mathbb{R}$ will be called *almost everywhere non-Gaussian* if

$$\text{there is a closed nullset} \quad \mathcal{N} \subset G \quad \text{s.t.} \quad \varsigma|_{(G\setminus\mathcal{N})} \quad \text{is strictly non-Gaussian}\,;$$

the notion of ς being *a.e. non-separable* is defined analogously.
Finally, a twice continuously differentiable function $\varsigma\,[= \varsigma(x,y)] : \tilde{U} \times \tilde{U} \to \mathbb{R}_{>0}$, with $\tilde{U} \subseteq \mathbb{R}$ open, will be called *regularly non-separable* if

$$\varsigma \quad \text{is a.e. non-separable} \quad \text{and} \quad (\partial_x \partial_y \log \varsigma)\big|_{\Delta_{\tilde{U}}} \neq 0 \quad \text{a.e. on } \Delta_{\tilde{U}} \tag{3.13}$$

where $\Delta_{\tilde{U}} := \{(x,x) \mid x \in \tilde{U}\}$ denotes the diagonal over $\tilde{U}$.

(Clearly, if ς is [strictly/a.e.] non-Gaussian then it is also [strictly/a.e.] non-separable.)

It will be convenient for us to have an analytical characterisation of these 'pathological' types of densities at hand. Such a characterisation is provided by Lemma 3.2.12 below.

Remark 3.2.11. (i) In light of Lemma 3.2.12 (ii), the assumption of regular non-separability can be regarded as a minimal extension of the above notion of strict non-separability. The necessity for this extension will become clear in Section 3.3.2.

(ii) The log-derivative condition of (3.13) is non-vacuous as there are (strictly) non-separable functions whose mixed log-derivatives vanish on the diagonal. As an example of such a function which is strictly non-separable but not regularly non-separable, consider $\varphi_0 \equiv \varphi_0(x,y)$ given by

$$\varphi_0(x,y) = \begin{cases} \exp\big(-1/(x-y)^2\big), & x < y \\ 0, & x = y \\ \exp\big(-1/(x-y)^2\big), & y < x, \end{cases} \quad \text{and define} \quad \varphi := e^{\varphi_0}.$$

Then clearly $\varphi \in C^2(\tilde{U}^{\times 2}; \mathbb{R}_{>0})$ for $\tilde{U} := (0,1)$, and as the mixed-log-derivatives $\partial_x \partial_y \log(\varphi) = \partial_x \partial_y \varphi_0$ vanish nowhere on the dense subset $\tilde{U}^2 \setminus \Delta_{\tilde{U}} \subset \tilde{U}^{\times 2}$ the function φ is also strictly non-separable on $\tilde{U}^2$ by Lemma 3.2.12 (ii). However, since $(\partial_x \partial_y \varphi_0)|_{\Delta_{\tilde{U}}} = 0$ everywhere on $\tilde{U}$, the function φ is clearly not regularly non-separable. ◆

Let us prepare for the following lemma by declaring what we mean by a symmetric set:

Writing $(u, v) \equiv (u_1, \cdots, u_d, v_1, \cdots, v_d)$ for the coordinates on $\mathbb{R}^{2d} \cong \mathbb{R}^d \times \mathbb{R}^d$, a given subset $A \subseteq \mathbb{R}^{2d}$ will be called *symmetric* if

$$A = \tau(A) \qquad \text{for the transposition} \qquad \tau(u, v) := (v, u).$$

A function $\varphi : G \to \mathbb{R}$, $G \subseteq \mathbb{R}^{2d}$, will be called *symmetric* if $\varphi \circ \tau = \varphi$.

(Since $\tau^2 = \mathrm{id}$, it is clear that A is symmetric iff $\tau(A) \subseteq A$. Also, if $A \subseteq \mathbb{R}^{2d}$ is symmetric then $A \subseteq \pi_{[d]}(A) \times \pi_{[d]}(A)$.)

Lemma 3.2.12. *Let $\zeta : G \to \mathbb{R}_{>0}$, with $G \subseteq \mathbb{R}^2$ open, be twice continuously differentiable. Then the following holds.*

(i) *Provided that G is convex, we have that:*

$$\partial_x \partial_y \log \zeta \equiv 0 \qquad \textit{if and only if} \qquad \zeta \ \textit{ is separable};$$

(ii) *ζ is strictly non-separable if and only if the open set*

$$G' := \{z \in G \mid \partial_x \partial_y \log \zeta(z) \neq 0\} \quad \textit{is a dense subset of } G;$$

(iii) *provided that $\mathcal{O} \subseteq G'$ is symmetric, open and convex, we have that:*

$$[\partial_x \partial_y \log \zeta]\big|_{\mathcal{O}} \ \textit{ is separable and symmetric} \qquad \textit{iff} \qquad \zeta\big|_{\mathcal{O}} \ \textit{ is pseudo-Gaussian.}$$

Proof. We use the global abbreviations $\xi := \partial_x \partial_y \log \zeta$ and $\phi := \log \zeta$.

(i): The ‘if’-direction is clear, so suppose that $\partial_x \partial_y \phi = 0$. Then, as G is convex, $\phi \equiv \phi(x, y) = \phi_1(x) + \phi_2(y)$ and hence $\zeta = \exp(\phi) = \zeta_1(x) \cdot \zeta_2(y)$ for $\zeta_i := \exp(\phi_i)$, as claimed.

(ii): Since ξ is continuous, the set $\{\xi = 0\}$ is closed, whence the set $G' = G \cap \{\xi = 0\}^c$ is open. To see that G' is dense in G, take any $z \in G$ and note that, as G is open, there is some z-centered open ball $B_z \subseteq G$. Since ζ is strictly non-separable, $\zeta|_{B'_z}$ is not separable for any open z-centered sub-ball $B'_z \subseteq B_z$, whence by (i) there must be some $z' \in B'_z$ with $\xi(z') \neq 0$, implying $B'_z \cap G' \neq \emptyset$.

The (contrapositive of the) converse implication in (ii) follows via (i).

(iii): Let $\mathcal{O} \subseteq G'$ be symmetric, open and convex. ($\Leftarrow$) is clear by Def. 3.2.10.

($\Rightarrow$): Suppose that $\tilde{\xi} := \xi|_{\mathcal{O}}$ is separable and symmetric, i.e. assume that

$$\tilde{\xi} \equiv \tilde{\xi}(x, y) = f(x) \cdot g(y) \quad \text{ and } \quad \tilde{\xi} \circ \tau = \tilde{\xi}$$

for some $f, g : \mathcal{O}_1 \to \mathbb{R}$, with $\mathcal{O}_1 := \pi_1(\mathcal{O})$. Then $\tilde{\xi} \equiv \tilde{\xi}(x,y) = \operatorname{sgn}(\tilde{\xi}) \cdot \eta(x) \cdot \eta(y)$ for some function $\eta : \mathcal{O}_1 \to \mathbb{R}$, where $\epsilon := \operatorname{sgn}(\tilde{\xi})$ denotes the sign of $\tilde{\xi}$ (i.e., $\operatorname{sgn}(\tilde{\xi}) = \mathbb{1}_{(0,\infty)}(\tilde{\xi}) - \mathbb{1}_{(-\infty,0)}(\tilde{\xi})$). Indeed, the symmetry of $\tilde{\xi}$ implies that $\tilde{\xi}^2 = \tilde{\xi}(x,y) \cdot \tilde{\xi}(y,x) = \tilde{\eta}(x) \cdot \tilde{\eta}(y)$ for the map $\tilde{\eta}(z) := f(z) \cdot g(z)$; consequently, $\tilde{\xi} = \epsilon \cdot \sqrt{\tilde{\xi}^2} = \epsilon \cdot \eta(x) \cdot \eta(y)$ for the map $\eta \equiv \eta(z) := \sqrt{|\tilde{\eta}(z)|}$. Now since $\mathcal{O}$ is a connected subset of G', the sign of $\tilde{\xi}$ is constant, i.e. $\epsilon = \pm 1$. Integrating $\xi = \partial_x \partial_y \phi$ thus implies that

$$\begin{aligned} \phi &= \int \partial_y \phi \, \mathrm{d}y + f_1(x) = \iint \tilde{\xi} \, \mathrm{d}x + f_2(y) \, \mathrm{d}y + f_1(x) \\ &= \epsilon \cdot f_3(x) \cdot f_3(y) + \tilde{f}_2(y) + \tilde{f}_1(x) \end{aligned} \tag{3.14}$$

for $f_3 \equiv f_3(z) := \int_{z_0}^{z} \eta(s) \, \mathrm{d}s$ (some priorly fixed $z_0 \in \mathcal{O}_1$) and some additional continuous functions $f_i, \tilde{f}_i : \mathcal{O}_1 \to \mathbb{R}$. Note that since $\mathcal{O}$ is convex, the integrated identity of functions (3.14) holds pointwise on all of $\mathcal{O}$. Exponentiating (3.14) now yields the claim. □

3.3 An Identifiability Theorem for Nonlinearly Mixed Independent Sources

We are now ready to present the mathematical core behind our identifiability results for nonlinearly mixed time-dependent sources. Following an overview of our strategy (Section 3.3.1), we state and prove our main results (Sections 3.3.2 and 3.3.3); a comparison with related work (Section 3.3.4) concludes this section.

Throughout, let S and X be two continuous stochastic processes in $\mathbb{R}^d$ that are related via

$$X = f(S) \tag{3.15}$$

for a mixing transformation f which is C^2-invertible on some open superset of D_S.

(Recall that (3.15) reads as the pointwise identity: $X_t = f(S_t)$ for each $t \in \mathbb{I}$.)

Here, we say that $f : \mathbb{R}^d \to \mathbb{R}^d$ is C^k-*invertible* on an open set G of $\mathbb{R}^d$, in symbols: $f \in C^k(G)$, if the restriction $f|_G$ is a C^k-diffeomorphism (with C^k-inverse $f^{-1} : f(G) \to G$).

Throughout the rest of this thesis, we operate under the convenience assumption that every connected component of D_S is convex. This assumption can be dropped immediately at the price of having the last identity in (3.20) hold almost surely with the permutation P depending on the connected component of D_S that a given realisation $\mathfrak{s} \equiv S(\omega)$ of S is [almost surely] contained in *and*, and this is the price, with the scales $h_i \equiv h_i(\mathfrak{s})$ varying with each maximally convex subset of D_S that the trace of $\mathfrak{s}$ passes through; cf. Sect. 3.4.1.

3.3.1 Main Idea and Overview

Starting from (3.15) with the coordinates $(S_t^1)_{t\in\mathbb{I}}, \ldots, (S_t^d)_{t\in\mathbb{I}}$ of the source $S = (S_t^1, \cdots, S_t^d)_{t\in\mathbb{I}}$ assumed mutually independent, we seek to identify S from X by exploiting the main dimensions of our model, space and time, via their statistical synthesis (3.6), the temporal structure of S. This will be done along the following lines (see Section 3.3.2 for details).

Given $(s,t) \in \Delta_2(\mathbb{I})$ we first double the available spatial degrees of freedom by lifting the mixing identity (3.15) to an associated identity in the factor-space $\mathbb{R}^d \times \mathbb{R}^d$, namely

$$(X_s, X_t) = (f \times f)(S_s, S_t) \tag{3.16}$$

(recall page vii for the Cartesian product notation). The lifted mixing identity (3.16), which directly involves the temporal structure (3.6) of the source, now allows for the following statistical comparison in the spirit of [83]:

For X_t^* an independent copy of X_t, consider the intertemporal features $Y := (X_s, X_t)$ and $Y^* := (X_s, X_t^*)$ of the observable X at a fixed (s,t) together with their random combination

$$\bar{Y} := C \cdot Y + (1 - C) \cdot Y^*$$

for an equiprobable $\{0,1\}$-valued random variable C independent of Y, Y^*. Combining (3.16) with the fact that S is IC, we obtain for the (deterministic) functional $L(Y, Y^*) := \psi \circ \rho$ with $\rho(y) := \mathbb{E}[C \mid \bar{Y} = y]$ and $\psi(p) := \log(p/(1-p))$ a contrast identity of the form

$$L(Y, Y^*) = R(f, (S_s, S_t)) \tag{3.17}$$

for a function $R \equiv R(f, (S_s, S_t))$ which depends exclusively on f and the distribution of (S_s, S_t). In other words, (3.17) relates X to S via the source's temporal structure (3.6).

Since the LHS $L \equiv L(Y, Y^*)$ in (3.17) is a function of the (joint) distribution of (Y, Y^*) – and thus of the mixture X – only, we for any alternative pair $(\tilde{f}, \tilde{S})$ with $\tilde{f}(\tilde{S}) = X$ and $\tilde{f} \in C^2$ and $\tilde{S}$ IC analogously obtain that $L(Y, Y^*) = \tilde{R}(\tilde{f}, (\tilde{S}_s, \tilde{S}_t))$ and hence

$$\tilde{R}(\tilde{f}, (\tilde{S}_s, \tilde{S}_t)) = R(f, (S_s, S_t)) \tag{3.18}$$

by (3.17), where again $\tilde{R} \equiv \tilde{R}(\tilde{f}, (\tilde{S}_s, \tilde{S}_t))$ is some function which depends only on $\tilde{f}$ and the distribution of $(\tilde{S}_s, \tilde{S}_t)$. Using the C^2-invertibility of $\tilde{f}$, the IC-properties of both $\tilde{S}$ and S allow us to derive from (3.18) via (3.9) a (deterministic) system of functional equations

$$\Gamma(\varrho, (\tilde{S}_s, \tilde{S}_t), (S_s, S_t)) = 0 \qquad \text{for} \qquad \varrho := (\tilde{f}^{-1} \circ f)\Big|_{D_S} \tag{3.19}$$

which involves the partial derivatives of the 'mixing residual' ϱ and is otherwise completely determined by the distributions of $(\tilde{S}_s, \tilde{S}_t)$ and (S_s, S_t).

The assumed distributional properties of (S_s, S_t), i.e. the temporal structure of S as specified by Definition 3.3.2, together with the required IC-property of $\tilde{S}$ are then sufficient to infer from (3.19) that the residual ϱ must be 'monomial' in the sense of Definition 3.3.1.

In other words, we obtained the following: Given a C^2-invertible map $\tilde{f}$, we have that:

$$(\tilde{S}^1, \cdots, \tilde{S}^d) \equiv \tilde{S} = \tilde{f}^{-1}(X) \; = \; \big[P \circ (h_1 \times \cdots \times h_d)\big](S) \tag{3.20}$$

$$\text{for some } P \in \mathrm{P}_d \text{ and monotone } h_1, \ldots, h_d \quad \textit{if and only if} \tag{3.21}$$

$$\text{the component processes} \quad \tilde{S}^1, \ldots, \tilde{S}^d \quad \text{are mutually independent.}$$

The characterisation (3.21), formulated as Theorem 3.3.3, can thus be read as a natural extension of Comon's classical independence criterion (3.1) to nonlinear mixtures of IC stochastic processes whose temporal structure is sufficiently regular.

Additional source conditions that qualify S for the characterisation (3.21) are obtained by 'unfreezing' the above time pair $(s,t) \in \Delta_2(\mathbb{I})$, see Theorem 3.3.7 in Section 3.3.3.

Analogous to how Comon's criterion (3.1) became practically applicable by way of (3.2), our extended criterion (3.21) is clearly equivalent to the optimisation-based procedure

$$\begin{gathered}\left[\operatorname*{arg\,min}_{\tilde{g}\in\Theta} \phi(\tilde{g}(X))\right] \cdot X \;\subseteq\; \mathrm{DP}_d \cdot S \quad \text{ for any } \quad \phi : \mathcal{M}_1(\mathcal{C}_d) \to \mathbb{R}_+ \\ \text{such that:} \quad \phi(\mu) = 0 \quad \text{iff} \quad \mu = \mu^1 \otimes \cdots \otimes \mu^d,\end{gathered} \tag{3.22}$$

for Θ some 'large enough' family of C^2-invertible candidate transformations, and DP_d a nonlinear analogon of the family of monomial matrices M_d (Def. 3.3.1), see Theorem 4.2.3. Based on a 'moment-like' coordinate description for (the laws of) stochastic processes, we propose an efficiently computable such objective ϕ that generalises Comon's original contrast (3.3) from random vectors to stochastic processes (Section 4.2).

3.3.2 Main Theorem

This section forms the heart of our identifiability theory.

We seek to recover the source $S = (S^1, \cdots, S^d)$ from its nonlinear mixture X in (3.15) up to a minimal deviation, namely a permutation and monotone scaling of its components $S^1, \ldots, S^d$. The following nonlinear analogue of the family of monomial matrices makes this precise.

Definition 3.3.1 (Monomial Transformations)**.** Given a subset G of $\mathbb{R}^d$, a map $\varrho : \mathbb{R}^d \to \mathbb{R}^d$ will be called *monomial on* G if for each connected component $\tilde{G}$ of G we have that

$$\varrho\big|_{\tilde{G}} \;=\; P \circ (h_1 \times \cdots \times h_d) \quad \text{for} \;\; P \in \mathrm{P}_d \;\; \text{and} \;\; h_i \in \mathrm{Diff}^1(\pi_i(\tilde{G})). \tag{3.23}$$

(The above differentiability condition is considered void at isolated points of $\pi_i(\tilde{G})$.) We write $\mathrm{DP}_d(G)$ for the family of all functions on $\mathbb{R}^d$ which are monomial on G.

Accordingly, we say that any two paths $\tilde{x}$ and x in $\mathcal{C}_d$ coincide up to a permutation and monotone scaling of their coordinates, in symbols:

$$\tilde{x} \;\in\; \mathrm{DP}_d \cdot x, \tag{3.24}$$

if $(\tilde{x}_t)_{t\in\mathbb{I}} = (\varrho(x_t))_{t\in\mathbb{I}}$ for some $\varrho \in \mathrm{DP}_d(\mathrm{tr}(x))$, where $\mathrm{tr}(x) \equiv \bigcup_{t\in\mathbb{I}} x_t$ is the *trace* of x.

Definition 3.2.10 describes analytical forms that need to be avoided by 'sufficiently many' of the distributions constituting its temporal structure (3.6) if the source S is to be identifiable from X up to a monomial transformation. Sources for which this is the case will be given the following label of regularity (or 'non-degeneracy').

Definition 3.3.2 (α-Contrastive)**.** A stochastic process $S \equiv (S_t^1, \cdots, S_t^d)_{t\in\mathbb{I}}^\intercal$ in $\mathbb{R}^d$ with spatial support D_S will be called *α-contrastive* if S is IC and there is a collection of time-pairs $\mathcal{P}$ in $\Delta_2(\mathbb{I})$ and an associated collection $(D_\mathfrak{p})_{\mathfrak{p}\in\mathcal{P}}$ of open subsets of $\mathbb{R}^d$ such that

(i) the union $\bigcup_{(s,t)\in\mathcal{P}} D_{(s,t)}$ is dense in D_S, and

(ii) for each $(i,(s,t)) \in [d] \times \mathcal{P}$ it holds that S^i is C^2-regular at (s,t) with density $\zeta^i_{s,t}$, and

$$\begin{aligned} &\zeta^i_{s,t}\Big|_{D^{\times 2}_{(s,t)}} \text{ is regularly non-separable for all } i \in [d], \quad \text{and} \\ &\zeta^i_{s,t}\Big|_{D^{\times 2}_{(s,t)}} \text{ is almost everywhere non-Gaussian for all but at most one } i \in [d], \end{aligned}$$

where the above restrictions of the densities are understood wrt. the abuse of notation $\zeta^i_{s,t}(x) := \zeta^i_{s,t}(x_i, x_{i+d})$ for $x = (x_\nu) \in \mathbb{R}^{2d}$. (For notational convenience, this abuse of notation is kept throughout the following.)

Notice that the conditions in Definition 3.3.2 (ii) reflect the classical pathologies (ii) and (i) from p. 43. Further below we will see how the assumptions of Definition 3.3.2 are linked to related works (Section 3.3.4) and that they are satisfied for a number of popular copula-based time series models (Section 3.5.1).

Theorem 3.3.3. *Let the process S in (3.15) be α-contrastive. Then, for any transformation h which is C^2-invertible on some open superset of D_X, we have with probability one that:*

$$h(X) \;\in\; \mathrm{DP}_d \cdot S \quad \textit{if and only if} \quad h(X) \;\; \textit{has independent components.} \tag{3.25}$$

Proof. The 'only-if'-direction in (3.25) is clear, so we only need to show the converse implication. To this end, we in total prove the slightly stronger assertion that

$$\text{If} \quad h(X) \text{ is IC} \quad \text{and} \quad D \equiv D_{(s,t)} \text{ as in Def. 3.3.2,} \qquad \text{then} \qquad \{J_{h\circ f}(u) \mid u \in D\} \subseteq \mathrm{M}_d\,. \tag{3.26}$$

Given (3.26) (and Definition 3.3.2 (i)), the assertion (3.25) follows by way of Lemma 3.4.1 (ii) and the fact that the trace of almost every realisation of S is contained in a connected component of D_S (Lemma 3.2.5 (ii)) which is convex by assumption.

Let now $(s,t) \in \Delta_2(\mathbb{I})$ (fixed) be as in Definition 3.3.2 (ii), i.e. suppose that $(S_s, S_t) = \pi_{(s,t)}(S)$ admits a (joint) C^2-density $\zeta = \zeta_1 \cdots \zeta_d$ (where $\zeta_i \equiv \zeta^i_{s,t}$) with a support $\bar{D} := \mathrm{supp}(\zeta) \subseteq \mathbb{R}^{2d}$ whose boundary $\partial\bar{D}$ is a Lebesgue nullset (cf. Remark 3.2.8).

Moreover, let X^*_t be a copy of X_t which is independent of (X_s, X_t), and denote

$$Y := (X_s, X_t) \qquad \text{and} \qquad Y^* := (X_s, X^*_t). \tag{3.27}$$

For $C \sim \mathrm{Ber}(1/2)$ and independent of Y and Y^*, consider further

$$\bar{Y} := C \cdot Y + (1 - C) \cdot Y^*$$

(so that $\mathbb{P}_{\bar{Y}} = \frac{1}{2}\mathbb{P}_Y + \frac{1}{2}\mathbb{P}_{Y^*}$) together with the associated regression function

$$\rho \,:\, \mathbb{R}^{2d} \to [0,1] \quad \text{given by} \quad \rho(y) := \mathbb{E}[C \,|\, \bar{Y} = y]. \tag{3.28}$$

The function ρ then satisfies the following central equation.

Lemma 3.3.4. *For μ the probability density of Y, and μ^* the probability density of Y^*,*

$$\psi \circ \rho \;=\; \log\mu - \log\mu^* \qquad \textit{a.e. on} \qquad \tilde{D} := \mathrm{supp}(\mu) \tag{3.29}$$

for the logit-function $\psi(p) := \log(p/(1-p))$.

Lemma 3.3.4 is proved in Section 3.4.2. Recalling now that the components of S are mutually independent, we obtain from the transformation formula for densities (3.9) that for the inverse $g \equiv (g_1, \cdots, g_d) := f^{-1}$ and the density ζ^i_1 of S^i_s, resp. the density ζ^i_2 of S^i_t,

$$\log\mu - \log\mu^* \;=\; \sum_{i=1}^{d} \left[\, \log\zeta_i \circ (g_i \times g_i) - \log\zeta^i_1 \circ g_i(u) - \log\zeta^i_2 \circ g_i(v) \right] \tag{3.30}$$

almost everywhere on $\tilde{D}\,(=(f\times f)(\bar{D}))$. Using (3.29), it follows that

$$\psi\circ\rho \;=\; \sum_{i=1}^{d} P_i\circ(g_i\times g_i) \quad \text{for} \quad P_i := \log\zeta_i - \sum_{\nu=1,2}\log\zeta^i_\nu\circ\pi_\nu. \tag{3.31}$$

Let now $h\equiv(h_1,\cdots,h_d)\in \mathrm{Diff}^2(\mathcal{O}_X)$, for some $\mathcal{O}_X\supseteq D_X$ open, be such that the process $\tilde{S}:=h(X)$ has independent components. Using that the above function $\psi\circ\rho$ depends on the observable X only, we due to $(\tilde{S}_s,\tilde{S}_t)=(h\times h)(X_s,X_t)$ and (3.9) obtain that

$$\psi\circ\rho \;=\; \sum_{i=1}^{d} Q_i\circ(h_i\times h_i) \quad \text{a.e. on } \; \tilde{D} \tag{3.32}$$

analogous to (3.31), where the functions[12] $Q_i\in C^2(\bar{D}')$, $i\in[d]$, are given as

$$Q_i := \log\tilde{\zeta}_i - \sum_{\nu=1,2}\log\tilde{\zeta}^i_\nu\circ\pi_\nu \quad \text{with} \quad \tilde{\zeta}_i := \frac{\mathrm{d}\mathbb{P}_{(\tilde{S}_s{}^i,\tilde{S}_t{}^i)}}{\mathrm{d}(u,v)} \;\text{ and }\; \tilde{\zeta}^i_\nu := \frac{\mathrm{d}\mathbb{P}_{\tilde{S}_{r_\nu}{}^i}}{\mathrm{d}u} \tag{3.33}$$

for $r_1:=s$ and $r_2:=t$, and where $\bar{D}'\subseteq\mathbb{R}^{2d}$ denotes the support of $\tilde{\zeta}\equiv\tilde{\zeta}_1\cdots\tilde{\zeta}_d$.

Note that the Q_i are indeed twice continuously differentiable: By (3.9) we have

$$\tilde{\zeta} = \frac{\mathrm{d}\mathbb{P}_{(\tilde{S}_s,\tilde{S}_t)}}{\mathrm{d}(u,v)} = |\det(J_\phi)|\cdot[\zeta\circ\phi] \;\in\; C^1(\bar{D}')$$

for the C^2-density ζ and for $\phi:=((h\circ f)\times(h\circ f))^{-1}\in\mathrm{Diff}^2(\mathcal{O}_{\tilde{S}}^{\times 2};\mathcal{O}_S^{\times 2})$, with $\mathcal{O}_{\tilde{S}}:=h(\mathcal{O}_S)$; reading off the marginal densities $\tilde{\zeta}_i$, $\tilde{\zeta}^i_\nu$, cf. (3.10), we see that the Jacobians appearing in (3.33) cancel out as they did in (3.30), giving us $Q_i\in C^2(\bar{D}')$ as desired.

Combining the identities (3.31) and (3.32) yields that

$$\sum_{i=1}^{d} Q_i\circ(h_i\times h_i) \;=\; \sum_{i=1}^{d} P_i\circ(g_i\times g_i) \tag{3.34}$$

everywhere on the dense open subset $D_\mu:=\{\mu>0\}$ of $\tilde{D}$.

Therefore, the desired implication (3.26) – and hence the assertion of the theorem (see the initial remarks of this proof) – holds if (3.34) implies that for $\varrho:=h\circ f$ we have

$$\{J_\varrho(u)\mid u\in D\} \;\subseteq\; \mathrm{M}_d \quad \text{for each open } \; D\subseteq D_S \; \text{ as in Def. 3.3.2 (ii),} \tag{3.35}$$

i.e. for any (non-empty) open subset D of $\mathbb{R}^d$ for which $\zeta^i|_{D^{\times 2}}$ is regularly non-separable for all $i\in[d]$, and a.e. non-Gaussian for all but at most one $i\in[d]$. Let any such D be fixed.

The remainder of this proof is aimed at deriving (3.35) from (3.34). To this end, notice that since (3.34) can be equivalently written as

$$Q\circ(h\times h) \;=\; P\circ(g\times g)$$

[12] Note that here, we employ the abuse of notation $Q_i(x)\equiv Q_i(x_i,x_{i+d})$ for $x=(x_\nu)\in\bar{D}'$.

for $Q := \varsigma \circ (Q_1 \times \cdots \times Q_d) \circ \tau$ and $P := \varsigma \circ (P_1 \times \cdots \times P_d) \circ \tau$ with $\varsigma(y_1, \ldots, y_d) := \sum_{i=1}^d y_i$ and $\tau(x_1, \ldots, x_{2d}) := (x_1, x_{d+1}, x_2, x_{d+2}, \ldots, x_d, x_{2d})$, we obtain that (3.34) is equivalent to $Q \circ (\varrho \times \varrho) = P$, i.e. to the ($D_\zeta := \{\zeta > 0\}$–everywhere) identity[13]

$$\sum_{i=1}^{d} Q_i \circ (\varrho_i \times \varrho_i) \;=\; \sum_{i=1}^{d} P_i. \tag{3.36}$$

The above is an identity between two twice-continuously-differentiable functions in the arguments $(u_1, \ldots, u_d, v_1, \ldots, v_d) \in D_\zeta \subseteq \mathbb{R}^{2d}$, so we can apply the cross-derivatives $\partial_{u_j}\partial_{v_k}$ to both sides of (3.36) to arrive at the identities

$$\sum_{i=1}^{d} [q_i \circ (\varrho_i \times \varrho_i)] \cdot \partial_{u_j}\varrho_i \cdot \partial_{v_k}\varrho_i \;=\; \sum_{i=1}^{d} \xi_i \cdot \delta_{ijk} \qquad (j, k \in [d]) \tag{3.37}$$

where the ϱ_i are the components of (3.35) and the functions q_i and ξ_i are given as

$$q_i \;:=\; \partial_{u_i}\partial_{v_i} Q_i \qquad \text{and} \qquad \xi_i \;:=\; \partial_{u_i}\partial_{v_i} P_i = \partial_{u_i}\partial_{v_i} \log \zeta_i, \tag{3.38}$$

respectively. (Note that $\partial_{u_j}\partial_{v_k} R_i = r_i \cdot \delta_{ijk}$ ($(R, r) \in \{(Q, q), (P, \xi)\}$) by the Cartesian product-form of the functions (3.31) and (3.33).) Observe now that the system of equations (3.37) can be equivalently expressed as the congruence relation

$$J_\varrho^\intercal \cdot \Lambda_q \cdot J_\varrho \;=\; \Lambda_\xi \qquad \Big(:\Leftrightarrow\; J_\varrho^\intercal(u) \cdot \Lambda_q(u, v) \cdot J_\varrho(v) = \Lambda_\xi(u, v)\Big)$$

for J_ϱ the Jacobian of ϱ and for Λ_q, Λ_ξ defined as the matrix-valued functions

$$\Lambda_q := \operatorname{diag}_{i=1,\ldots,d}[q_i \circ (\varrho_i \times \varrho_i)] \quad \text{and} \quad \Lambda_\xi := \operatorname{diag}_{i=1,\ldots,d}[\xi_i].$$

Since ϱ is a diffeomorphism over $\bar{D}$, its Jacobian J_ϱ is invertible and hence

$$\Lambda_q \;=\; B_\varrho^\intercal \cdot \Lambda_\xi \cdot B_\varrho \quad \text{on} \quad D_\zeta, \qquad \text{for} \quad B_\varrho := J_\varrho^{-1}. \tag{3.39}$$

Since $B_\varrho = J_{\varrho^{-1}} \circ \varrho$ by the inverse function theorem, the matrix-valued function B_ϱ is clearly continuous. Hence[14] we can apply Lemma 3.4.2 below to from (3.39) and the assumptions of Definition 3.3.2 (ii) obtain as desired that

$$\{J_\varrho(u) \mid u \in D\} \;\subseteq\; \mathrm{M}_d\,. \tag{3.40}$$

Indeed, since the above open set $D \subseteq D_S$ has been chosen such that the (positive) functions $\zeta^i\big|_{D^{\times 2}}$ are regularly non-separable for each $i \in [d]$ and a.e. non-Gaussian for all but at most one $i \in [d]$ (Definition 3.3.2 (ii)), Lemma 3.4.2 is clearly applicable to the system (3.39), providing (3.40) as required. But since the above set D was chosen without further restrictions, (3.40) amounts to (3.35) and hence proves Theorem 3.3.3 as desired. □

The following section extends the above line of argument to additional types of sources.

[13] Once more, we abuse notation by writing $P_i(x) \equiv P_i(x_i, x_{i+d})$ ($x \in D$) for the RHS of (3.36).

[14] Notice that $D^{\times 2} \subset \bar{D} \equiv \operatorname{supp}(\zeta)$ (and hence $D^{\times 2} \subseteq D_\zeta$, as D is open) since $\zeta\big|_{D^{\times 2}} > 0$ a.e. by the fact that $\zeta\big|_{D^{\times 2}}$ is a.e. non-separable (and hence a.e. non-zero in particular).

3.3.3 An Extension to Sources of Alternative Temporal Structures

We can generalise the strategy behind Theorem 3.3.3 by 'unfreezing' its usage of the temporal structure (3.6), that is by allowing the considered time-pairs (s, t) to 'vary more freely' across $\Delta_2(\mathbb{I})$; see Lemma 3.4.3. This qualifies additional source classes for nonlinear identification via the characterisation (3.25). As before, the technical key for this is to make the Jacobian J_ϱ of the mixing residual (cf. (3.39)) serve as change of basis for a source-dependent matrix function with non-degenerate eigenspectrum. The next definition formulates two sufficient conditions for this.

Define $\psi(x, y, z) := x^{-2}yz$, and denote by $\nabla^\times := \{(\lambda_\nu) \in \mathbb{R}^d \mid \exists\, i, j \in [d],\, i \neq j \,:\, \lambda_i = \lambda_j\}$ the set of all vectors in $\mathbb{R}^d$ whose coordinates are not pairwise distinct.

Definition 3.3.5 ($\{\beta, \gamma\}$-Contrastive)**.** A continuous stochastic process $S = (S_t^1, \ldots, S_t^d)_{t\in\mathbb{I}}$ in $\mathbb{R}^d$ with independent components and spatial support D_S will be called

- *β-contrastive* if D_S is the closure of its interior and for any open subset U of D_S there is

$$\begin{aligned} &\text{an open subset } \tilde{U} \text{ of } U \quad \text{and} \quad \mathfrak{p} \equiv (s,t),\, \mathfrak{p}' \in \Delta_2(\mathbb{I}) \quad \text{such that,} \quad \text{for all } i \in [d],\\ &\text{the density } \zeta^i_{s,t} \text{ of } (S^i_s, S^i_t),\ \text{likewise } \zeta^i_{\mathfrak{p}'},\ \text{exists with } \zeta^i_{\mathfrak{p}}, \zeta^i_{\mathfrak{p}'} \in C^2(\tilde{U}^{\times 2}) \ \text{and}\\ &\xi^{i\,|\,\tilde{U}}_{s,t} := \left[\partial_{x_i}\partial_{x_{i+d}} \log \zeta^i_{s,t}\right] \circ \iota_{\tilde{U}} \neq 0 \quad \text{and} \quad \xi^{i\,|\,\tilde{U}}_{\mathfrak{p}'} \neq 0 \quad \text{(a.e.),} \qquad \text{and} \qquad (3.41)\\ &\xi^{i\,|\,\tilde{U}}_{\mathfrak{p}'} \notin \langle \xi^{i\,|\,\tilde{U}}_{\mathfrak{p}} \rangle_{\mathbb{R}} := \{c \cdot \xi^{i\,|\,\tilde{U}}_{\mathfrak{p}} \mid c \in \mathbb{R}\} \qquad (3.42)\end{aligned}$$

  with $\iota_{\tilde{U}} : \tilde{U} \ni u \mapsto (u, u) \in \Delta_{\tilde{U}}$ and both $U, \tilde{U}$ non-empty;[15]

- *γ-contrastive* if there is a dense open subset $\mathcal{U}$ of D_S for which the following holds:

$$\begin{aligned} &\text{for each } u \in \mathcal{U} \quad \text{there exists } (v, \mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2) \in \mathbb{R}^d \times \Delta_2(\mathbb{I})^{\times 3} \quad \text{such that}\\ &S \text{ is } C^2\text{-regular around } (\mathfrak{p}_0, (u, v)), (\mathfrak{p}_1, (u, u)) \text{ and } (\mathfrak{p}_2, (v, v)), \quad \text{and}\\ &\left(\psi(\xi^i_{\mathfrak{p}_0}(u, v), \xi^i_{\mathfrak{p}_1}(u, u), \xi^i_{\mathfrak{p}_2}(v, v))\right)_{i\in[d]} \in (\mathbb{R}^d \setminus \nabla^\times), \qquad (3.43)\end{aligned}$$

  where $\xi^i_{\mathfrak{p}} := \partial_{x_i}\partial_{x_{i+d}} \log \zeta_{\mathfrak{p}}$ is the mixed log-derivatives of the C^2-density $\zeta^i_{\mathfrak{p}}$ of (S^i_s, S^i_t).

We will see that the assumptions of γ-contrastivity are satisfied for a number of popular stochastic processes (Section 3.5.2).

Remark 3.3.6 (Relation Between α-, β- and γ-Contrastive Sources)**.** Notice that every α-contrastive process is also γ-contrastive (for $\mathfrak{p}_0 = \mathfrak{p}_1 = \mathfrak{p}_2$, as the proof of Thm. 3.3.3 shows), while β-contrastivity does not imply, nor is it implied by, either α- or γ-contrastivity.

[15] Here as before, we abuse notation by writing $\zeta^i_{s,t}(x) = \zeta^i_{s,t}(x_i, x_{i+d})$ for $x = (x_\nu) \in \mathbb{R}^{2d}$.

Theorem 3.3.7. *Let the source S in* (3.15) *be β- or γ-contrastive. Then for any map h which is C^2-invertible on an open superset of D_X, we have with probab. one that:*

$$h(X) \in \mathrm{DP}_d \cdot S \quad \textit{if and only if} \quad h(X) \textit{ has independent components.} \tag{3.44}$$

Proof. Let h be C^2-invertible on some open superset of D_X and such that $h(X)$ has independent coordinates; the proof of (3.44) is an extension of the proof of Theorem 3.3.3, so let us adopt the set-up and notation of the latter (as done in Lemma 3.4.3). Recall from there (cf. (3.35)) that (3.44) follows if we can find a dense open subset $\mathcal{D}$ of D_S such that

$$B_\varrho(u) \in \mathrm{M}_d \quad \text{for each } \; u \in \mathcal{D}. \tag{3.45}$$

Let us first suppose that S is β-contrastive. In this case, we consider the set

$$\mathcal{D}_0 := \{u \in \mathrm{int}(D_S) \mid \exists\, (\delta, \mathfrak{p}) \in \mathbb{R}_{>0} \times \Delta_2(\mathbb{I}) \;\text{ satisfying }\; (3.47) \text{ and } (3.48)\} \tag{3.46}$$

with properties (3.47), (3.48) that for a given $(\delta, \mathfrak{p}) \in \mathbb{R}_{>0} \times \Delta_2(\mathbb{I})$ are defined as

$$B_\delta(u) \subset D_S \quad \text{with} \quad \Lambda_{\xi_\mathfrak{p}}\big|_{B_\delta(u)\times B_\delta(u)} \subset \mathrm{GL}_d(\mathbb{R}), \qquad \text{and} \tag{3.47}$$

$$\begin{gathered}\text{there is } \; \mathfrak{p}' \in \Delta_2(\mathbb{I}) \;\text{ s.t. for }\; \mathcal{U} := B_\delta(u) \text{ and each } i \in [d], \;\text{ the diagonal restriction} \\ \xi_{\mathfrak{p}'}^{i|\mathcal{U}} \;\text{ vanishes nowhere and is such that }\; \xi_{\mathfrak{p}'}^{i|\mathcal{U}} \notin \langle \xi_{\mathfrak{p}}^{i|\mathcal{U}} \rangle_\mathbb{R}.\end{gathered} \tag{3.48}$$

Let us show first that $\mathcal{D}_0$ is dense in the interior $D_S^\circ := \mathrm{int}(D_S)$. Indeed: Since D_S° is open, assuming that $\mathcal{D}_0$ is not dense in D_S° implies that there exists $(u_*, r) \in D_S^\circ \times \mathbb{R}_{>0}$ with $B_r(u_*) \subseteq D_S^\circ \setminus \mathcal{D}_0$. Now since S is β-contrastive, there will be some $(\tilde{u}_*, r_1) \in B_r(u_*) \times (0, r)$ with $B_{r_1}(\tilde{u}_*) \subseteq B_r(u_*)$ such that both (3.41) and (3.42) hold everywhere on $\tilde{U} := B_{r_1}(\tilde{u}_*)$ for some $\mathfrak{p}, \tilde{\mathfrak{p}}' \in \Delta_2(\mathbb{I})$; thus also (3.48) holds for $\mathcal{U} = \tilde{U}$ and $(\mathfrak{p}, \mathfrak{p}') := (\tilde{\mathfrak{p}}, \tilde{\mathfrak{p}}')$. And since the functions $\xi_{\tilde{\mathfrak{p}}}^i$ are continuous at $(\tilde{u}_*, \tilde{u}_*)$, there (due to (3.41)) will further be some $r_2 > 0$ such that $\xi_{\tilde{\mathfrak{p}}}^i$ vanishes nowhere on $B_{r_2}(\tilde{u}_*)^{\times 2} \subset D_S^{\times 2}$ for each $i \in [d]$; hence also (3.47) holds for $(\delta, \mathfrak{p}) := (r_2, \tilde{\mathfrak{p}})$. But this yields that both (3.47) and (3.48) hold for $u := \tilde{u}_*$ and $\delta := \min(r_1, r_2)$ and $(\mathfrak{p}, \mathfrak{p}') := (\tilde{\mathfrak{p}}, \tilde{\mathfrak{p}}')$, which implies that $\tilde{u}_* \in D_S^\circ \setminus \mathcal{D}_0$ is an element of $\mathcal{D}_0$.

As this is obviously a contradiction, the set (3.46) must be dense in D_S°.

Now since the interior D_S° is dense in D_S by assumption, the theorem's assertion follows if we can show that (3.45) holds for $\mathcal{D} := D_S^\circ$. But since in turn $\mathcal{D}_0$ is dense in D_S°, we obtain that (3.45) holds for $\mathcal{D} := D_S^\circ$ if we can show that (3.45) holds for $\mathcal{D} := \mathcal{D}_0$.[16]

Let to this end $u \in \mathcal{D}_0$ be fixed with $\mathfrak{p}, \mathfrak{p}' \in \Delta_2(\mathbb{I})$ and $\mathcal{U} \equiv B_\delta(u) \subset D_S$ as in (3.47) and (3.48). Then by Lemma 3.4.3 we have that

$$B_\varrho(\tilde{u})^{-1} \cdot \bar{\Lambda}_\nu(\tilde{u}, \tilde{v}) \cdot B_\varrho(\tilde{u}) \;=\; \tilde{\Lambda}_\nu(\tilde{u}, \tilde{v}) \qquad \text{for each } \; (\tilde{u}, \tilde{v}) \in \mathcal{U} \times \mathcal{U} \tag{3.49}$$

[16] Cf. (3.69) and the argument around (3.64) for details.

with $\nu = 1, 2$ and diagonal matrices $\bar{\Lambda}_1, \bar{\Lambda}_2, \tilde{\Lambda}_1, \tilde{\Lambda}_2 \in \mathrm{GL}_d(\mathbb{R})$ given by

$$\bar{\Lambda}_1 := \bar{\Lambda}_{\mathfrak{p},\mathfrak{p},\mathfrak{p}}, \quad \tilde{\Lambda}_1 := \tilde{\Lambda}_{\mathfrak{p},\mathfrak{p},\mathfrak{p}} \qquad \text{and} \qquad \bar{\Lambda}_2 := \bar{\Lambda}_{\mathfrak{p},\mathfrak{p},\mathfrak{p}'}, \quad \tilde{\Lambda}_2 := \tilde{\Lambda}_{\mathfrak{p},\mathfrak{p},\mathfrak{p}'}$$

with matrices $\bar{\Lambda}_{\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2}$ and $\bar{\Lambda}_{\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2}$ as defined in (3.80). Combining the cases $\nu = 1$ and $\nu = 2$ of (3.49), we for any $C \in \mathbb{R}$ obtain that

$$B_\varrho(\tilde{u})^{-1} \cdot \Big[\bar{\Lambda}_1(\tilde{u}, \tilde{v}) + C \cdot \bar{\Lambda}_2(\tilde{u}, \tilde{v})\Big] \cdot B_\varrho(\tilde{u}) \;=\; \tilde{\Lambda}_1(\tilde{u}, \tilde{v}) + C \cdot \tilde{\Lambda}_2(\tilde{u}, \tilde{v}) \tag{3.50}$$

for each $(\tilde{u}, \tilde{v}) \in \mathcal{U} \times \mathcal{U}$. Hence (and since $u \in \mathcal{D}_0$ was chosen arbitrarily), the identity (3.50) implies (3.45) if there is a pair $(C, \tilde{v}) \in \mathbb{R} \times \mathcal{U}$ for which the diagonal entries of $[\bar{\Lambda}_1(u, \tilde{v}) + C \cdot \bar{\Lambda}_2(u, \tilde{v})] =: \mathrm{diag}[\lambda^1_{u,\tilde{v}}, \cdots, \lambda^d_{u,\tilde{v}}]$ are pairwise distinct. We now prove this, i.e. we show that

$$\begin{gathered}\text{there is } \; C \in \mathbb{R} \; \text{ for which we can find some } \; \tilde{v} \in \mathcal{U} \; \text{ s.t. the diagonal entries} \\ \lambda^i_{u,\tilde{v}} = \frac{\xi^i_{\mathfrak{p}}(u,u)}{\xi^i_{\mathfrak{p}}(u,\tilde{v})^2} \cdot \left(\xi^i_{\mathfrak{p}}(\tilde{v},\tilde{v}) + C \cdot \xi^i_{\mathfrak{p}'}(\tilde{v},\tilde{v})\right), \; i \in [d], \;\; \text{are pw. distinct.}\end{gathered} \tag{3.51}$$

Notice that, as detailed in the proof of Lemma 3.4.2, the fact that by construction each of the functions $q_i : \mathcal{U} \times \mathcal{U} \ni (\tilde{u}, \tilde{v}) \mapsto \lambda^i_{\tilde{u},\tilde{v}}$ ($i \in [d]$) are continuous implies that (3.51) holds if

$$\exists\, C \in \mathbb{R} \;\text{ such that }\; \vartheta_i \,:\, \mathcal{U} \ni \tilde{v} \,\mapsto\, q_i(u, \tilde{v}) \;\text{ is non-constant for each }\; i \in [d]. \tag{3.52}$$

To prove (3.52), notice that since for each $i \in [d]$ we have the decomposition

$$\vartheta_i = \theta_i + C \cdot \theta'_i \qquad \text{with} \qquad \theta_i(\tilde{v}) := \frac{\xi^i_{\mathfrak{p}}(u,u)\xi^i_{\mathfrak{p}}(\tilde{v},\tilde{v})}{\xi^i_{\mathfrak{p}}(u,\tilde{v})^2}$$

and $\theta'_i : \mathcal{U} \times \mathcal{U} \to \mathbb{R}$ defined likewise but with the right factor in the above enumerator replaced by $\xi^i_{\mathfrak{p}'}(\tilde{v}, \tilde{v})$, we find that if, for each $i \in [d]$, the functions

$$\theta_i \;\text{ or }\; \theta'_i \;\text{ are non-constant,} \qquad \text{then} \qquad (3.52) \;\text{ holds.} \tag{3.53}$$

Indeed: If *either* θ_i *or* θ'_i is non-constant in $\tilde{v}$, then clearly their linear combination ϑ_i will be non-constant in $\tilde{v}$ for each $C \neq 0$. If θ_i or θ'_i are *both* non-constant in $\tilde{v}$, then there might be some $C_i \in \mathbb{R}$ such that $\theta_i + C_i \cdot \theta'_i$ is constant in $\tilde{v}$ (define $C_i := 1$ otherwise); in this case, setting $C := \max_{i \in [d]} C_i + 1$ implies that $\vartheta_i = \theta_i + C \cdot \theta'_i = (\theta_i + C_i \cdot \theta'_i) + (C - C_i) \cdot \theta'_i$ is non-constant in $\tilde{v}$ for each $i \in [d]$, as desired.

To see that the premise of (3.53) holds, assume otherwise that there is $i \in [d]$ for which the function $\theta_i : \mathcal{U} \to \mathbb{R}$ is constant in $\tilde{v}$, say

$$\theta_i(\tilde{v}) =: \varsigma_i \qquad \text{for all } \; \tilde{v} \in \mathcal{U}.$$

Then, as θ_i vanishes nowhere in consequence of (3.47), we find that

$$\left[\xi_{\mathfrak{p}}^i(u,\cdot)\right]^2 = \frac{\xi_{\mathfrak{p}}^i(u,u)\cdot\xi_{\mathfrak{p}}^{i|\mathcal{U}}}{\theta_i} = c_i\cdot\eta \qquad \text{on} \quad \mathcal{U} \tag{3.54}$$

for the constant $c_i := \xi_{\mathfrak{p}}^i(u,u)\cdot\varsigma_i^{-1}$ and the function $\eta:\mathcal{U}\to\mathbb{R}$ given by $\eta(\tilde{v}) := \xi_{\mathfrak{p}}^i(\tilde{v},\tilde{v})$. Now if the function θ_i' were constant as well, say $\theta_i' \equiv \varsigma_i' (\neq 0)$, then we would likewise obtain that $[\xi_{\mathfrak{p}}^i(u,\cdot)]^2 = c_i'\cdot\eta'$ on $\mathcal{U}$, for the non-zero constant $c_i' := \xi_{\mathfrak{p}}^i(u,u)\cdot(\varsigma_i')^{-1}$ and the function $\eta':\mathcal{U}\to\mathbb{R}$ given by $\eta'(\tilde{v}) := \xi_{\mathfrak{p}'}^i(\tilde{v},\tilde{v})$. Combined with (3.54), we find that

$$c_i'\cdot\eta' = c_i\cdot\eta \quad \text{ and hence } \quad \xi_{\mathfrak{p}'}^{i|\mathcal{U}} = \text{const.}\cdot\xi_{\mathfrak{p}}^{i|\mathcal{U}},$$

the latter contradicting (3.48). This proves the premise of (3.53) and hence (3.45) for $\mathcal{D} = \mathrm{int}(D_S)$.

Suppose next that S is γ-contrastive. In this case, we claim that (3.45) holds for the dense subset $\mathcal{D}\equiv\mathcal{U}$ of D_S postulated by Def. 3.3.5. To see that this is true, fix any $u\in\mathcal{U}$ and recall that, by Lemma 3.4.3 and the previous discussions, the Jacobian $B_\varrho(u)$ of $\varrho\equiv h\circ f$ at u is monomial if there is $v\in\mathbb{R}^d$ and $\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2\in\Delta_2(\mathbb{I})$ with $(u,v)\in\{\xi_{\mathfrak{p}_0}^1\neq 0,\ldots,\xi_{\mathfrak{p}_0}^d\neq 0\}$ for which the diagonal matrix $\bar{\Lambda}_{u,v}\equiv\bar{\Lambda}_{\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2}(u,v)$ given by (3.80) has pairwise distinct eigenvalues. (Note that $(u,v)\in\{\xi_{\mathfrak{p}_0}^1\neq 0,\ldots,\xi_{\mathfrak{p}_0}^d\neq 0\}$ if $\{\psi(\xi_{\mathfrak{p}_0}^i(u,v),\alpha,\beta)\mid i\in[d]\}\subset\mathbb{R}$ for some $\alpha,\beta\in\mathbb{R}$.) Since the diagonal of $\bar{\Lambda}_{u,v}$ equals the vector $(\psi(\xi_{\mathfrak{p}_0}^i(u,v),\xi_{\mathfrak{p}_1}^i(u,u),\xi_{\mathfrak{p}_2}^i(v,v))_{i\in[d]}$, choosing $(v,\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2)$ as in (3.3.5), (3.43) thus yields $B_\varrho(u)\in\mathrm{M}_d$ as claimed. As $u\in\mathcal{U}$ was arbitrary, we obtain $B_\varrho\big|_{\mathcal{U}}\subset\mathrm{M}_d$ as desired in (3.45). This shows (3.44) for γ-contrastive sources. □

3.3.4 Related Work

We remark that the above assumption of α-contrastivity is strictly weaker than the earlier identifiability conditions given in [83] which served us as motivation. Indeed: the latter are defined for densities on $G=\mathbb{R}^2$ only, and if such a density ς is "uniformly dependent" in the sense of [83, Def. 1] then ς is also strictly (and regularly) non-separable on $G=\mathbb{R}^2$ by Lemma 3.2.12 (ii); as to [83]'s complementary notion of ς being "quasi-Gaussian" [83, Def. 2], we thank one of our referees for drawing our attention to the fact[17] that, as stated in loc. cit., the corresponding non-separability condition [83, Thm. 1, Assmpt. 3.] fails to ensure the validity of [83, Thm. 1], see Remark 3.3.8; this deficiency can be remedied, however, if one weakens the excluding notion of quasi-Gaussianity [83, Def. 2] by imposing its defining factorisation condition [83, Def. 2 eq. (4)] to hold merely on some open subset of $\mathbb{R}^2$ instead of globally on all of $\mathbb{R}^2$ (cf. Remark 3.3.8), as is done – upon logical negation

[17] The insufficiency of [83, Thm. 1, Assmpt. 3.] (for [83, Def. 2 eq. (4)] as stated) was also conjectured in [65, (end of) Section 4.3]; we prove this conjecture true in Remark 3.3.8below.

– in Definition 3.2.10, eq. (3.12), and also in the later work [65, Theorem 2]. With [83, Def. 2] thus weakened, the (thus strengthened) identifiability condition [83, Thm. 1, Assmpt. 3.] then becomes a special case of our assumption of pseudo-Gaussianity (Definition 3.2.10) by Lemma 3.2.12 (iii). Consequently, if a source $S \equiv (S^i)$ satisfies [83, Hypotheses 1., 2. & 3. of Theorem 1] – that is if S is stationary and C^2-regular at some point $(s_0, t_0) \equiv (t-1, t)$ with $D_{(s_0,t_0)} = \mathbb{R}^{2d}\,(= D_S)$ such that the densities $\varsigma^i \equiv \varsigma^i_{s_0,t_0}$ of S^i are all uniformly dependent (hence all regularly non-separable) and none quasi-Gaussian in the above, corrected sense (cf. also [65, Assmpt. B2]) – then S is clearly α-contrastive in particular, and the converse is clearly not true in general.

When contrasted with the few prior works in the area that allow for a theoretical comparison, most notably [82, 83], we see that our approach provides a strict generalisation of previously attained results, see above, or yields stronger conclusions while operating under assumptions which are much less restrictive; for example, we do not require the source to belong to a predefined distributional family as, e.g., in [82].

With regards to methodology, we recall that [83] propose to estimate the demixing nonlinearity by training a universal approximator (typically a neural network) to distinguish between vectors excerpting originally-ordered data and vectors excerpting data whose initial sequential order has undergone a random permutation. By implementing this classification task via logistic regression, an approximation of the demixing transformation is then obtained as an optimally trained configuration of the classifying universal approximator provided that the employed regression function is of a certain composite functional form.

In contrast, our approach approximates the demixing nonlinearity more directly via a dependence minimisation task in the classical spirit of Comon [37], which we propose to perform by optimising an explicitly defined, universally applicable contrast function derived from novel signature-based statistics for multidimensional stochastic processes (Chapter 4). Not only is our method thus guaranteed to work under much weaker assumptions than [83] — see the above discussion and the facts that our method is fully applicable to the (non-stationary) discrete- and continuous-time case and free of assumptions on (the functional form of) any approximating auxiliary nonlinearities; its equivalence to a simple-to-formulate optimisation problem also makes our method straightforward to implement and more directly accessible to a theoretical analysis of its statistical properties, cf. Sections 4.4 and 4.8.

We also note that a slightly weaker technical modification of our assumptions (i) & (ii) from Definition 3.3.2 is given and used in the later work [65, Theorem 2], where the problem of nonlinear blind source separation is studied in the presence of independent additive noise.

To the best of our knowledge, our notions of β- or γ-constrastivity (Definition 3.3.5 (and 3.7.1)) bear no evident resemblance to the conditions proposed in this or other works.

Remark 3.3.8 (An Error in the Proof of [83, Theorem 1])**.** The (fixable) error occurs in the proof of [83, Lemma 2 (Supplement)] (we use their notation for the following): Lemma 2 in [83] requires that for each $\bar{\mathbf{u}}^1 \in \mathbb{R}^d$ there exists an $\bar{\mathbf{u}}^2 \in \mathbb{R}^d$ such that the diagonal entries $\psi_i(\bar{u}_i^1, \bar{u}_i^2) \equiv \psi_i(a, b)$ ($i = 1, \ldots, d$; cf. [83, Eq. (42) (Supplement)]) of the matrix $\mathbf{D}_{11}\mathbf{D}_{12}^{-2}\mathbf{D}_{22}(\bar{\mathbf{u}}^1, \bar{\mathbf{u}}^2)$ are pairwise distinct. For this, [83] rely on (the continuity of the ψ_i and) an indirect proof – via contradiction to the exclusion of [83, Def. 2 eq. (4)] – of the assertion that (A) : "The function $\psi_i(a, \cdot)$ is non-constant for [almost] every given $a \in \mathbb{R}$." Yet instead of leading the negation of (A) – i.e., $(\neg A)$: "The function $\psi_i(a, \cdot)$ is constant for some $a \in \mathbb{R}$" – to a contradiction, their proof by contradiction departs from the *stronger* assumption that (B) : "The function $\psi_i(a, \cdot)$ is constant *for each* $a \in \mathbb{R}$" (which implies, but is not generally implied by, $(\neg A)$). The reductio ad absurdum of (B) provided in [83] is therefore insufficient to prove (A). The (negation [83, Thm. 1 Assmpt. 3.] of the) given factorisation property [83, Def. 2 eq. (4)], on the other hand, is *too weak* to contradict $(\neg A)$ as would be required for the (indirect) proof of (A). To see this, it suffices to note the existence of functions $q : \mathbb{R}^2 \to \mathbb{R}_\times$ which are not of the (excluded) global product form [83, Def. 2 eq. (4)] but factorize locally on a given 'strip' $R := I \times \mathbb{R}$, for some $I \subset \mathbb{R}$. Indeed: Given such a function[18] q, the assumption of $(\neg A)$ (in the indirect proof of (A)), with ψ_i and q related by [83, eq. (42)] and $\psi_i(a, \cdot)$ constant iff $a \in I$ (with $I \neq \mathbb{R}$ in general), leads to the R-*local* identity

$$q(a, b) = c\alpha(a)\alpha(b) \quad \text{for} \quad (a, b) \in I \times \mathbb{R} \equiv R \qquad \text{(cf. [83, derivation of (45)]).} \tag{3.55}$$

Now since $(\neg A)$ does not imply $R \neq \mathbb{R}^2$ in general, the identity (3.55) does generally *not* contradict the (global) non-factorizability assumption [83, Thm. 1 Assumption 3.], leaving (A) – thus [83, Lemma 2] and hence [83, Theorem 1] – unproved. Clearly however, as mentioned in Remark 3.3.4, the desired contradiction can be re-obtained by allowing for only such q for which the factorisation property (3.55) does not hold on any open subset R in $\mathbb{R}^2$.

Suitable functions q as required above can be easily constructed via cut-off functions. Indeed: Assuming for simplicity that I is compact, take any (continuous) $\alpha : \mathbb{R} \to \mathbb{R}_\times$ and let $\chi : \mathbb{R}^2 \to (0, 1]$ be a smooth function with $\chi|_{\tilde{R}} \equiv 1$ for $\tilde{R}$ the (closure of) some bounded non-rectangular open superset of R. Then $q \equiv q(x, y) := \alpha(x)\alpha(y) \cdot \chi(x, y)$ satisfies (3.55) but not [83, Def. 2 eq. (4)]. ◆

[18] And assuming $q \equiv q_{x,y}$ to be of the form [83, Def. 1 eq. (3)] for a probability density $p_{x,y}$ on $\mathbb{R}^2$, which can be established by normalisation and straightforward decay modifications.

3.4 Some Lemmas for Section 3.3

This section contains some of the more technical proofs for Section 3.3.

3.4.1 A Lemma on Monomial Transformations

The following lemma helps us to infer the desired recovery of a source from the existence of a 'well-behaved' subset of its spatial support.

Lemma 3.4.1. *Let $\varrho \in C^1(G)$ for some $G \subseteq \mathbb{R}^d$ open. Then the following holds.*

(i) Let G also be connected and such that $G \cap \pi_i^{-1}(\{\eta\})$ is connected for each $\eta \in \pi_i(G)$ and all $i \in [d]$. Then $\varrho \in \mathrm{DP}_d(G)$ if and only if the Jacobian J_ϱ of ϱ is monomial on G, i.e. such that $\{J_\varrho(u) \mid u \in G\} \subseteq \mathrm{M}_d$.

(ii) Let D be a convex subset of G an such that D is the closure of some dense subset $\mathcal{O}$ of D. If J_ϱ is invertible on D and monomial on $\mathcal{O}$, then ϱ is monomial on D.

Proof. (i) : The 'only-if' direction is clear, so let us prove that $\varrho \in \mathrm{DP}_d(G)$ if

$$J_\varrho(u) \;=\; \big(\beta_\mu(u) \cdot \delta_{\nu,\sigma^u(\mu)}\big)_{\mu,\nu\in[d]} \;\in\; \mathrm{M}_d, \qquad \forall\, u \equiv (u_\mu) \in G, \tag{3.56}$$

for some $\beta_\mu : G \to \mathbb{R}_\times$ and $\{\sigma^u \mid u \in G\} \subseteq S_d$. To this end, we first note that

$$\sigma^u = \sigma^{u'} \quad \text{ for any } \quad u, u' \in G. \tag{3.57}$$

Indeed: Fix any $u_0 \in G$ and note that since the Jacobian $J_\varrho \equiv (J_\varrho^{ij})_{ij} : G \to \mathrm{GL}_d$ of ϱ is continuous, we for each $\varepsilon_{u_0} > 0$ can find a $\delta_{u_0} > 0$ such that ($B_{\delta_{u_0}}(u_0) \subseteq G$ and)

$$\|J_\varrho(u_0) - J_\varrho(u)\| < \varepsilon_{u_0} \quad \text{ for all } \quad u \in B_{\delta_{u_0}}(u_0).$$

Taking $\|\cdot\|$ as the infinity-norm $\|A\| \equiv \|A\|_\infty := \max_{1\le i\le d} \sum_{j=1}^d |a_{ij}|$ and choosing $\varepsilon_{u_0} := \min_{1\le 1\le d} \sum_{j=1}^d |J_\varrho^{ij}(u_0)| > 0$, the fact that $J_\varrho(u_0), J_\varrho(u) \in \mathrm{M}_d$ then readily implies that

$$\sigma^u = \sigma^{u_0} \quad \text{ for all } \quad u \in B_{\delta_{u_0}}(u_0). \tag{3.58}$$

Let now $u, u' \in G$ be arbitrary. Since, as a connected subset of $\mathbb{R}^d$, the domain G is also path-connected, the points u and u' can be joined by a continuous path γ with $u = \gamma(0)$ and $u' = \gamma(1)$ and trace $\bar\gamma := \gamma([0,1]) \subseteq G$. We claim that

$$\sigma^{u_1} = \sigma^{u_2} \quad \text{ for any } \quad u_1, u_2 \in \bar\gamma, \tag{3.59}$$

yielding (3.57) in particular. And indeed: Since $\bar\gamma \subset \bigcup_{u\in\bar\gamma} B_{\delta_u}(u)$ for $\{\delta_u\}$ as in (3.58), the compactness of $\bar\gamma$ yields that, for an $m \in \mathbb{N}$,

$$\bar\gamma \;\subseteq\; \bigcup_{j\in[m]} B_{\delta_{u_j}}(u_j) \qquad \text{for some } \; u_1, \ldots, u_m \in \bar\gamma. \tag{3.60}$$

But since the trace $\bar{\gamma}$ is connected, the $\{u_j\}$ from (3.60) can be renumerated such that $B_{\delta_{u_i}}(u_i) \cap B_{\delta_{u_{i+1}}}(u_{i+1}) \neq \emptyset$ for each $1 \leq i < m$, which implies (3.59) (and hence (3.57)) by way of (3.58).

Hence on G, the Jacobian (3.56) of ϱ is in fact of the form

$$J_\varrho\big|_G \;=\; (\beta_\mu \cdot \delta_{\nu,\sigma(\mu)})_{\mu,\nu\in[d]} \quad \text{for some } \; \beta_\mu \in C(G;\mathbb{R}_\times) \; \text{ and } \; \sigma \in S_d. \tag{3.61}$$

The assertion that $\varrho \equiv (\varrho_i) \in \mathrm{DP}_d(G)$ now follows from (3.61) and the mean value theorem (MVT): Given any $u_0 = (u_0^1, \ldots, u_0^d) \in G$, the MVT implies[19] that for each $u = (u_1, \ldots, u_d) \in G$ which is connected to u_0 via the line segment $\overline{u_0, u} \equiv \{u_0 + t \cdot (u - u_0) \mid t \in [0,1]\} \subset G$ and any $v \in \mathbb{R}^d$, there exists a point $\xi \in \overline{u_0, u}$ such that

$$v \cdot (\varrho(u) - \varrho(u_0)) \;=\; v \cdot J_\varrho(\xi) \cdot (u - u_0). \tag{3.62}$$

Hence if for any fixed $i \in [d]$ we take u with $u_{\sigma(i)} = u_0^{\sigma(i)}$ and choose $v = e_i$ (for $(e_i)_{i\in[d]}$ the standard basis of $\mathbb{R}^d$), then by way of (3.62) and (3.61) we find that

$$\varrho_i(u) - \varrho_i(u_0) = \big[J_\varrho(\xi) \cdot (u - u_0)\big]_i = 0 \qquad (i \in [d]).$$

This implies that for any given $u_0 \equiv (u_0^1, \ldots, u_0^d) \in G$ we have $\varrho_i(u) = \varrho_i(u_0^{\sigma(i)})$ for all $u \in G_{u_0|i} := \{u \in G \mid \exists$ polygonal path in $\pi_{\sigma(i)}^{-1}(\{u_0^{\sigma(i)}\})$ connecting u and $u_0\}$. But since by assumption the slices $G_\eta^j := G \cap \pi_j^{-1}(\{\eta\})$ are each (polygonally-)connected for any $\eta \in \mathbb{R}$ and $j \in [d]$, we have that $G_{u_0|i} = G^{\sigma(i)}_{u_0^{\sigma(i)}}$. As $u_0 \in G$ was arbitrary, we thus find that

$$\varrho_i(u) = \varrho_i(u_{\sigma(i)}) \quad \text{ for each } \; u \equiv (u_1, \ldots, u_d) \in G \qquad (i \in [d]), \tag{3.63}$$

hence the diffeomorphism[20] $\varrho \equiv (\varrho_1, \cdots, \varrho_d)$ is monomial on G as claimed.

(ii) : This is a corollary to the above proof of (i). Indeed, let $\mathcal{O} \subseteq D$ be dense with

$$J_\varrho(u) \in \mathrm{M}_d \quad \text{for each } \; u \in \mathcal{O}.$$

Then for any fixed $z \in D$, the fact that $\mathcal{O}$ is dense in D ensures that there will be a sequence $(u^{(k)})_{k\in\mathbb{N}} \subset \mathcal{O}$ with $\lim_{k\to\infty} u^{(k)} = z$, implying that

$$J_\varrho(z) \;=\; \lim_{k\to\infty} J_\varrho(u^{(k)}) \tag{3.64}$$

due to ϱ being continuously differentiable on G. Hence and because $J_\varrho(z) \in \mathrm{GL}_d$, we obtain that in fact $J_\varrho(z) \in \mathrm{M}_d$ by (3.64) and the fact that the subset M_d is closed in GL_d. The claim now follows from (the proof of) (i) [for $G := D$] upon noting that, due to its convexity, the set D is polygonally-connected and satisfies the slice requirements of statement (i). □

[19] Indeed: For u, u_0, v as above, define $\varphi(t) := v \cdot \varrho(u_0 + t \cdot \eta)$ for $t \in I_\delta \equiv (-\delta, 1+\delta)$ and $\eta := u - u_0$ and $\delta > 0$ s.t. $\{u_0 + t \cdot \eta \mid t \in I_\delta\} \subset G$ (such a δ exists as G is open). Then $\varphi \in C^1(I_\delta)$, whence (3.62) follows from the (classical) MVT applied to the difference $\varphi(1) - \varphi(0)$.

[20] Note that since by (3.61) and (3.63) each ϱ_i is a continuously differentiable map from $\pi_{\sigma(i)}(G)(\subseteq \mathbb{R})$ to $\mathbb{R}$ with nowhere-vanishing derivative, each ϱ_i is a univariate local diffeomorphism and thus in fact a global diffeomorphism on $\pi_{\sigma(i)}(G)$ (cf. e.g. [62, Ex. 1.3.3]).

3.4.2 Proof of Lemma 3.3.4

Recall that Y, Y^* are defined by (3.27) and ρ is given by (3.28).

Lemma 3.3.4. *For μ the probability density of Y, and μ^* the probability density of Y^*,*

$$\psi \circ \rho \;=\; \log \mu - \log \mu^* \qquad a.e.\ on \qquad \tilde{D} := \mathrm{supp}(\mu) \tag{3.29}$$

for the logit-function $\psi(p) := \log(p/(1-p))$.

Proof. We note first that since the support $\mathrm{supp}(\nu)$ of a (Borel) probability measure $\nu : \mathcal{B}(E) \to [0,1]$ is defined as the smallest closed set $C \subseteq E$ having total mass $\nu(C) = 1$, it is easy to see that $\mathrm{supp}(\tilde{f}_*\nu) = \overline{\tilde{f}(\mathrm{supp}(\nu))}$ for any $\tilde{f}$ continuous (cf. the proof of Lem. 3.2.5 (i)). For the above case, this implies $\mathrm{supp}(\mu) = \overline{(f \times f)(\mathrm{supp}(\zeta))} = \overline{(f \times f)(\bar{D})}$, whence $\partial(\mathrm{supp}(\mu))$ is a Lebesgue-nullset (as is $\partial\bar{D}$, by assumption, and hence also the boundary of its C^2-image $(f \times f)(\bar{D})$; the latter boundary in turn contains $\partial\overline{(f \times f)(\bar{D})}$ (as the boundary of the closure of a set is always contained in the boundary of that set) and hence $\mu > 0$ a.e. on $\mathrm{supp}(\mu)$. Since also $\mathrm{supp}(\mu^*) = \mathrm{supp}(\mathbb{P}_{X_s}) \times \mathrm{supp}(\mathbb{P}_{X_t})$, we further obtain $\mathrm{supp}(\mu) = \mathrm{supp}(\mathbb{P}_{(X_s, X_t)}) \subseteq \mathrm{supp}(\mu^*)$ by (3.11), which implies that the RHS of (3.29) is defined a.e. on $\mathrm{supp}(\mu)$ indeed.

Note now that since by definition the function ρ equals the conditional probability of the event $\{C = 1\}$ given $\bar{Y}$, we have

$$\rho \cdot \frac{\mathrm{d}\mathbb{P}_{\bar{Y}}}{\mathrm{d}y} = \mathbb{P}(C = 1 \,|\, \bar{Y}) \cdot \frac{\mathrm{d}\mathbb{P}_{\bar{Y}}}{\mathrm{d}y} = \frac{\mathrm{d}\mathbb{P}_{\bar{Y}}(\,\cdot\,|\, C = 1)}{\mathrm{d}y} \cdot \mathbb{P}(C = 1) \tag{3.65}$$

almost everywhere, where the first factor on the RHS of (3.65) denotes (a regular version of) the conditional density of $\bar{Y}$ given $C = 1$.[21] Next we observe that

$$\frac{\mathrm{d}\mathbb{P}_{\bar{Y}}(\,\cdot\,|\, C = 1)}{\mathrm{d}y} \cdot \mathbb{P}(C = 1) \;=\; \tfrac{1}{2}\mu \qquad \text{(a.e.)}. \tag{3.66}$$

Indeed, denote by η the LHS of (3.66) and let $A \in \mathcal{B}(\mathbb{R}^{2d})$ be arbitrary. Then, since by construction $\mathbb{P}_{(C,\bar{Y})} = \mathbb{P}_C \otimes \mathbb{P}^{\bar{Y}}_C$ and $\mathbb{P}^{\bar{Y}}_{C=1} \equiv \mathbb{P}_{\bar{Y}}(\,\cdot\,|\, C = 1) = \mathbb{P}_Y$ and $\mathbb{P}(C = 1) = \frac{1}{2}$,

$$\begin{aligned}\int_{\mathbb{R}^{2d}} \eta \cdot \mathbb{1}_A \,\mathrm{d}y \;&=\; \mathbb{P}(\bar{Y} \in A \,|\, C = 1) \cdot \mathbb{P}(C = 1) = \mathbb{P}(\bar{Y} \in A,\ C = 1) \\ &=\; \mathbb{P}_{(C,\bar{Y})}(\{1\} \times A) = \mathbb{P}(C = 1) \cdot \mathbb{P}_Y(A) = \int_{\mathbb{R}^{2d}} \tfrac{1}{2}\mu \cdot \mathbb{1}_A \,\mathrm{d}y\end{aligned}$$

[21] Indeed, abbreviating $\ell := \rho \cdot \frac{\mathrm{d}\mathbb{P}_{\bar{Y}}}{\mathrm{d}y}$ and $r := \frac{\mathrm{d}\mathbb{P}(\,\cdot\,|\,C=1)}{\mathrm{d}y} \cdot \mathbb{P}(C = 1)$, we for any $A \in \mathcal{B}(\mathbb{R}^{2d})$ find

$$\begin{aligned}\int_{\mathbb{R}^{2d}} r \cdot \mathbb{1}_A \,\mathrm{d}y \;&= \mathbb{P}(C = 1) \int_A \mathbb{P}_{\bar{Y}}(\mathrm{d}y \,|\, C = 1) = \mathbb{P}((C, \bar{Y}) \in \{1\} \times A) \\ &= \mathbb{P}((\bar{Y}, C) \in A \times \{1\}) = \int_A \mathbb{P}(C = 1 \,|\, \bar{Y} = y)\, \mathbb{P}_{\bar{Y}}(\mathrm{d}y) = \int_{\mathbb{R}^{2d}} \ell \cdot \mathbb{1}_A \,\mathrm{d}y\end{aligned}$$

which implies $r = \ell$ (a.e.) by the fundamental lemma of calculus of variations. (Note that the second and the fourth of the above equations hold by definition of conditional distributions.)

from which (3.66) follows by the fundamental lemma of calculus of variations. Combining (3.65) with the fact that $\frac{\mathrm{d}\mathbb{P}_{\bar{Y}}}{\mathrm{d}y} = \frac{1}{2}\mu + \frac{1}{2}\mu^*$ and (3.66) now yields the identity $(\mu + \mu^*) \cdot \rho = \mu$ (a.e.), from which equation (3.29) follows immediately. □

3.4.3 A Separation Lemma

The following is a core lemma for the proof of Theorem 3.3.3.

Lemma 3.4.2. *Let $U \subseteq \mathbb{R}^d$ be open and $\varphi_i \in C^2(U^{\times 2}; \mathbb{R}_{>0})$, $i \in [d]$, with $\varphi_i(x) \equiv \varphi_i(x_i, x_{i+d})$, be a family of regularly non-separable, positive functions of which all but at most one are a.e. non-Gaussian. Set $\xi_i := \partial_{x_i}\partial_{x_{i+d}} \log \varphi_i$ for each $i \in [d]$. Then for any continuous $B : U \to \mathrm{GL}_d(\mathbb{R})$ for which the composition $\Lambda : U^{\times 2} \to \mathbb{R}^{d \times d}$ given by*

$$\Lambda(u, v) \;:=\; B(u)^\intercal \cdot \mathrm{diag}_{i \in [d]}\big[\xi_i(u_i, v_i)\big] \cdot B(v) \tag{3.67}$$

(in the coordinates $(u, v) \equiv (u_1, \ldots, u_d, v_1, \ldots, v_d) \in U^{\times 2}$) has identically-vanishing off-diagonal elements, it holds that the function B is monomial on U, i.e. that

$$B(u) \;\in\; \mathrm{M}_d \qquad \textit{for each} \;\; u \in U. \tag{3.68}$$

Proof. Set $\check{U} := U \times U$, and for a given $(u, v) \in \check{U}$, denote $\hat{\Lambda}_{u,v} := \mathrm{diag}_{i \in [d]}\big[\xi_i(u_i, v_i)\big]$ and $\Lambda_{u,v} := \Lambda(u, v)$ and $B_u := B(u)$, and assume (wlog, upon re-enumeration) that φ_i is a.e. non-Gaussian for each $i \in [d-1]$.

Note that by the fact that B is GL_d-valued and continuous, the identity (3.68) holds if

$$\exists\, \tilde{U} \subseteq U \text{ dense} \quad \text{such that} \quad B_u \in \mathrm{M}_d \quad \text{for all } u \in \tilde{U} \tag{3.69}$$

(cf. the argument around (3.64) for details). Our proof consists of constructing a set $\tilde{U}$ for which (3.69) holds. Let to this end $i \in [d]$ be fixed, and recall that φ_i being regularly non-separable implies that there is a closed nullset[22] $\mathcal{N} \subset \check{U}$ s.t. for the open and dense[23] subset $\check{U}_\circ := \check{U} \setminus \mathcal{N}$ of $\check{U}$, each restriction $\varphi_i\big|_{\check{U}_\circ}$ is such that

$$\varphi_i\big|_{\check{U}_\circ} \text{ is strictly non-Gaussian for } i \neq d, \qquad \text{and for each } i \in [d]: \tag{3.70}$$

$$\varphi_i\big|_{\check{U}_\circ} \text{ is strictly non-separable} \qquad \text{with} \quad \xi_i\big|_{(\Delta_U \cap \check{U}_\circ)} \neq 0 \;\; \text{everywhere.} \tag{3.71}$$

[22] Notice that if φ_i is regularly non-separable and a.e. non-Gaussian, there (by Definition 3.2.10) will be a closed nullset $\tilde{\mathcal{N}}_i \subset \pi_{(i,i+d)}(\check{U}) \subseteq \mathbb{R}^2$ (s.t. $\tilde{\mathcal{N}}_i \cap \{(x, x) \mid x \in \mathbb{R}\}$ has Hausdorff-measure zero on the diagonal $\Delta_\mathbb{R} := \{(x, x) \mid x \in \mathbb{R}\}$) on whose complement φ_i is strictly non-Gaussian and non-separable and s.t. $\varphi_i\big|_{\Delta_\mathbb{R} \setminus \tilde{\mathcal{N}}_i}$ vanishes nowhere. Hence for the (relatively) closed nullsets $\mathcal{N}_i := \pi^{-1}_{(i,i+d)}(\tilde{\mathcal{N}}_i) \cap \check{U} \subset \mathbb{R}^{2d}$, the (relatively) closed union $\mathcal{N} := \bigcup_{i \in [d]} \mathcal{N}_i$ works as desired.

[23] Recall that for (any) $\check{U} \subseteq \mathbb{R}^m$ open and $\mathcal{N} \subseteq \mathbb{R}^m$ a Lebesgue nullset, the complement $\check{U} \setminus \mathcal{N}$ is dense in $\check{U}$. Indeed: If for $\check{U}_\circ := \check{U} \setminus \mathcal{N}$ we had $\mathrm{clos}(\check{U}_\circ) \subsetneq \check{U}$, then there would be $u \in \check{U}$ with $B_\delta(u) \subseteq \check{U} \setminus \mathrm{clos}(\check{U}_\circ) \subseteq \mathcal{N}$ for some $\delta > 0$, contradicting that the Lebesgue measure of $\mathcal{N}$ is zero.

Given (3.71), Lemma 3.2.12 (ii) (together with the elementary topological facts that (a) : a subset which lies densely inside a dense subspace is itself dense again, and (b) : the intersection of two open dense subsets is again an open dense subset) implies that the intersection

$$\check{U}_* := \bigcap_{i\in[d]} \{z \in \check{U}_\circ \mid \xi_i(z) \neq 0\} \quad \text{is an open dense subset of } \check{U}.$$

Consequently, the coordinate-projections $U' := \pi_{[d]}(\check{U}_*)$ and $V' := \pi_{(d+1,\dots,2d)}(\check{U}_*)$ are open and dense subsets of U $(= \pi_{[d]}(\check{U}))$. We now claim that the sets

$$\tilde{U}_i := \Big\{u \in U' \mid \exists\,(\emptyset \neq)\, \mathcal{V}_u \subseteq V' \text{ open}^{24} \ :\ q^i_u\big|_{\mathcal{V}_u} \text{ is non-constant}\Big\},$$
$$\text{defined by the function} \qquad q^i_u(v) := \frac{\xi_i(u,u)\cdot\xi_i(v,v)}{\xi_i(u,v)^2}, \tag{3.72}$$

are dense in U' — and hence (cf. fact (a)) are also dense in U — for each $i \in [d-1]$.

To see that this holds, we proceed via proof by contradiction and assume that $\tilde{U}_i$ is not dense in U'. In this case, there exists $(\bar{u}, r) \in U' \times \mathbb{R}_{>0}$ with $B_r(\bar{u}) \subset U' \setminus \tilde{U}_i$ (recall that U' is open). Moreover: Since we have $\Delta_U \cap \check{U}_\circ \subset \check{U}_*$ (by (3.71)) and $\check{U}_*$ is open, we can even find a convex open neighbourhood $\mathcal{V}_{\bar{u}} \subseteq B_r(\bar{u})$ of $\bar{u}$ such that the whole square $\mathcal{Q}_{\bar{u}} \equiv \mathcal{V}_{\bar{u}} \times \mathcal{V}_{\bar{u}}$ is contained in $\check{U}_*$. Now by construction, we for each slice $\{u\} \times \mathcal{V}_{\bar{u}} \subset \mathcal{Q}_{\bar{u}}$, $u \in \mathcal{V}_{\bar{u}}$, must have that $q^i_u|_{\mathcal{V}_{\bar{u}}}$ is a constant function, say $q^i_u|_{\mathcal{V}_{\bar{u}}} \equiv c_u$ for some $c_u \in \mathbb{R}$, so that in particular (for $\varrho : \mathcal{V}_{\bar{u}} \ni u \mapsto c_u \in \mathbb{R}$)

$$\xi_i(u,u)\cdot\xi_i(v,v) = \varrho(u)\cdot\xi_i(u,v)^2, \quad \forall\,(u,v) \in \mathcal{Q}_{\bar{u}}. \tag{3.73}$$

But since the square $\mathcal{Q}_{\bar{u}}$ contains its diagonal, i.e.: $\mathcal{Q}_{\bar{u}} \supset \{(u,u) \mid u \in \mathcal{V}_{\bar{u}}\} =: \Delta_{\mathcal{V}_{\bar{u}}}$, we can evaluate the relation (3.73) for the points in $\Delta_{\mathcal{V}_{\bar{u}}}$, yielding that $\varrho \equiv 1$. Consequently,

$$\xi_i\big|_{\mathcal{Q}_{\bar{u}}}(u,v) = \epsilon\cdot\varsigma(u)\cdot\varsigma(v), \quad \forall\,(u,v) \in \mathcal{Q}_{\bar{u}},$$

for $\varsigma : \mathcal{V}_{\bar{u}} \ni x \mapsto \varsigma(x) := \sqrt{|\xi_i(x,x)|}$ and with ϵ denoting the sign of $\xi_i|_{\mathcal{Q}_{\bar{u}}}$. Notice that since ξ_i is continuous and $\mathcal{Q}_{\bar{u}}$ is connected, ϵ will be constant (i.e. $\epsilon \equiv \pm 1$). But since $\mathcal{Q}_{\bar{u}} \subset \check{U}_*$ is symmetric, open and convex, Lemma 3.2.12 (iii) then implies that

$$\varphi_i\big|_{\mathcal{Q}_{\bar{u}}} \quad \text{is pseudo-Gaussian,} \quad \text{in contradiction to (3.70).}$$

This proves that each of the above sets $\tilde{U}_i$, $i \in [d-1]$, must be dense in U.

But since each of the dense subsets $\tilde{U}_i$ is also open by the fact that the quotients in (3.72) are continuous in (u,v), we (once more by the above fact (b)) find that their intersection

$$\tilde{U} := \bigcap_{i\in[d-1]} \tilde{U}_i \quad \text{is a dense subset of } U.$$

[24] Note that for $\tilde{U}_i$ to be well-defined, we in addition require each $\mathcal{V}_u$ to be s.t. $\{u\} \times \mathcal{V}_u \subset \check{U}_*$.

We claim that the above set $\tilde{U}$ satisfies (3.69).

To see this, note first that (3.67) yields[25] the system of matrix equations

$$\begin{cases} B_u^\intercal \cdot \hat{\Lambda}_{u,u} \cdot B_u &= \Lambda_{u,u} \\ B_u^\intercal \cdot \hat{\Lambda}_{u,v} \cdot B_v &= \Lambda_{u,v} \\ B_v^\intercal \cdot \hat{\Lambda}_{v,v} \cdot B_v &= \Lambda_{v,v} \end{cases} \qquad \text{for each} \quad (u,v) \in U^{\times 2}; \tag{3.74}$$

we prove (3.69) by defining a map $\eta : \tilde{U} \to V'$ with the property that (3.74) when evaluated at $(u,v) := (u, \eta(u))$ yields $B_u \in \mathrm{M}_d$ by necessity. Let to this end $u \in \tilde{U}$ be arbitrary. Then by the definition of $\tilde{U}$ (recalling (3.72)), we for each $i \in [d]$ can find some open $\mathcal{V}_u^i \subset V'$ such that the intersection $\mathcal{V}_u := \bigcap_{i\in[d]} \mathcal{V}_u^i$ is non-empty and the continuous map

$$q_u := q_u^1 \times \cdots \times q_u^d \;:\; \mathcal{V}_u \;\to\; \mathbb{R}^d \qquad (q_u^i \text{ as in (3.72)}) \tag{3.75}$$

is non-constant in its first $(d-1)$ components. Hence by the intermediate-value theorem, the set $\pi_{[d-1]}(q_u(\mathcal{V}_u)) \subseteq \mathbb{R}^{d-1}$ has non-empty interior, which implies that for $\nabla^\times := \{(v_i) \in \mathbb{R}^d \mid \exists\, i,j \in [d], i \neq j \;:\; |v_i| = |v_j|\}$ (the closed nullset of all vectors in $\mathbb{R}^d$ having two components differing at most up to a sign), the preimage

$$\tilde{\mathcal{V}}_u := q_u^{-1}(\mathbb{R}^d \setminus \nabla^\times) \;\subset\; \mathcal{V}_u \tag{3.76}$$

will be non-empty. This observation gives rise to maps of the form

$$\eta \;:\; \tilde{U} \to V', \quad u \mapsto \eta(u) \in \tilde{\mathcal{V}}_u,$$

and as we will now see, any such map is of the desired type that we announced above. Indeed: Taking any $(u,v) \in \bigcup_{\tilde{u}\in\tilde{U}} \{\tilde{u}\} \times \tilde{\mathcal{V}}_{\tilde{u}} \subseteq \check{U}_*$, we obtain from (3.74) that[26]

$$B_u^{-1} \cdot \bar{\Lambda}_{u,v} \cdot B_u \;=\; \tilde{\Lambda} \qquad \text{for} \qquad \bar{\Lambda}_{u,v} \equiv \mathrm{diag}_{i\in[d]}[\lambda_{u,v}^i] := \hat{\Lambda}_{u,u} \cdot \hat{\Lambda}_{u,v}^{-2} \cdot \hat{\Lambda}_{v,v} \tag{3.77}$$

and the diagonal matrix $\tilde{\Lambda} := \Lambda_{u,v}^{-1} \cdot \Lambda_{v,v} \cdot \Lambda_{u,v}^{-1} \cdot \Lambda_{u,u}$. Observing now (recalling (3.72)) that

$$(\lambda_{u,v}^1, \cdots, \lambda_{u,v}^d) \;=\; q_u(v) \qquad \text{for the function } \; q_u \; \text{ defined in (3.75),}$$

we immediately obtain by the definition (3.76) of $\tilde{\mathcal{V}}_u$ that

$$\text{the eigenvalues} \quad \lambda_{u,v}^1, \ldots, \lambda_{u,v}^d \;\; \text{of} \;\; \bar{\Lambda}_{u,v} \quad \text{are pairwise distinct.} \tag{3.78}$$

Hence by the elementary fact that diagonal matrices with pairwise distinct eigenvalues are stabilised by monomial matrices only, observation (3.78) by way of the similarity equation (3.77) finally implies $B_u \in \mathrm{M}_d$ — and hence (3.69) — as desired. □

[25] Cf. the proof of Lemma 3.4.3 for details.

[26] Note that the invertibility of $\Lambda_{u,v}$ is obtained from the choice of (u,v).

3.4.4 A Second Separation Lemma

The following lemma underlies the proof of Theorem 3.3.7.

Lemma 3.4.3. *Suppose that in addition to* (3.15) *there are* $\mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2 \in \Delta_2(\mathbb{I})$ *and* $(u, v) \in \mathbb{R}^{2d}$ *such that* S *is* C^2*-regular at* $(\mathfrak{p}_0, (u, v))$, $(\mathfrak{p}_1, (u, u))$ *and* $(\mathfrak{p}_2, (v, v))$ *with density* $\zeta_{\mathfrak{p}_0}$, $\zeta_{\mathfrak{p}_1}$ *and* $\zeta_{\mathfrak{p}_2}$*, respectively. Then for any map* h *which is* C^2*-invertible on some open superset of* D_X *and such that* $h(X)$ *has independent components, we have that*

$$B_\varrho(u)^{-1} \cdot \bar{\Lambda}_{\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2}(u, v) \cdot B_\varrho(u) \ = \ \tilde{\Lambda}_{\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2}(u, v) \tag{3.79}$$

for B_ϱ *the inverse Jacobian of* $\varrho := h \circ f$ *and the diagonal matrices*

$$\begin{aligned} \bar{\Lambda}_{\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2}(u, v) \ &:= \ \Lambda_{\xi_{\mathfrak{p}_1}}(u, u) \cdot \Lambda_{\xi_{\mathfrak{p}_0}}(u, v)^{-2} \cdot \Lambda_{\xi_{\mathfrak{p}_2}}(v, v) \qquad \textit{and} \\ \tilde{\Lambda}_{\mathfrak{p}_0,\mathfrak{p}_1,\mathfrak{p}_2}(u, v) \ &:= \ \Lambda_{q_{\mathfrak{p}_1}}(u, u) \cdot \Lambda_{q_{\mathfrak{p}_0}}(u, v)^{-2} \cdot \Lambda_{q_{\mathfrak{p}_2}}(v, v), \end{aligned} \tag{3.80}$$

where $\Lambda_{\xi_{\mathfrak{p}_\nu}}(x) := \operatorname{diag}[\xi^1_{\mathfrak{p}_\nu}(x), \cdots, \xi^d_{\mathfrak{p}_\nu}(x)]$ *for* $\xi^i_{\mathfrak{p}_\nu}(x) := \partial_{x_i}\partial_{x_{i+d}} \log \zeta_{\mathfrak{p}_\nu}(x)$ *and* $\Lambda_{q_{\mathfrak{p}_\nu}}(x) := \operatorname{diag}[q^1_{\mathfrak{p}_\nu}(x), \cdots, q^d_{\mathfrak{p}_\nu}(x)]$ *given by the LHS of* (3.38) *(in dependence of* $\mathfrak{p}_\nu$*).*

Proof. Copying the argumentation that led to (3.39), we obtain the congruence relations

$$\Lambda_{q_{\mathfrak{p}_\nu}}(\tilde{u}, \tilde{v}) \ = \ B^\intercal_\varrho(\tilde{u}) \cdot \Lambda_{\xi_{\mathfrak{p}_\nu}}(\tilde{u}, \tilde{v}) \cdot B_\varrho(\tilde{v}) \quad \text{for each} \ \ (\tilde{u}, \tilde{v}) \in \{\zeta_{\mathfrak{p}_\nu} > 0\}$$

and $\nu = 0, 1, 2$. Evaluating these at the points $(u, v), (u, u)$ and (v, v), we arrive at the system

$$B_\varrho(u)^\intercal \cdot \Lambda_{\xi_{\mathfrak{p}_0}}(u, v) \cdot B_\varrho(v) \ = \ \Lambda_{q_{\mathfrak{p}_0}}(u, v) \tag{3.81}$$

$$B_\varrho(u)^\intercal \cdot \Lambda_{\xi_{\mathfrak{p}_1}}(u, u) \cdot B_\varrho(u) \ = \ \Lambda_{q_{\mathfrak{p}_1}}(u, u) \tag{3.82}$$

$$B_\varrho(v)^\intercal \cdot \Lambda_{\xi_{\mathfrak{p}_2}}(v, v) \cdot B_\varrho(v) \ = \ \Lambda_{q_{\mathfrak{p}_2}}(v, v). \tag{3.83}$$

From (3.81) we then find that

$$B_\varrho(v) = \Lambda^{-1}_{\xi_{\mathfrak{p}_0}}(u, v) \cdot \left(B_\varrho(u)^\intercal\right)^{-1} \cdot \Lambda_{q_{\mathfrak{p}_0}}(u, v),$$

which, when plugged into (3.83), yields

$$B_\varrho(u)^{-1} \cdot \Lambda_{\xi_{\mathfrak{p}_2}}(v, v) \cdot \Lambda^{-2}_{\xi_{\mathfrak{p}_0}}(u, v) \cdot \left(B_\varrho(u)^\intercal\right)^{-1} = \Lambda_{q_{\mathfrak{p}_2}}(v, v) \cdot \Lambda_{q_{\mathfrak{p}_0}}(u, v)^{-2} \tag{3.84}$$

and hence, upon left-multiplying both sides of (3.82) by the matrix product (3.84),

$$B_\varrho(u)^{-1} \cdot \Lambda_{\xi_{\mathfrak{p}_2}}(v, v)\Lambda^{-2}_{\xi_{\mathfrak{p}_0}}(u, v)\Lambda_{\xi_{\mathfrak{p}_1}}(u, u) \cdot B_\varrho(u) = \Lambda_{q_{\mathfrak{p}_2}}(u, u)\Lambda_{q_{\mathfrak{p}_2}}(v, v)\Lambda_{q_{\mathfrak{p}_0}}(u, v)^{-2}.$$

This last equation is identical to (3.79), as desired. $\square$

3.5 Examples of Applicable Sources

The statistical non-degeneracy assumptions of α-, β- or γ-contrastivity hold for a number of well-established models for stochastic signals, among them most popular copula-based time series models (Section 3.5.1) as well as a variety of Gaussian processes and Geometric Brownian Motion (Section 3.5.2).

3.5.1 Popular Copula-Based Source Models Are α-Contrastive

It is well-known (e.g. [114, Sect. 2.10], [44]) that the temporal structure (3.6) of a scalar stochastic process $S = (S_t)_{t\in\mathbb{I}}$ can be given an analytical representation of the form

$$\zeta_{s,t}(x, y) \;=\; \zeta_s(x)\zeta_t(y) \cdot c_{s,t}(F_s^S(x), F_t^S(y)) \qquad ((s,t) \in \Delta_2(\mathbb{I})), \tag{3.85}$$

where $\zeta_{s,t}$ is the probability density of (S_s, S_t), F_r^S is the cdf of the vector S_r with ζ_r its density, and $c_{s,t} : [0,1]^{\times 2} \to \mathbb{R}$ is the uniquely determined copula density of (S_s, S_t).

Proposition 3.5.1. *Let $S \equiv (S_t)_{t\in\mathbb{I}} \equiv (S^1, \cdots, S^d)$ be an IC stochastic process in $\mathbb{R}^d$ such that S_t admits a C^2-density ζ_t for each $t \in \mathbb{I}$ with the property that $t \mapsto \zeta_t(x)$ is continuous for each $x \in \mathbb{R}^d$. Suppose further that for some $\mathcal{P} \subseteq \Delta_2(\mathbb{I})$ with $\bigcup_{(s,t)\in\mathcal{P}}\{\zeta_s \cdot \zeta_t > 0\}$ dense in D_S,[27] it holds that the copula densities $\{c^i_{s,t} \mid (s,t) \in \mathcal{P}\}$ of S^i (cf. (3.85)) are such that*

$$c^i_{s,t} \;\; \textit{are positive and strictly non-Gaussian} \quad \textit{and} \quad \partial_x\partial_y \log c^i_{s,t} \;\; \textit{vanishes nowhere,} \tag{3.86}$$

for each $i \in [d]$. Then the process S is α-contrastive.

Proof. We verify that S satisfies the conditions of Definition 3.3.2. Take any $i \in [d]$ and $(s,t) \in \mathcal{P}$. By assumption the density $\zeta^i_{s,t}$ of (S^i_s, S^i_t) exists and satisfies (3.85), hence

$$\xi^i_{s,t} := \partial_x\partial_y \log \zeta^i_{s,t} = \zeta^i_s \cdot \zeta^i_t \cdot \left[(\partial_x\partial_y \log c_i) \circ \phi^i_{s,t}\right] \quad \text{on} \quad \{\zeta^i_{s,t} > 0\} \supseteq \tilde{D}^{\times 2}_{s,t} \tag{3.87}$$

for $\tilde{D}_{s,t} := \dot{D}_s \cap \dot{D}_t$ and the map $\phi^i_{s,t} := \mathfrak{s}^i_s \times \mathfrak{s}^i_t : \tilde{D}^{\times 2}_{s,t} \to [0,1]^{\times 2}$ with $\mathfrak{s}^i_r := F^{S^i}_r$.[28] Notice that $\phi^i_{s,t}$ is a differentiable injection since the function $\mathfrak{s}^i_r = \mathfrak{s}^i_r(x) \overset{\text{def}}{=} \int_{-\infty}^{x_i} \zeta^i_r(u)\,\mathrm{d}u$ $(r \in \mathbb{I})$ has positive derivative on $\dot{D}_r$. Hence and since $(\zeta^i_s \cdot \zeta^i_t)|_{\tilde{D}^{\times 2}_{s,t}} > 0$ by construction, the α-contrastivity of S follows by way of (3.85) and (3.87) and assumption (3.86). Indeed: Setting $D_{(s,t)} := \tilde{D}_{s,t}$ for $(s,t) \in \mathcal{P}$, we see that Definition 3.3.2 (i) holds by the assumption on $\mathcal{P}$, while Definition 3.3.2 (ii) is immediate by (3.85), (3.87) and (3.86) and the above-noted fact that $\phi^i_{s,t} : \tilde{D}^{\times 2}_{s,t} \to \phi^i_{s,t}(\tilde{D}^{\times 2}_{s,t})$ is a diffeomorphism. ☐

[27] Lemma 3.2.5 (v) guarantees that such a set $\mathcal{P}$ exists.

[28] Recall that, by convention, we write $\zeta^i_{s,t}(x) \equiv \zeta^i_{s,t}(x_i, x_{i+d})$ for $x = (x_1, \cdots, x_{2d}) \in \mathbb{R}^{2d}$.

A popular approach in finance, insurance economy and other fields is to read (3.85) as a semi-parametric stationary model for $S = (S_t)_{t\in\mathbb{I}}$ by assuming the existence of some $\mathcal{I} \subset \mathbb{I}$ discrete ('set of observations') such that $\zeta_r \equiv \zeta$ with cdf F_ζ for each $r \in \mathcal{I}$, and $D_S = \operatorname{supp}(\zeta)$ and $c_{s,t} \equiv c_\theta$ uniformly parametrized for all $(s,t) \in \mathcal{P} := \mathcal{I}^{\times 2} \cap \Delta_2(\mathbb{I})$, e.g. [32, Sect. 2], [50]:

$$\zeta_{s,t}(x,y) \;=\; \zeta(x)\zeta(y)\cdot c_\theta(F_\zeta(x), F_\zeta(y)), \qquad (s,t) \in \mathcal{P}. \tag{3.88}$$

We verify exemplarily that a source $S = (S^1, \cdots, S^d)$ in $\mathbb{R}^d$ whose components S^i are modelled according to (3.88) is α-contrastive for a number of popular copula densities c_θ.

Corollary 3.5.2. *Let $S = (S^1, \cdots, S^d)$ be a stochastic process whose independent components S^i are modelled according to* (3.88) *for each $i \in [d]$ with copula-density c_i belonging to one of the following popular classes:*

(i) (Clayton) $\qquad c_i(x,y) = (1+\theta)(xy)^{(-1-\theta)}(-1 + x^{-\theta} + y^{-\theta})^{(-2-1/\theta)}$

where $\theta \in (-1,\infty) \setminus \{0, -\frac{1}{2}\}$;

(ii) (Gumbel) $\qquad c_i(x,y) = 1 + \theta(1-2x)(1-2y), \qquad \theta \in [-1,1] \setminus \{0\}$;

(iii) (Frank) $\qquad c_i(x,y) = \dfrac{\theta e^{\theta(x+y)}(e^\theta - 1)}{(e^\theta - e^{\theta x} - e^{\theta y} + e^{\theta(x+y)})^2}, \qquad \theta \in \mathbb{R} \setminus \{0\}$.

Then S is α-contrastive.

Proof. This is a direct consequence of Proposition 3.5.1 upon checking that each of the copula densities (i), (ii) and (iii) satisfies (3.86). This, however, follows from inspection and a straightforward computational verification. □

3.5.2 Popular Gaussian Processes and Geometric Brownian Motion are γ-Contrastive

Given an interval $\mathbb{I}$ and functions $\mu : \mathbb{I} \to \mathbb{R}^d$ and $\kappa : \mathbb{I}^{\times 2} \to \mathrm{GL}_d(\mathbb{R})$, we write $S \sim \mathcal{GP}_\mathbb{I}(\mu,\kappa)$ to denote that $S = (S_t)_{t\in\mathbb{I}}$ is a Gaussian Process in $\mathbb{R}^d$ with mean $\mu = (\mu_i)$ and covariance $\kappa = (\kappa^{ij})$. We assume that any pair (μ,κ) we consider in the following is such that each process $S \sim \mathcal{GP}(\mu,\kappa)$ admits a version with continuous sample paths.

Lemma 3.5.3. *Let $S \sim \mathcal{GP}_\mathbb{I}(\mu,\kappa)$ be a (continuous) Gaussian process in $\mathbb{R}^d$ with diagonal covariance function $\kappa \equiv (\kappa^{ij}) = (\kappa^{ij}\delta_{ij})$. Then S is γ-contrastive if and only if there exist pairs $\mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2 \in \Delta_2(\mathbb{I})$ such that*

$$\left(\frac{\kappa^i_{\mathfrak{p}_1}\cdot\kappa^i_{\mathfrak{p}_2}\cdot[k^i_{\mathfrak{p}_0} - (\kappa^i_{\mathfrak{p}_0})^2]^2}{[k^i_{\mathfrak{p}_1} - (\kappa^i_{\mathfrak{p}_1})^2]\cdot[k^i_{\mathfrak{p}_2} - (\kappa^i_{\mathfrak{p}_2})^2]\cdot(\kappa^i_{\mathfrak{p}_0})^2}\right)_{i\in[d]} \;\in\; (\mathbb{R}^d \setminus \nabla^\times) \tag{3.89}$$

for the auxiliary functions

$$\kappa^i_{s,t} := \kappa^{ii}(s,t) \quad \textit{and} \quad k^i_{s,t} := \kappa^{ii}(s,s)\cdot\kappa^{ii}(t,t). \tag{3.90}$$

Proof. Let $i \in [d]$ be fixed, and $(s,t) \in \Delta_2(\mathbb{I})$ be arbitrary. Since $S^i \sim \mathcal{GP}(\mu_i, \kappa^{ii})$, we find that[29] $(S^i_s, S^i_t) \sim \mathcal{N}(\mu^i_{s,t}, \Sigma^i_{s,t})$ with $\mu^i_{s,t} := (\mu_i(s), \mu_i(t))$ and $\Sigma^i_{s,t} := (\kappa^{ii}(t_\nu, t_{\tilde{\nu}}))_{\nu,\tilde{\nu}=1,2}$ for $t_1 := s$ and $t_2 := t$, so the density $\zeta^i_{s,t}$ of (S^i_s, S^i_t) exists, is smooth on $\mathbb{R}^2$ and reads

$$\zeta^i_{s,t}(x,y) = c_i(s,t)\cdot\exp(\varphi_i(s,t,x,y)) \qquad \text{with}$$
$$\varphi_i(s,t,x,y) = \frac{\rho_{s,t}}{(1-\rho^2_{s,t})\sigma_s\sigma_t}\cdot xy \ + \ \eta_i(s,t,x) \ + \ \tilde{\eta}_i(s,t,y)$$

for $c_i(s,t) = (4\pi^2\sigma_s^2\sigma_t^2(1-\rho^2_{s,t}))^{-1/2}$ and $-\frac{1}{2}((x,y)-\mu^i_{s,t})^\intercal\cdot[\Sigma^i_{s,t}]^{-1}\cdot((x,y)-\mu^i_{s,t}) - \frac{\rho_{s,t}}{(1-\rho^2_{s,t})\sigma_s\sigma_t}\cdot xy =: \eta_i(s,t,x) + \tilde{\eta}_i(s,t,y)$ and with the (auto-)correlations

$$\sigma_r := \sqrt{\kappa^{ii}(r,r)} \qquad \text{and} \qquad \rho_{s,t} := \frac{\kappa^{ii}(s,t)}{\sqrt{\kappa^{ii}(s,s)\kappa^{ii}(t,t)}} = \frac{\kappa^i_{s,t}}{\sqrt{k^i_{s,t}}}.$$

Consequently, the mixed log-derivatives $\xi^i_{s,t} := \partial_x\partial_y \log(\zeta^i_{s,t})$ are given as

$$\xi^i_{s,t}(x,y) = \partial_x\partial_y\varphi_i(s,t,x,y) = \frac{\kappa^i_{s,t}}{k^i_{s,t} - (\kappa^i_{s,t})^2}. \tag{3.91}$$

Since by Definition 3.3.5 the process S is γ-contrastive iff there are $\mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2 \in \Delta_2(\mathbb{I})$ with $\left(\psi(\xi^i_{\mathfrak{p}_0}, \xi^i_{\mathfrak{p}_1}, \xi^i_{\mathfrak{p}_2})\right)_{i\in[d]} = \left(\frac{\xi^i_{\mathfrak{p}_1}\xi^i_{\mathfrak{p}_2}}{(\xi^i_{\mathfrak{p}_0})^2}\right)_{i\in[d]} \in (\mathbb{R}^d \setminus \nabla^\times)$, the lemma now follows from (3.91). □

Remark 3.5.4. The above lemma asserts that IC Gaussian processes are 'generically identifiable', namely if the function (3.89) of their autocovariances avoids the nullset $\nabla^\times$ for some time pairs $\mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2$. Compare this to the well-known result [14] that an IC Gaussian process S is identifiable from its linear mixtures – via joint diagonalisation of the covariance matrices of such mixtures at one or several time lags – if the (vector whose components are the) autocovariances of S themselves avoids the nullset $\nabla^\times$ at one of these time lags.

We verify the above contrastivity condition for a number of popular Gaussian processes.

Proposition 3.5.5. *Let $S = (S^1_t, \cdots, S^d_t)_{t\in\mathbb{I}}$ be an IC stochastic process in $\mathbb{R}^d$ with $S^i \sim \mathcal{GP}(\mu_i, \kappa_i)$ for each $i \in [d]$. Then S is γ-contrastive in each of these four classical cases.*

(i) *For each $i \in [d]$, the componental autocovariance functions* (3.90) *of S are of the form*

$$\kappa^i(s,t) = \exp\left(-\left[\frac{|t-s|}{\alpha_i}\right]^{\gamma_i}\right)$$

with $\gamma \equiv (\gamma_i)_{i\in[d]} \in (0,2]^d$ and $\alpha \equiv (\alpha_i)_{i\in[d]} \in (\mathbb{R}_\times)^{\times d} \setminus \mathcal{N}_\gamma$, where $\mathcal{N}_\gamma \subset \mathbb{R}^d$ is a Lebesgue nullset defined in the proof below.[30]

[29] We write $Z \sim \mathcal{N}(\mu, \Sigma)$ to say that Z is normally distributed with mean $\mu \in \mathbb{R}^d$ and covariance $\Sigma \in \mathbb{R}^{d\times d}$.

[30] This includes the family of γ-exponential processes, cf. [128, Sect. 4.2 (pp. 84 ff.)].

(ii) *Each component process S^i of S is an Ornstein-Uhlenbeck process*

$$\mathrm{d}S_t^i = \theta_i \cdot (\mu_i - S_t^i)\,\mathrm{d}t + \sigma_i\,\mathrm{d}B_t^i, \quad S_0^i = a_i, \qquad (i \in [d]) \tag{3.92}$$

with $a_i, \mu_i \in \mathbb{R}$ and $\sigma \equiv (\sigma_i)_{i\in[d]} \in \mathbb{R}^d_{>0}$ and $\theta \equiv (\theta_i)_{i\in[d]} \in \mathbb{R}^d_{>0} \setminus \tilde{\mathcal{N}}$, where $\tilde{\mathcal{N}} \subset \mathbb{R}^d$ is a Lebesgue nullset defined in the proof below.

(iii) *The component processes of S are fractional Brownian motions with pairwise distinct Hurst indices, that is their autocovariance functions* (3.90) *take the form*

$$\kappa^i(s,t) = \frac{1}{2}(|t|^{2H_i} + |s|^{2H_i} - |t-s|^{2H_i}) \qquad (i \in [d])$$

for some $(H_i)_{i\in[d]} \in (0,1)^d \setminus \nabla^\times$.

(iv) *Denoting $s \wedge t := \min(s,t)$, the autocovariance functions* (3.90) *of the S^i are of the form*

$$\kappa^i(s,t) = \int_0^{s\wedge t} \eta_i(r)\,\mathrm{d}r \qquad \textit{for each } \; i \in [d],$$

with functions $\eta_1, \ldots, \eta_d : \mathbb{I} \to \mathbb{R}$ for which there are $r_0, r_1 \in \mathbb{I}$ such that the products $\{\eta_i(r_0)\cdot\eta_j(r_1) \mid i,j \in [d]\}$ are pairwise distinct. This includes deterministic signals perturbed by white noise, i.e. signals $S = (S_t^1, \cdots, S_t^d)_{t\in\mathbb{I}}$ which, for $(B_t^i)_{t\geq 0}$ some standard Brownian motion in $\mathbb{R}^d$, are given by

$$\mathrm{d}S_t^i = \mu_i(t)\,\mathrm{d}t + \sigma_i(t)\,\mathrm{d}B_t^i \qquad \textit{for each } \; i \in [d]$$

with $\mu_i, \sigma_i : \mathbb{I} \to \mathbb{R}$ integrable and continuous such that the entries of $(\sigma_i^2(r_0)\cdot\sigma_j^2(r_1))_{i,j\in[d]}$ are pairwise distinct for some $r_0, r_1 \in \mathbb{I}$.

Proof. We apply Lemma 3.5.3 by showing that for each case there are $\mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2 \in \Delta_2(\mathbb{I})$ for which (3.89) holds. Write Ξ_i for the i^{th} component of (3.89) and let $|(s,t)| := |t-s|$.

(i): Fix any $\mathfrak{p}_0, \mathfrak{p}_1 \in \Delta_2(\mathbb{I})$ with $|\mathfrak{p}_0| \neq |\mathfrak{p}_1|$ and take $\mathfrak{p}_2 := \mathfrak{p}_0$. Then for each $i \in [d]$ we have $\Xi_i = \tilde{\Xi}_i^{(1)} / \tilde{\Xi}_i^{(0)}$ for the factors

$$\tilde{\Xi}_i^{(\nu)} = \frac{\kappa^i_{\mathfrak{p}_\nu}}{1-(\kappa^i_{\mathfrak{p}_\nu})^2} = \left([\kappa^i_{\mathfrak{p}_\nu}]^{-1} - \kappa^i_{\mathfrak{p}_\nu}\right)^{-1} = \tfrac{1}{2}\left(\sinh\left(\left[\frac{|\mathfrak{p}_\nu|}{\alpha_i}\right]^{\gamma_i}\right)\right)^{-1}.$$

Hence we have the parametrisation $\alpha_i \mapsto \Xi_i \equiv \Xi_i(\alpha_i)$ given by

$$\Xi_i = \sinh\left(\left[\frac{|\mathfrak{p}_0|}{\alpha_i}\right]^{\gamma_i}\right)\cdot\left[\sinh\left(\left[\frac{|\mathfrak{p}_1|}{\alpha_i}\right]^{\gamma_i}\right)\right]^{-1},$$

which for $|\mathfrak{p}_1| \neq |\mathfrak{p}_0|$ and γ_i fixed is differentiable and strictly monotone in $\alpha_i > 0$. Denoting by ϕ_i the associated (differentiable) inverse of $\alpha_i \mapsto \Xi_i(\alpha_i)$ (for γ and $\mathfrak{p}_0, \mathfrak{p}_1$ fixed), we find that $(\Xi_i(\alpha_i))_{i\in[d]} \in (\mathbb{R}^d \setminus \nabla^\times)$ for any $(\alpha_i)_{i\in[d]}$ not contained in $\mathcal{N}_\gamma := (\phi_1 \times \cdots \times \phi_d)(\nabla^\times)$.

(ii) : It is well-known that the covariance function κ^i of (3.92) reads

$$\kappa^i(s,t) = \gamma_i\big(e^{-\theta_i|s-t|} - e^{-\theta_i(s+t)}\big) \quad \text{for} \quad \gamma_i := \frac{\sigma_i^2}{2\theta_i}.$$

Suppose $\mathbb{I} = [0,1]$ wlog. Then for $\mathfrak{p}_\nu \equiv (t_\nu, 2t_\nu) \in \Delta_2(\mathbb{I})$ ($\nu = 0,1$) with $t_0 \neq t_1$ and $\mathfrak{p}_2 := \mathfrak{p}_0$, we obtain $\kappa^i_{\mathfrak{p}_\nu} = \gamma_i(e^{-\theta_i t_\nu} - e^{-3\theta_i t_\nu})$ and $k^i_{\mathfrak{p}_\nu} = \kappa^i_{\mathfrak{p}_\nu} \cdot \gamma_i e^{\theta_i t_\nu} \cdot (1 - e^{-4\theta_i t_\nu})$ and, thus, $\Xi_i = \tilde{\Xi}_i^{(1)} / \tilde{\Xi}_i^{(0)}$ for each $i \in [d]$, with the factors

$$\tilde{\Xi}_i^{(\nu)} = \frac{\kappa^i_{\mathfrak{p}_\nu}}{k^i_{\mathfrak{p}_\nu} - (\kappa^i_{\mathfrak{p}_\nu})^2} = \Big(\gamma_i e^{\theta_i t_\nu} \cdot (1 - e^{-4\theta_i t_\nu}) - \kappa^i_{\mathfrak{p}_\nu}\Big)^{-1} = \left(\frac{\sigma_i^2}{\theta_i}\sinh(\theta_i t_\nu)\right)^{-1}.$$

We hence have the parametrisation $\theta_i \mapsto \Xi_i \equiv \Xi_i(\theta_i) = \frac{\sinh(\theta_i t_0)}{\sinh(\theta_i t_1)}$, which due to $t_0 \neq t_1$ is strictly monotone and differentiable in $\theta_i > 0$. Denoting by $\tilde{\phi}_i$ the associated (differentiable) inverse of $\theta_i \mapsto \Xi_i(\theta_i)$, we obtain that $(\Xi_i(\theta_i))_{i\in[d]} \in (\mathbb{R}^d \backslash \nabla^\times)$ provided the parameter vector $(\theta_i)_{i\in[d]} \in \mathbb{R}^d_{>0}$ is not contained in the nullset $\tilde{\mathcal{N}} := (\tilde{\phi}_1 \times \cdots \times \tilde{\phi}_d)(\nabla^\times)$.

(iii) : Choosing again $\mathfrak{p}_\nu \equiv (t_\nu, 2t_\nu) \in \Delta_2(\mathbb{I})$ ($\nu = 0,1$) with $t_0 \neq t_1$ and $\mathfrak{p}_2 := \mathfrak{p}_0$, we for each $i \in [d]$ find that $\Xi_i = \tilde{\Xi}_i^{(1)} / \tilde{\Xi}_i^{(0)}$ for the factors

$$\tilde{\Xi}_i^{(\nu)} = \frac{4^{H_i - \frac{1}{2}} \cdot t_\nu^{2H_i}}{4^{H_i} \cdot t_\nu^{4H_i} - 4^{2H_i - 1} \cdot t_\nu^{4H_i}} = \Big(2(1 - 4^{H_i - 1}) \cdot t_\nu^{2H_i}\Big)^{-1},$$

whence it holds that

$$\Xi_i = \left(\frac{t_0}{t_1}\right)^{2H_i} \qquad \text{for each} \quad i \in [d].$$

But since due to $t_0 \neq t_1$ the assignment $h \mapsto (\frac{t_0}{t_1})^{2h}$ is clearly injective, we clearly obtain that $(\Xi_i)_{\in[d]}$ is not in $\nabla^\times$ whenever $(H_i)_{i\in[d]}$ is not in $\nabla^\times$, as claimed.

(iv) : As the numbers $\theta_{ij} := \eta_i(r_0) \cdot \eta_j(r_1)$, $(i,j) \in [d] \times [d]$, are pairwise distinct and (thus) non-zero, the continuity of the functions $\vartheta_{ij} : \mathbb{I}^{\times 2} \ni (s,t) \mapsto \eta_i(s) \cdot \eta_j(t)$ allows us to find pairs $(s_0, t_0), (s_1, t_1) \in \Delta_2(\mathbb{I})$ such that for the rectangle $R := [s_0, t_0] \times [s_1, t_1] \subseteq \mathbb{I}^{\times 2}$ the associated integrals

$$\int_R \vartheta_{ij} \, \mathrm{d}s \, \mathrm{d}t, \quad (i,j) \in [d]^{\times 2}, \quad \text{are pairwise distinct and non-zero.} \tag{3.93}$$

Clearly then, (3.93) implies that for $\iota_i(s,t) := \int_s^t \eta_i(r) \, \mathrm{d}r$, the numbers

$$\iota_i(s_0, t_0) \cdot \iota_j(s_1, t_1), \ (i,j) \in [d] \times [d], \ \text{are pairwise dinstinct} \tag{3.94}$$

(Note further that by (3.93), s_0 and s_1 may be chosen such that in addition to (3.94) it holds $\iota_i(0, s_\nu) \neq 0$ for each $i \in [d]$.) Now by setting $\mathfrak{p}_2 := \mathfrak{p}_0$ with $\mathfrak{p}_\nu := (s_\nu, t_\nu)$ for $\nu = 0,1$

(notice that $\mathfrak{p}_0 \neq \mathfrak{p}_1$ by (3.94)), we once more find that $\Xi_i = \tilde{\Xi}_i^{(1)} / \Xi_i^{(0)}$ for each $i \in [d]$, this time for the factors

$$\tilde{\Xi}_i^{(\nu)} = \frac{\kappa^i_{\mathfrak{p}_\nu}}{k^i_{\mathfrak{p}_\nu} - (\kappa^i_{\mathfrak{p}_\nu})^2} = \frac{\iota_i(0,s_\nu)}{\iota_i(0,s_\nu) \cdot \iota_i(0,t_\nu) - \iota_i(0,s_\nu)^2} = (\iota_i(s_\nu,t_\nu))^{-1}\,.$$

Consequently, the entries of $(\Xi_i)_{i\in[d]} = \left(\frac{\iota_i(s_0,t_0)}{\iota_i(s_1,t_1)}\right)_{i\in[d]}$ are pairwise distinct. Indeed, assuming otherwise that $\Xi_i = \Xi_j$ for some $i \neq j$, we find that

$$\iota_i(s_0,t_0) \cdot \iota_j(s_1,t_1) = \iota_j(s_0,t_0) \cdot \iota_i(s_1,t_1), \qquad \text{contradicting (3.94).}$$

Hence $(\Xi_i)_{i\in[d]} \in (\mathbb{R}^d \setminus \nabla^\times)$ as desired. $\square$

The proposition below concludes our short compilation of applicable source models.

Proposition 3.5.6. *Let $S = (S_t)_{t\geq 0} = (S^1, \cdots, S^d)$ be an IC geometric Brownian motion in $\mathbb{R}^d$, i.e. suppose that there is a standard Brownian motion $B = (B^1_t, \cdots, B^d_t)_{t\geq 0}$ such that*

$$\mathrm{d}S^i_t \;=\; S^i_t \cdot (\mu_i(t)\,\mathrm{d}t \;+\; \sigma_i(t)\,\mathrm{d}B^i_t), \quad S^i_0 = s^i_0 \qquad (i \in [d])$$

for some $s^i_0 > 0$ and continuous functions $\mu_i : \mathbb{I} \to \mathbb{R}$ and $\sigma_i : \mathbb{I} \to \mathbb{R}_{>0}$. Then S has spatial support $D_S = \mathbb{R}^d_+$, and S is γ-contrastive if there are $r_0, r_1 \in \mathbb{I}$ for which the numbers $\{\sigma_i^2(r_0) \cdot \sigma_j^2(r_1) \mid (i,j) \in [d]\times[d]\}$ are pairwise distinct.

Proof. A straightforward application of Itô's lemma yields that for any $(s,t) \in \Delta_2(\mathbb{I})$, the density $\zeta^i_{s,t}$ of (S^i_s, S^i_t) is given by

$$\zeta^i_{s,t}(x,y) = \rho^i_{s,t}(\log(x),\log(y)) \cdot (xy)^{-1} = c_i(s,t,x,y) \cdot \exp\big(\varphi_i(s,t,x,y)\big)$$

for the functions $c_i : \Delta_2(\mathbb{I}) \times \mathbb{R}^2_{>0} \to \mathbb{R}$ and $\varphi_i : \Delta_2(\mathbb{I}) \times \mathbb{R}^2_{>0} \to \mathbb{R}$ defined by

$$\begin{aligned} c_i(s,t,x,y) &= \frac{(4\pi^2 \cdot \det(\mathfrak{s}^i_{s,t}))^{-1/2}}{xy} \qquad\qquad \text{and} \\ \varphi_i(s,t,x,y) &= -\frac{1}{2}\left([\phi^{-1}(x,y) - \mathfrak{m}^i_{s,t}]^\intercal \cdot (\mathfrak{s}^i_{s,t})^{-1} \cdot [\phi^{-1}(x,y) - \mathfrak{m}^i_{s,t}]\right) \\ &= \beta^i_{s,t} \log(x)\log(y) + \eta_i(s,t,x) + \tilde{\eta}_i(s,t,y) \end{aligned}$$

with $\eta_i, \tilde{\eta}_i$ given by $\eta_i(s,t,x) + \tilde{\eta}_i(s,t,y) := \varphi_i(s,t,x,y) - \beta^i_{s,t}$ and

$$\beta^i_{s,t} := \frac{\kappa_i(s,t)}{\kappa_i(s,s)\kappa_i(t,t) - \kappa_i^2(s,t)} = \frac{\int_0^s \sigma_i^2(r)\,\mathrm{d}r}{\left(\int_0^s \sigma_i^2(r)\,\mathrm{d}r\right)\left(\int_0^t \sigma_i^2(r)\,\mathrm{d}r\right) - \left(\int_0^s \sigma_i^2(r)\,\mathrm{d}r\right)^2}.$$

Consequently, the spatial support of S is $D_S = \mathbb{R}^d_+ = \pi_{[d]}\big(\mathrm{supp}[\mathbb{R}^{2d} \ni (u,v) \mapsto \zeta^i_{s,t}(u_i, v_i)]\big)$ (any $i \in [d]$), and the mixed log-derivatives of $\zeta^i_{s,t}$ read

$$\xi^i_{s,t} := \partial_x\partial_y \log(\zeta^i_{s,t}) = \partial_x\partial_y\varphi^i(s,t,\cdot,\cdot) = \frac{\beta^i_{s,t}}{xy}.$$

Hence by Definition 3.3.5, the process S is γ-contrastive iff

$$\exists\, \mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2 \in \Delta_2(\mathbb{I}) \quad \text{with} \quad (\Xi_i)_{i\in[d]} := \left(\frac{\beta^i_{\mathfrak{p}_1}\beta^i_{\mathfrak{p}_2}}{(\beta^i_{\mathfrak{p}_0})^2}\right)_{i\in[d]} \in (\mathbb{R}^d \setminus \nabla^\times).$$

Having $\mathfrak{p}_\nu \equiv (s_\nu, t_\nu) \in \Delta_2(\mathbb{I})$ $(\nu = 0, 1)$ arbitrary and $\mathfrak{p}_2 := \mathfrak{p}_0$ hence yields that

$$\Xi_i = \frac{\beta^i_{\mathfrak{p}_1}}{\beta^i_{\mathfrak{p}_0}} = \frac{\int_{s_0}^{t_0}\sigma_i^2(r)\,\mathrm{d}r}{\int_{s_1}^{t_1}\sigma_i^2(r)\,\mathrm{d}r} \qquad \text{for each} \;\; i \in [d].$$

Thus by choosing $\mathfrak{p}_0$ and $\mathfrak{p}_1$ as done in the proof of Proposition 3.5.5 (iv), we obtain $(\Xi_i)_{i\in[d]} \in (\mathbb{R}^d \setminus \nabla^\times)$ as desired. □

3.6 Nonlinear ICA via Marginal Distributions

One of the main insights behind our approach to Nonlinear ICA is that $\mathcal{C}_d$-valued random variables carry enough structure to allow for the disentanglement of their d nonlinearly mixed independent $\mathcal{C}_1$-valued components by applying to such a mixture those transformations of $\mathbb{R}^d$ for which the component processes of the transformed mixture are independent again (Theorems 3.3.3 and 3.3.7).

Roughly speaking, the reason for this is that by being stochastic objects in $(\mathbb{R}^d)^{\mathbb{I}}$ instead of just in $\mathbb{R}^d$, random processes have an additional degree of structure ('inter-temporality') that allows them to carry an infinitude of sequentially ordered, statistically encoded [as probability distributions] spatial information which, via inter-temporal comparisons ('relating joint distributions at different points in time', cf. Section 3.3.1), can then be unlocked to infer about the structure of a given transformation of $\mathbb{R}^d$ that acts on this process.

A single random vector in $\mathbb{R}^d$, on the other hand,[31] carries spatial information only, which turns out to be enough to infer about linear maps on $\mathbb{R}^d$ but is in general insufficient for the inversion of nonlinear transformations of $\mathbb{R}^d$ (cf. Example 3.1.4).

[31] Similar to how a moving picture ($\sim$ stochastic process in $\mathbb{R}^d$) contains more information than each of its frames ($\sim$ random vectors in $\mathbb{R}^d$).

However: By asserting that a given map[32] $h : \mathbb{R}^d \to \mathbb{R}^d$ can be identified as a monomial transformation (Definition 3.3.1) if its induced pushforward map $h_* : \mathcal{M}_1(\mathbb{R}^d) \to \mathcal{M}_1(\mathbb{R}^d)$,

$$\mu \;\mapsto\; h_*(\mu) \equiv \mu \circ h^{-1}, \tag{3.95}$$

preserves the product structure of its arguments, the following proposition shows that spatial information (i.e. laws on $\mathbb{R}^d$, as opposed to (3.6)) alone *can be sufficient* to invert a given transformation on $\mathbb{R}^d$, namely if a *vast amount* of this information – specifically, the validity of the implication "μ is product $\Rightarrow h_*(\mu)$ is product" for every $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ – is available.

This opens up another perspective on the conceptional significance of stochastic processes in the context of ICA, relating this time to their marginal distributions (i.e. their 1st-order fdds, cf. Section 3.2.1): While a single law on $\mathbb{R}^d$ ($\sim$ random vector in $\mathbb{R}^d$) is generally not enough to ensure the independence-based inversion of a nonlinear transformation (3.95) of this law (up to monomial ambiguity), the totality $\mathcal{M}_1(\mathbb{R}^d)$ of all laws on $\mathbb{R}^d$ is sufficient for this purpose (Proposition 3.6.1); as a parametrised family of infinitely many laws on $\mathbb{R}^d$ (its marginal distributions), a stochastic process may bridge this 'transition from just a single law on $\mathbb{R}^d$ to all the laws on $\mathbb{R}^d$' by providing a family of marginals that is "sufficiently rich/varied" to imply an (3.96)-based blind inversion of its nonlinear diffeomorphic transformations.

Similar to how our identifiability conditions of α-, β- or γ-contrastivity are specifications on the 2nd-order fdds (3.6) of a process to be 'sufficiently non-degenerate', we expect that according specifications on what it means for the marginals of the process to be 'sufficiently varied' (so as to weaken the sufficient identifiability condition (3.96)) may yield alternative conditions for an independence-based recovery of this process from its nonlinear mixtures (implementable still via Theorem 4.2.3). We leave the concrete specification of such conditions for future research, with the following proof offered as a source of potentially relevant ideas.

Denote by $\mathcal{S}_\perp$ the set of all random vectors in $\mathbb{R}^d$ with mutually independent components.

Proposition 3.6.1. *Let $h \in C^{1,1}(\mathbb{R}^d;\mathbb{R}^d)$ be a map with the property that*

$$h(U) \in \mathcal{S}_\perp \quad \textit{for each} \quad U \in \mathcal{S}_\perp. \tag{3.96}$$

Then h is necessarily a monomial transformation.

[32] The assumption of the map h being defined on all of $\mathbb{R}^d$ is for convenience only; generalising the proof of Proposition 3.6.1 to maps defined on open subdomains of $\mathbb{R}^d$ is straightforward.

Proof. Let $h \equiv (h_1, \cdots, h_d) \in C^1(\mathbb{R}^d; \mathbb{R}^d)$ be a C^1-diffeomorphism satisfying (3.96). Noting that the premise (3.96) is equivalent to the identity of Borel measures on $\mathbb{R}^d$

$$\mathbb{P}_{h(U)} = \mathbb{P}_{h_1(U)} \otimes \cdots \otimes \mathbb{P}_{h_d(U)} \quad \forall\, U \in \mathcal{S}_\perp, \tag{3.97}$$

our proof of the asserted inclusion $h \in \mathrm{DP}_d$, i.e. that h be monomial, is based on an algebraic comparison of the (inverse) Fourier transforms (i.e. the characteristic functions) of both sides of the identity (3.97). More precisely: Using that for each random vector $U \equiv (U_1, \cdots, U_d)$ in $\mathbb{R}^d$ whose components $U_j \equiv \langle U, e_j \rangle_2$ are independent we by way of (3.97) obtain the following product identity (of functions in $\xi \equiv (\xi_1, \ldots, \xi_d) \in \mathbb{R}^d$),

$$\mathbb{E}[e^{\mathrm{i}\xi^\intercal \cdot h(U)}] = \mathbb{E}[e^{\mathrm{i}\xi_1 \cdot h_1(U)}] \cdots \mathbb{E}[e^{\mathrm{i}\xi_d \cdot h_d(U)}], \tag{3.98}$$

we show that for each $k \in [d]$ and arbitrary $\eta \equiv (\eta_1, \ldots, \widehat{\eta_k}, \ldots \eta_d) \in \mathbb{R}^{d-1}$ it holds that

$$\forall\, \alpha, \beta \in \mathbb{R} \text{ with } \alpha \neq \beta \;:\; \exists!\, j_0 \equiv j_0(k) \in [d] \;:\; h_{j_0}(\eta_\alpha) \neq h_{j_0}(\eta_\beta) \tag{3.99}$$

with $\eta_\delta\, [\equiv \eta_\delta^{(k)}] := (\eta_1, \ldots, \eta_{k-1}, \delta, \eta_{k+1}, \ldots, \eta_d) \in \mathbb{R}^d$, $\delta \in \{\alpha, \beta\}$, which directly implies that h is monomial. Indeed, given (3.99) we for any fixed $k \in [d]$ obtain by the uniqueness of $j_k := j_0(k)$ (and its independence of the choice of α, β) that

$$\forall\, j \neq j_k \;:\; h_j(\eta_\alpha^{(k)}) = h_j(\eta_\beta^{(k)}) \text{ for all } \alpha, \beta \in \mathbb{R},$$

which (as η is arbitrary) implies that h_j (for $j \neq j_k$) does not depend on the k^{th}-coordinate of its argument, i.e. that for each $j \in [d]$ and any $k \in [d]$ with $j_0(k) \neq j$ it holds

$$\forall\, u \equiv (u_1, \ldots, u_d) \in \mathbb{R}^d \;:\; \alpha \mapsto h_j(u_1, \ldots, u_{k-1}, \alpha, u_{k+1}, \ldots, u_d) \text{ is constant.} \tag{3.100}$$

This, together with (3.99) and h being onto, implies that the map $\sigma_h : [d] \ni k \mapsto j_k \in [d]$ is a bijection,[33] yielding that $\forall\, j \in [d] : \exists!\, k \equiv k_j\, [:= \sigma_h^{-1}(j)] \in [d] \;:\; j = j_0(k_j)$ and hence (by (3.100))

$$\forall\, j \in [d] \;:\; h_j(u) = \tilde{h}_j(u_{\sigma(j)}) \quad \text{for} \quad \tilde{h}_j : \mathbb{R} \ni x \mapsto h_j(\eta_x^{(\sigma(j))}) \quad \text{and} \quad \sigma := \sigma_h^{-1} \in S_d$$

(where, as above, $\eta_x^{(k)} := (\eta_1, \ldots, \eta_{k-1}, x, \eta_{k+1}, \ldots, \eta_d)$ for $\eta \equiv (\eta_1, \ldots, \widehat{\eta_k}, \ldots \eta_d) \in \mathbb{R}^{d-1}$; note that in consequence of (3.100), $\tilde{h}_j$ does not depend on the choice of $\eta \in \mathbb{R}^{d-1}$). This shows that if (3.99) holds then h is monomial (in the sense of (3.23)), as claimed.

[33] Indeed: Assume $\exists\, k, \ell \in [d]$ such that $k \neq \ell$ (say $k < \ell$) and $j_k = j_\ell =: j'$. Then (3.100) implies that for each $j \in [d]$ with $j \neq j'$ it holds $h_j(u) \equiv h_j(u_1, \cdots, \widehat{u_k}, \cdots, \widehat{u_\ell}, \cdots, u_d)$, whence by the surjectivity of h (and the pigeonhole principle) there must be two *distinct* $j_1, j_2 \in [d] \setminus \{j'\}$ and $k_0 \in [d]$ such that the maps $\mathbb{R} \ni \alpha \mapsto h_{j_1}(\eta_\alpha^{(k_0)})$ and $\mathbb{R} \ni \alpha \mapsto h_{j_2}(\eta_\alpha^{(k_0)})$ are non-constant (otherwise there would be an index $\tilde{j} \in [d]$ with $h_{\tilde{j}} \equiv$ const., contradicting that h is onto). But (3.99) then implies that $j_1 = j_0(k_0) = j_2$ in contradiction to $j_1 \neq j_2$. This shows that $\sigma_h : [d] \to [d]$ is injective and thus an element of S_d (the symmetric group of degree d).

Let us now prove (3.99). Choose any $k \in [d]$ and $\alpha, \beta \in \mathbb{R}$ with $\alpha \neq \beta$. Since by the injectivity of h there clearly exists an index $j_0 \equiv j_0(k, \alpha, \beta, \eta) \in [d]$ with $h_{j_0}(\eta_\alpha^{(k)}) \neq h_{j_0}(\eta_\beta^{(k)})$, assertion (3.99) follows if we can show that (for each $k \in [d]$) such an index is uniquely determined and does not depend on the choice of α, β and η. To this end, take a Bernoulli-distributed random variable $C \sim \mathrm{Ber}(\frac{1}{2})$ together with an arbitrary $\eta \equiv (\eta_1, \ldots, \widehat{\eta_k}, \ldots, \eta_d) \in \mathbb{R}^{d-1}$ and define

$$U_\eta^{(k)} := (\alpha \cdot C + \beta \cdot (1 - C)) \cdot e_k + \sum_{j \neq k} \eta_j \cdot e_j. \tag{3.101}$$

The random variable $U_\eta^{(k)}$ then has independent components, whence by (3.98) it holds that

$$\frac{1}{2}\left[e^{\mathrm{i}\xi^\intercal h(\eta_\alpha)} + e^{\mathrm{i}\xi^\intercal h(\eta_\beta)}\right] = \mathbb{E}\left[e^{\mathrm{i}\xi^\intercal h(U_\eta^{(k)})}\right] = \frac{1}{2^d} \prod_{j \in [d]} \left[e^{\mathrm{i}\xi_j h_j(\eta_\alpha)} + e^{\mathrm{i}\xi_j h_j(\eta_\beta)}\right]. \tag{3.102}$$

The uniqueness of the above index $j_0 \equiv j_0(k, \alpha, \beta, \eta)$ will be derived by comparing both sides of (3.102). For a closer inspection of the right-hand side of (3.102), consider the number $\gamma \equiv \gamma(\alpha, \beta, \eta) \in [d]$ of those components of h whose values at η_α and η_β coincide,

$$\gamma := \sum_{j \in [d]} \mathbb{1}_{h_j(\eta_\beta)}(h_j(\eta_\alpha)).$$

For the functions $a_j \equiv a_j(\xi_j) := e^{\mathrm{i}\xi_j h_j(\eta_\alpha)}$ and $b_j \equiv b_j(\xi_j) := e^{\mathrm{i}\xi_j h_j(\eta_\beta)}$ it holds $a_j = b_j$ iff $h_j(\eta_\alpha) = h_j(\eta_\beta)$, and thus

$$\forall\, \nu \equiv (\nu_1, \ldots, \nu_d) \in \{0,1\}^{\times d} \ : \ c_\nu := \prod_{j \in [d]} a_j^{\nu_j} b_j^{1-\nu_j} = \kappa \cdot \prod_{j \in [d]_\gamma} a_j^{\nu_j} b_j^{1-\nu_j} \tag{3.103}$$

for $[d]_\gamma := \{j \in [d] \mid a_j \neq b_j\}$ and $\kappa := \prod_{j \in [d] \setminus [d]_\gamma} b_j$. Since $\#[d]_\gamma = d - \gamma$ by construction, and because for any two of the above multivariate functions (3.103) it holds $c_\nu = c_{\nu'}$ if and only if[34] $a_j^{\nu_j} b_j^{1-\nu_j} = c_\nu(0_{\xi_j}^{(j)}) = c_{\nu'}(0_{\xi_j}^{(j)}) = a_j^{\nu'_j} b_j^{1-\nu'_j}$ for all $j \in [d]$, we clearly have that

$$n(\alpha, \beta, \eta) := \left|\left\{ \prod\nolimits_{j \in [d]_\gamma} a_j^{\nu_j} b_j^{1-\nu_j} \text{ as in (3.103)} \,\middle|\, (\nu_1, \ldots, \nu_d) \in \{0,1\}^{\times d} \right\}\right| = 2^{|[d]_\gamma|} = 2^{d-\gamma}.$$

Note that since the family of functions $\{\xi \mapsto e^{\mathrm{i}\xi^\intercal u} \mid u \in \mathbb{R}^d\}$ are linearly independent over $\mathbb{R}$ (which is immediate by the injectivity of the Fourier transform (over tempered distributions)), the number $n(\alpha, \beta, \eta)$ counts the number of additively irreducible summands in the linear combination $\sum_{\nu \in \{0,1\}^{\times d}} c_\nu$. Hence and since (in light of the multinomial theorem and (3.103)) the identity (3.102) can be written as

$$\frac{1}{2}\left[e^{\mathrm{i}\xi^\intercal h(\eta_\alpha)} + e^{\mathrm{i}\xi^\intercal h(\eta_\beta)}\right] = \frac{\kappa}{2^d} \sum_{\nu \in \{0,1\}^{\times d}} \prod_{j \in [d]_\gamma} a_j^{\nu_j} b_j^{1-\nu_j}, \tag{3.104}$$

[34] In accordance with the above notation, we denote $0_{\xi_j}^{(j)} := \tilde{\eta}_{\xi_j}^{(j)}$ for $\tilde{\eta} \equiv 0 \in \mathbb{R}^{d-1}$.

a comparison of the number of (additively irreducible) terms (i.e. functions of ξ) on both sides of (3.104) is permissible and (due to $\alpha \neq \beta$ and h injective) yields that $2 = n(\alpha, \beta, \eta) = 2^{d-\gamma}$ and hence $\gamma = d - 1$. Recalling the definition of γ, this implies that

$$\exists !\, j_0 \equiv j_0(k, \alpha, \beta, \eta) \in [d] \;:\; h_{j_0}(\eta_\alpha) \neq h_{j_0}(\eta_\beta). \tag{3.105}$$

In order to from (3.105) derive assertion (3.99) and thus conclude the proof, it remains to show that j_0 is independent of the choice of α, β and η, i.e. that for any $k, \ell \in [d]$ it holds

$$\forall\, ((\alpha, \beta), \eta), ((\alpha', \beta'), \eta') \in \mathbb{R}^2_{\setminus\Delta} \times \mathbb{R}^{d-1} \;:\; j_0(k, \alpha, \beta, \eta) \neq j_0(\ell, \alpha', \beta', \eta') \implies k \neq \ell \tag{3.106}$$

where we denoted $\mathbb{R}^2_{\setminus\Delta} := \mathbb{R}^2 \setminus \{(x, x) \,|\, x \in \mathbb{R}\}$. We show (3.106) in two steps, first by establishing that the index j_0 from (3.105) is independent of $\alpha, \beta \in \mathbb{R}$ and then proving its independence of $\eta \in \mathbb{R}^{d-1}$. To this end, let $k \in [d]$ and $\eta \in \mathbb{R}^{d-1}$ be fixed and choose any $\alpha', \alpha, \beta \in \mathbb{R}$ with $\alpha \neq \beta$ and $\alpha' \neq \alpha$. Assume now that $j_0 := j_0(k, \alpha, \beta, \eta) \neq j_0(k, \alpha', \beta, \eta) =: j_0'$. Then by (3.105) it holds

$$\begin{aligned} h_{j_0}(\eta_\alpha) \neq h_{j_0}(\eta_\beta) \quad &\text{and} \quad h_{j_0'}(\eta_\alpha) = h_{j_0'}(\eta_\beta), \qquad \text{as well as} \\ h_{j_0'}(\eta_{\alpha'}) \neq h_{j_0'}(\eta_\beta) \quad &\text{and} \quad h_{j_0}(\eta_{\alpha'}) = h_{j_0}(\eta_\beta), \end{aligned}$$

which implies that

$$h_{j_0}(\eta_{\alpha'}) \neq h_{j_0}(\eta_\alpha) \;\text{ and }\; h_{j_0'}(\eta_{\alpha'}) \neq h_{j_0'}(\eta_\alpha). \tag{3.107}$$

But since $j_0 \neq j_0'$ by assumption, statement (3.107) (due to $\alpha' \neq \alpha$) contradicts the uniqueness of $j_0(k, \alpha', \alpha, \eta)$ asserted by (3.105), whence we must have $j_0(k, \alpha, \beta, \eta) = j_0(k, \alpha', \beta, \eta)$. Hence and since clearly $j_0(k, \alpha, \beta, \eta) = j_0(k, \beta, \alpha, \eta)$ for all $\alpha, \beta \in \mathbb{R}$, we for any $(\alpha, \beta), (\alpha', \beta') \in \mathbb{R}^2_{\setminus\Delta_\mathbb{R}}$ find

$$j_0(k, \alpha, \beta, \eta) = j_0(k, \alpha', \beta, \eta) = j_0(k, \beta, \alpha', \eta) = j_0(k, \beta', \alpha', \eta) = j_0(k, \alpha', \beta', \eta),$$

which shows that $j_0(k, \alpha, \beta, \eta) = j_0(k, \eta)$. In particular, (3.105) now reads as follows: for each $k \in [d]$ and any $\eta \in \mathbb{R}^{d-1}$,

$$\exists !\, j_0 \equiv j_0(k, \eta) \in [d] \;:\; \text{the function } \; h_{j_0}^{(k,\eta)} \;:\; \mathbb{R} \ni x \;\mapsto\; h_{j_0}(\eta_x^{(k)}) \;\text{ is non-constant,} \tag{3.108}$$

and (3.105) clearly implies that any such non-constant function $h_{j_0}^{(k,\eta)}$ is also injective.
To now prove that the index $j_0 \equiv j_0(k, \eta)$ in (3.108) does not depend on η, suppose that

$$\exists\, \eta_1, \eta_2 \in \mathbb{R}^{d-1} \text{ with } j_1 := j_0(k, \eta_1) \neq j_0(k, \eta_2) =: j_2. \tag{3.109}$$

Then with (3.108) and since h is assumed to be a C^1-diffeomorphism, it holds for each $x, y \in \mathbb{R}$ and $\tilde{\eta} := (\eta_1)_x^{(k)}$, $\tilde{\tilde{\eta}} := (\eta_2)_x^{(k)}$ that

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}x} h_{j_1}^{(k,\eta_1)}(x) \neq 0 \quad \text{and} \quad \frac{\mathrm{d}}{\mathrm{d}x} h_{j_2}^{(k,\eta_2)}(y) \neq 0, \quad \text{and} \\ &\exists\, k_1, k_2 \in [d] \;:\; \partial_{k_1} h_{j_2}(\tilde{\eta}_y^{(k_1)}) \neq 0 \quad \text{and} \quad \partial_{k_2} h_{j_1}(\tilde{\tilde{\eta}}_y^{(k_2)}) \neq 0, \end{aligned} \tag{3.110}$$

where we denoted $u_y^{(\ell)} := (u_1, \ldots, u_{\ell-1}, y, u_{\ell+1}, \ldots, u_d)$ for $u \in \{\tilde{\eta}, \tilde{\tilde{\eta}}\}$ and $\ell \in [d]$.

Indeed, since h is a C^1-diffeomorphism its Jacobian $J_h \equiv (\nabla h_i)_{i \in [d]}$ is globally invertible, whence the k^{th}-column of $J_h((\eta_\nu)_x^{(k)})$ $(\nu = 1, 2)$ — which by (3.108) is equal to $\partial_k h_{j_\nu}((\eta_\nu)_x^{(k)}) \cdot e_{j_\nu}$ — must be non-zero, implying the first part of (3.110). (The second part of (3.110) follows similarly by noting that otherwise the j_2^{th}-row of $J_h(\tilde{\eta}_y^{(k_1)})$ (that is, the gradient vector $\nabla h_{j_2}(\tilde{\eta}_y^{(k_1)})$) would be zero (likewise for j_1), again contradicting the global invertibility of J_h.) Note further that by (3.108) and $j_1 \neq j_2$ we have $k_\nu \neq k$. Let now $x \in \mathbb{R}$ be fixed and denote $\gamma_\nu := \operatorname{sgn}(\frac{\mathrm{d}}{\mathrm{d}x} h_{j_\nu}^{(k,\eta_\nu)}(x))$ as well as $\gamma_1' := \operatorname{sgn}(\partial_{k_1} h_{j_2}((\eta_1)_x^{(k)}))$ and $\gamma_2' := \operatorname{sgn}(\partial_{k_2} h_{j_1}((\eta_2)_x^{(k)}))$ (all of these signs are independent of $x \in \mathbb{R}$ due to (3.110) and the continuity of the partial derivatives of h) as well as $\tilde{\eta} := (\eta_1)_x^{(k)}$ and $\tilde{\tilde{\eta}} := (\eta_2)_x^{(k)}$, and for

$$\begin{aligned} \varphi_{j_1} &: \mathbb{R} \ni \beta \;\mapsto\; h_{j_1}(\tilde{\tilde{\eta}}_1, \ldots, \tilde{\tilde{\eta}}_{k_2-1}, \beta, \tilde{\tilde{\eta}}_{k_2+1}, \ldots, \tilde{\tilde{\eta}}_d) \quad \text{and} \\ \varphi_{j_2} &: \mathbb{R} \ni \alpha \;\mapsto\; h_{j_2}(\tilde{\eta}_1, \ldots, \tilde{\eta}_{k_1-1}, \alpha, \tilde{\eta}_{k_1+1}, \ldots, \tilde{\eta}_d) \end{aligned}$$

define the function

$$\tilde{h} \;:\; \mathbb{R}^2 \ni (\alpha, \beta) \;\mapsto\; \begin{pmatrix} \gamma_1 \cdot h_{j_1}^{(k,\eta_1)}(\alpha) - \gamma_2' \cdot \varphi_{j_1}(\beta) \\ \gamma_1' \cdot \varphi_{j_2}(\alpha) + \gamma_2 \cdot h_{j_2}^{(k,\eta_2)}(\beta) \end{pmatrix} =: \begin{pmatrix} \tilde{h}_1(\alpha, \beta) \\ \tilde{h}_2(\alpha, \beta) \end{pmatrix}.$$

Then $\tilde{h}$ is continuously differentiable, and since for the functional determinant of $\tilde{h}$ we have

$$\det(J_{\tilde{h}})(\alpha, \beta) = \Big[\gamma_1 \cdot \frac{\mathrm{d}}{\mathrm{d}x} h_{j_1}^{(k,\eta_1)}(\alpha)\Big] \cdot \Big[\gamma_2 \cdot \frac{\mathrm{d}}{\mathrm{d}x} h_{j_2}^{(k,\eta_2)}(\beta)\Big] + \Big[\gamma_2' \cdot \varphi_{j_1}'(\beta)\Big] \cdot \Big[\gamma_1' \cdot \varphi_{j_2}'(\beta)\Big] \overset{(3.110)}{>} 0$$

(with $\varphi_{j_\nu}' := \frac{\mathrm{d}}{\mathrm{d}x} \varphi_{j_\nu}$, $\nu = 1, 2$) for each $(\alpha, \beta) \in \mathbb{R}^2$, the inverse function theorem implies that $\tilde{h}$ is a local diffeomorphism. Moreover, it is clear by construction (and (3.108)) that $\tilde{h}$ inherits from h the property (3.96) of preserving the componental independence of any $\mathbb{R}^2$-valued random variable. This implies that (upon for any fixed $(\alpha_0, \beta_0) \in \mathbb{R}^2$ restricting $\tilde{h}$ to the open set $U_{(\alpha_0, \beta_0)} \subset \mathbb{R}^2$ on which it is diffeomorphic and adapting (in the obvious way) the definition of the random variables $\{U_\eta^{(k)} \mid (3.101)\}$ w.r.t. this restriction) the derivation of (3.105) also applies to $\tilde{h}$, whence for any $(\alpha_0, \beta_0) \in \mathbb{R}^2$, each $\tilde{k} = 1, 2$ and any $\alpha, \beta, \tilde{\eta} \in \mathbb{R}$ with $\tilde{\eta}_\alpha^{(k)}, \tilde{\eta}_\beta^{(k)} \in U_{(\alpha_0, \beta_0)}$ we must have

$$\exists!\, \tilde{j}_0 \equiv \tilde{j}_0(\tilde{k}, \alpha, \beta, \tilde{\eta}) \in \{1, 2\} \;:\; \tilde{h}_{\tilde{j}_0}(\tilde{\eta}_\alpha^{(\tilde{k})}) \neq \tilde{h}_{\tilde{j}_0}(\tilde{\eta}_\beta^{(\tilde{k})}). \tag{3.111}$$

By definition of $\tilde{h}$, however, we for each $(\alpha, \beta) \in \mathbb{R}^2$ find that (cf. (3.110))

$$\partial_1 \tilde{h}_1(\alpha, \beta) \neq 0 \;\; \text{and} \;\; \partial_1 \tilde{h}_2(\alpha, \beta) \neq 0,$$

which (for $(\tilde{k}, \tilde{\eta}) := (1, \beta_0)$ and some α, β close to α_0) clearly contradicts the uniqueness assertion of (3.111). Thus assumption (3.109) is overruled, proving (3.106) and hence (3.99). The proposition thus follows by the arguments given in the beginning of this proof. $\square$

3.7 Nonlinear ICA for Discrete-Time Signals

As detailed here, our approach towards the identifiability of nonlinearly mixed stochastic processes also covers the case of discrete-time signals with almost no further modifications.

Assume throughout that $X_* \equiv (X_j)_{j\in\mathbb{Z}}$ is some discrete time-series in $\mathbb{R}^d$ such that

$$X_* \;=\; f(S_*) \;\equiv\; (f(S_j))_{j\in\mathbb{Z}} \tag{3.112}$$

for some IC discrete time-series $S_* \equiv (S_j)_{j\in\mathbb{Z}}$ in $\mathbb{R}^d$ and $f \in C^{2,2}(D_{S_*};\mathbb{R}^d)$. Here, a time series $Y_* \equiv (Y_j)_{j\in\mathbb{Z}}$ in $\mathbb{R}^d$, with $Y_j \equiv (Y_j^1, \ldots, Y_j^d)$ for each $j \in \mathbb{Z}$, is called *IC* if its componental time-series $(Y_j^1)_{j\in\mathbb{Z}}, \ldots, (Y_j^d)_{j\in\mathbb{Z}}$ are mutually independent.

Denote further $D_{Y_*} := \overline{\bigcup_{j\in\mathbb{Z}} \operatorname{supp}(Y_j)}^{|\cdot|_2}$ for the *spatial support* of Y_*, and write $\Delta_2(\mathbb{Z}) := \{(j_1, j_2) \in \mathbb{Z}^2 \mid j_1 < j_2\}$ for the set of all strictly ordered pairs of integers.

Definition 3.7.1 ($\bar{\alpha}, \bar{\beta}, \bar{\gamma}$-Contrastive)**.** A discrete time-series $S_* \equiv (S_j)_{j\in\mathbb{Z}}$ in $\mathbb{R}^d$ with spatial support D_{S_*} will be called *$\bar{\alpha}$-contrastive* if S_* is IC and there exists $\mathcal{P} \subseteq \Delta_2(\mathbb{Z})$ together with a collection $(D_{\mathfrak{p}})_{\mathfrak{p}\in\mathcal{P}}$ of open subsets in $\mathbb{R}^d$ such that

(i) the union $\bigcup_{\mathfrak{p}\in\mathcal{P}} D_{\mathfrak{p}}$ is dense in D_{S_*}, and

(ii) for each $(i, (j_1, j_2)) \in [d] \times \mathcal{P}$, the vector $(S^i_{j_1}, S^i_{j_2})$ is C^2-distributed with density $\zeta^i_{j_1,j_2}$ such that

$$\begin{aligned} &\zeta^i_{j_1,j_2}\big|_{D^{\times 2}_{(j_1,j_2)}} \quad \text{is regularly non-separable for all } i \in [d], \quad \text{and} \\ &\zeta^i_{j_1,j_2}\big|_{D^{\times 2}_{(j_1,j_2)}} \quad \text{is almost everywhere non-Gaussian for all but at most one } i \in [d] \end{aligned}$$

(cf. Definition 3.3.2). The notions of $\bar{\beta}$*- and* $\bar{\gamma}$*-contrastive* time series are defined in analogous adaptation of Definition 3.3.5.

Theorem 3.7.2. *For X_* and S_* as in* (3.112) *with spatial supports D_{X_*} and D_{S_*} respectively, let the time-series S_* in be $\bar{\alpha}$-, $\bar{\beta}$- or $\bar{\gamma}$-contrastive. Then, for any transformation h which is C^2-invertible on some open superset of D_{X_*}, we have with probability one that:*

$$(h \circ f)\big|_{\tilde{Z}} \;\in\; \mathrm{DP}_d(\tilde{Z}),\ \forall\, Z \subseteq D_{X_*} \ \textit{connected} \qquad \textit{if and only if} \qquad h(X_*) \ \textit{is IC}, \tag{3.113}$$

where for any connected subset Z of D_{X_*} we denoted $\tilde{Z} := f^{-1}(Z)$.

Proof. Let S_* be $\bar{\alpha}$-contrastive, and $h \in C^{2,2}(D_{X_*}; \mathbb{R}^d)$ be such that $h(X_*)$ is IC. Then

$$(h \times h)(X_{j_1}, X_{j_2}) \text{ is IC} \quad \text{for any fixed } (j_1, j_2) \in \mathcal{P},$$

which in consequence of Def. 3.7.1 (ii) implies that, as derived in the proof of Thm. 3.3.3,

$$\text{the Jacobian of} \quad \varrho := h \circ f \quad \text{is monomial on } D_{(j_1, j_2)}.$$

The equivalence (3.113) thus follows from Definition 3.7.1 (i) and Lemma 3.4.1. The case of S_* being $\bar{\beta}$- or $\bar{\gamma}$-contrastive follows similarly via Theorem 3.3.7. □

Since a discrete time-series Y_* in $\mathbb{R}^d$ is IC if and only if its piecewise-linear interpolation[35] $\hat{Y}_*$ is IC in $\mathcal{C}_d$, the assertion of Theorem 4.2.3 remains valid as stated[36] if $(S, \alpha, \beta, \gamma, X)$ is replaced by $(S_*, \bar{\alpha}, \bar{\beta}, \bar{\gamma}, X_*)$ and the argument $h(X)$ in (4.9) is replaced by the piecewise-linear interpolation of $h(X_*)$. This shows that the identifiability theory of Sects. 3.3 and 4.2 directly applies to the discrete-time setting (3.112), as detailed in the following section.

3.7.1 Identifiability in the Discrete-Time Case

From Lemma 3.4.1 and Theorems 3.3.3, 3.3.7, we saw that the inversion (up to monomial ambiguity) of the transformation $f : D_S \to D_X$ given (nothing but) X is possible if there is a dense open subset $\mathcal{D}$ in D_S which is 'identifiability enforcing' in the sense that

$$J_{h \circ f}(u) \in \mathrm{M}_d \quad \text{for each } u \in \mathcal{D} \tag{3.114}$$

for any $h \in C^{2,2}(D_X)$ such that $h(X)$ is IC; see e.g. the proof of Theorem 3.3.3. As shown in the proofs of these theorems, such a set $\mathcal{D}$ can be induced by a subset of C^2-regular

$$\text{distributions of} \quad \{(S_s, S_t) \mid (s,t) \in \Delta_2(\mathbb{I})\} \tag{3.115}$$

if the source process S is α-, β- or γ-contrastive in the sense of Defs. 3.3.2 and 3.3.5. (The associated examples for $\mathcal{D}$ in these cases are $\bigcup_{\mathfrak{p} \in \mathcal{P}} D_{\mathfrak{p}}$ (Def. 3.3.2), and $\mathrm{int}(D_S)$ or $\mathcal{U}$ (Def. 3.3.5).) Now importantly, this approach towards identifiability makes no essential use of S being time-continuous: Both Definitions 3.3.2, 3.3.5 *and* their consequential derivations of (3.114) only use that

$$\begin{gathered} S = (S_t)_{t \in \mathbb{I}} \quad \text{is a family of random variables } S_t \equiv (S_t^i) \text{ in } \mathbb{R}^d \\ \text{with } \mathbb{I} \text{ a totally ordered subset of } \mathbb{R} \text{ and such that} \\ \text{the families} \quad (S_t^1)_{t \in \mathbb{I}}, \ldots, (S_t^d)_{t \in \mathbb{I}} \quad \text{are statistically independent,} \end{gathered} \tag{3.116}$$

[35] ... along any (countable) dissection of, say, $[0,1]$.

[36] With the addition that in (4.9), the monomial transformations α with $h(X) = \alpha(S)$ then depend on (j, ω) via the connected component that the given realisation of $S_j = S_j(\omega)$ is contained in (details below) — respectively (if the connected components of D_S are not assumed convex): via each convex component of D_{S_*} that the given realisation of S_j intersects with positive probability, see the beginning of Section 3.3 and Section 3.4.1.

as a quick inspection of the proofs of Theorems 3.3.3, 3.3.7 shows. For the special case where $\mathbb{I}$ is countable – i.e. $S = (S_t)_{t\in\mathbb{I}} \equiv (S_j)_{j\in\mathbb{J}}$ ($\mathbb{J} \subseteq \mathbb{Z}$) is a discrete time-series – we may thus 'abstract (3.6) from the topology on $\mathbb{I}$' without losing the power of our approach, as emphasised below. This means that we may simply consider the discrete 'lattice' of joint laws given by

$$\text{the distributions of} \quad \{(S_{j_1}, S_{j_2}) \mid (j_1, j_2) \in \Delta_2(\mathbb{Z})\} \tag{3.115'}$$

in lieu of the uncountable collection (3.115). In this case then still, regular dense subsets $\mathcal{D}$ of D_S with the desired property (3.114) can be induced from a C^2-distributed subselection[37] of (3.115') *by the exact same argumentation* that we have used in the proofs of Theorems 3.3.3 and 3.3.7, as it is evident that these proofs do *not* involve the topology on $\mathbb{I}$ but only its order. The associated (by direct analogy to the continuous-time case) premises for the existence of such a subselection of (3.115') are formulated as Definition 3.7.1, which thus appears as the natural 'discretization' of Definitions 3.3.2 and 3.3.5.

There are three points in our theory where the assumed continuity of the sample paths of S does make a subtle but not entirely trivial difference:

(i) From the 'pre-identifiability' property (3.114) for a (discrete- or continuous-time) stochastic process S of the general form (3.116), we obtain by Lemma 3.4.1 (ii) (recalling Definition 3.3.1) that the residual $h \circ f$ is monomial *on every connected component* of D_S. Consequently:

  (a) If the source S is time-continuous, then its sample path $S(\omega) \equiv (S_t(\omega))_{t\in\mathbb{I}}$ is a connected subset of D_S with probability one (cf. Lemma 3.2.5 (ii)). The identifiability equations (3.25) and (3.44) of Theorems 3.3.3 and 3.3.7 then state that, almost surely, the components of (the sample paths of) the estimated source $\hat{S}(\omega) := h(X(\omega))$ and those of the original source $S(\omega)$ coincide up to a pathwise-fixed permutation τ and some monotone scaling $(\alpha_1, \cdots, \alpha_d)$. That is, it states that with probability one there are τ and (α_i) such that

$$\hat{S}^i_t = \alpha_i(S^{\tau(i)}_t) \quad \textit{for each } t \in \mathbb{I} \quad (i \in [d]) \tag{3.117}$$

  where both τ and (α_i) are uniquely determined by (X and) h and ω via the connected component of D_S that the source realisation $S(\omega)$ is contained in.

[37] That is, a set of (the distributions of) vectors $\{(S_k, S_\ell) \mid (k, \ell) \in \mathcal{P}\}$, for some $\mathcal{P} \subseteq \Delta_2(\mathbb{Z})$, such that (S^i_k, S^i_ℓ) is C^2-distributed (cf. Definition 3.2.6) for each $(k, \ell) \in \mathcal{P}$ and each $i \in [d]$. (As the components of (S_j) are mutually independent, (S^i_k, S^i_ℓ) is C^2-distributed *for each* $i \in [d]$ iff the full vector (S_k, S_ℓ) is C^2-distributed.)

(b) If the source S is time-discrete, then its realisations $S(\omega) \equiv (S_j(\omega))_{j\in\mathbb{Z}}$ are generally not connected in $\mathbb{R}^d$ and might thus be spread over different connected components of D_S, again almost surely. Hence in this case we have with probability one that

$$\hat{S}_t^i = \alpha_i^{(t)}(S_t^{\tau_t(i)}) \quad \text{for each } t \in \mathbb{I} \quad (i \in [d]) \tag{3.118}$$

where the permutations τ_t and scalings $(\alpha_i^{(t)})$ are no longer pathwise-fixed but do now depend on (X and) h and ω *and* t via the connected component of D_S that each fixed-time realisation $S_t(\omega)$ is contained in;[38] this can be read off Theorem 3.7.2.

Clearly the 'minimal deviations' (3.117) and (3.118) between $\hat{S}$ and S coincide if D_S is connected, but if it is not connected they are generally different.[39]

So throughout the identifiability sections of this chapter (Sections 3.2 to 3.5), the assumption of sample continuity of S is merely a convenience assumption which we included because of its sufficiency for the pathwise (t-independent) identification (3.117). Except for this distinction between the 'connected case' (3.117) and the 'disconnected case' (3.118) our identifiability results (Theorems 3.3.3 and 3.3.7) apply without further changes, as summarised in Theorem 3.7.2.

Another modification for the discrete-time case, this time of a purely technical nature, concerns the optimisation in Theorem 4.2.3, specifically the usability of the contrast $\bar{\kappa}_{\mathrm{IC}}$:

(ii) By its definition the function $\bar{\kappa}_{\mathrm{IC}}$ can only take time-continuous processes as its arguments, but for those it is only the *order* of their time-indexed values that matters (cf. Lemma 2.3.5 (iii)). Consequently:

(a) If the source S is time-continuous then so is the candidate transformation $h(X)$, making $\bar{\kappa}_{\mathrm{IC}}(h(X))$ well-defined and Theorem 4.2.3 readily applicable as stated.

(b) If the source S is time-discrete, say $S = (S_j)_{j\in\mathbb{Z}}$, then we may perform the injection $S \mapsto \overline{S}$ where $\overline{S}$ is the piecewise-linear interpolation of S along any (fixed) strictly ordered bounded subset $\mathcal{I} \equiv \{t_j\} \subset \mathbb{R}$, see Section 4.7.3.[40] Likewise, let

[38] That is, $\big(\tau_t, (\alpha_i^{(t)})\big) = \big(\tau_r, (\alpha_i^{(r)})\big) \; \big[\equiv \big(\tau_r(\omega), \alpha_1^{(r)}(\omega), \ldots, \alpha_d^{(r)}(\omega)\big)\big]$ if $S_t(\omega)$ and $S_r(\omega)$ are in the same connected component of D_S, and possibly $\big(\tau_t, (\alpha_i^{(t)})\big) \neq \big(\tau_r, (\alpha_i^{(r)})\big)$ if not.

[39] Notice that, as specified in Prop. 4.8.2, the test statistic (4.102) from Definition 4.8.1 is originally tailored to the connected case (3.117), but upon straightforward modification it may of course be used for the case (3.118) as well.

[40] Remember that the choice of $\mathcal{I}$ is arbitrary up to order and cardinality (Lemma 2.3.5 (iii)), that is the interpolation $\overline{Y}$ only needs to preserve the time order of the data points $(Y_j) \equiv Y$; individual values $\overline{Y}_t$ for $t \notin \mathcal{I}$ are irrelevant. This also ensures (3.119) is well-defined, i.e. independent of the choice of interpolant of its arguments.

us for any $Y = (Y_j)_{j\in\mathbb{Z}}$ denote by $\overline{Y}$ the piecewise-linear interpolation of Y along $\mathcal{I}$, and define by

$$\bar{\kappa}_{\mathrm{IC}}(Y) := \bar{\kappa}_{\mathrm{IC}}(\overline{Y}) \tag{3.119}$$

an extension of $\bar{\kappa}_{\mathrm{IC}}$ to discrete-time processes. Note that $\overline{Y}$ is a time-continuous process with $Y_j = \overline{Y}_{t_j}$ for each $j \in \mathbb{Z}$, and further that Y is IC iff $\overline{Y}$ is IC. Thus if $X_j = f(S_j)$ on $\mathbb{Z}$ then also $\overline{X}_{t_j} = f(\overline{S}_{t_j})$ on $\mathcal{I}$, and if $S = (S^i_j)_j$ is $\bar{\alpha}$-, $\bar{\beta}$-, or $\bar{\gamma}$-contrastive then by Theorem 3.7.2,

$$\boxed{h_i(X_j) = \alpha_i^{(j)}(S_j^{\tau(i)}) \ \ (\forall\, j \in \mathbb{Z})} \quad \text{iff} \quad \bar{\kappa}_{\mathrm{IC}}\big((h(X_j))_j\big) = 0 \tag{3.120}$$

for any $h = (h_1, \ldots, h_d) \in C^{2,2}(D_X)$, where the boxed equation holds in the sense of (3.118). This is the discrete-time version of Theorem 4.2.3.

In summary, we emphasize that if the source S is time-discrete with its spatial support admitting a dense open subset $\mathcal{D}$ as in (3.114) — which, for instance, will be induced by the countably indexed reservoir of joint distributions (3.115') if S is $\bar{\alpha}$-, $\bar{\beta}$- or $\bar{\gamma}$-contrastive — then our ICA approach [Theorems 3.3.3, 3.3.7 and 4.2.3] is directly applicable [in the form of Theorem 3.7.2 and (3.120)] with no additional subtleties other than points (i) and (ii) above.[41]

Finally, let us explicate the practically important identifiability situation where the given data is a discrete-time approximation ("observation") of a continuous-time process $X = f(S)$, for S some α-, β- or γ-contrastive source in $\mathbb{R}^d$.

(iii) In this last setting, the given data is of the form $X_{\mathcal{I}} := (X_t)_{t\in\mathcal{I}}$ for $\mathcal{I}$ discrete. For the general case that the corresponding discrete time-series $S_{\mathcal{I}} := (S_t)_{t\in\mathcal{I}}$ is not itself $\bar{\alpha}$-, $\bar{\beta}$- or $\bar{\gamma}$-contrastive, we may not be able to exactly (i.e. up to minimal ambiguity) recover $S_{\mathcal{I}}$ from $X_{\mathcal{I}}$ as we did above. However, we are still guaranteed the asymptotic identification

$$\forall\, \varepsilon > 0 \ : \ \exists\, \delta > 0 \quad \text{s.t.} \quad \sup_{t\in\mathcal{I}} \big|\hat{\theta}_{\mathcal{I}}(X_t) - \tilde{\alpha}(S_t)\big| \ \leq \varepsilon \qquad \text{if} \ \ \|\mathcal{I}\| \leq \delta \tag{3.121}$$

for some monomial $\tilde{\alpha} \in \mathrm{DP}_d(D_S)$ depending on $\mathcal{I}$ and on the realisation of S (via the connected component of D_S that this realisation is contained in), cf. point (i), with $\hat{\theta}_{\mathcal{I}} \in \arg\min_{\theta\in\Theta} \bar{\kappa}_{\mathrm{IC}}\big(\theta(X_{\mathcal{I}})\big)$ [in the sense of (3.119) and (4.40)] and Θ as in Theorem 4.4.13, and where (3.121) holds on some ($\mathcal{I}$-dependent) $\mathbb{P}$-full set. (If Θ admits a unique minimizer of $\bar{\kappa}_{\mathrm{IC}}$, then the above $\tilde{\alpha}$ is independent of $\mathcal{I}$.) This is a special case of Theorem 4.4.13 for $(m_0, k, T) = (\infty, k, \infty)$, see also its proof in Section 4.4.6.

[41] In particular, the interpolation $Y \mapsto \overline{Y}$ ('discrete to continuous') used in (3.119) only serves to find dependence-minimising transformations h via $\min_h \bar{\kappa}_{\mathrm{IC}}(h(Y))$. This interpolation is thus merely 'operational' (for the use of $\bar{\kappa}_{\mathrm{IC}}$) and not related to the identifiability of S itself; in particular, it does not add any geometrical or topological intricacies or complications to the latter.

Chapter 4

Nonlinear ICA via Signature Cumulants

The identifiability theory from the previous chapter allows us to reformulate the problem of nonlinear blind source separation as an easy-to-state optimisation problem that involves the minimisation of statistical dependence between multiple stochastic processes, see Theorem 4.2.3. Unlike for vector-valued data, statistical dependence between stochastic processes can manifest itself inter-temporally, in the sense that different coordinates of the processes may exhibit statistical dependencies both instantaneously and over different points in time. In this chapter, we propose to quantify such complex dependency relations by using so-called signature cumulants [19] as objective functions. These signature cumulants can be seen as generalising the concept of cumulants from vector-valued data to path-valued data. Analogous to classical cumulants, signature cumulants then provide a graded, parsimonious, and efficiently computable quantification of the degree of statistical (in)dependence between stochastic processes. Joined with our optimisation approach, this combines to a widely applicable new and robust statistical method for the nonlinear blind source separation of time-dependent signals, see Theorem 4.2.3 and Section 4.6.

As this chapter will show further, this method is consistent with respect to time discretization and sample size (Section 4.4). When applying our methodology in practice, the following issues arise: firstly, although the underlying stochastic model is often formulated in continuous time, in practice one usually has access to time-discretized samples only, often taken over non-equally spaced time grids. Secondly, oftentimes only a single (time-discretized) sample path of the process is available rather than many independent realisations, for example in the classical cocktail party problem. We address both of these issues and show that our method is statistically consistent even if only a single, time-discretized and finite sample of the observable is given, see Theorem 4.4.13. This is also the setting in which our numerical experiments are carried out in Section 4.8.

This chapter is structured as follows. Following the introduction of signature cumulants as the log-compressed signature representation for the laws of stochastic processes (Section 4.1), we enable the practical applicability of our ICA-approach by proposing a novel independence criterion for time-dependent data (Section 4.2) that leads to a practical and statistically consistent separation algorithm (Sections 4.4 – 4.7) whose practical utility we demonstrate in a series of numerical experiments (Section 4.8). The chapter ends with a brief conclusion and an outlook on future directions (Section 4.9).

Remark 4.0.1. In this chapter, we restrict our exposition to stochastic processes whose sample paths are smooth [i.e., of bounded variation[1]], and further assume that the expected signature of these processes exists and characterizes their law. These assumptions can be avoided by using rough integration and tensor normalization, but since this requires background in rough path theory and is not central to our methodology, we simply refer the interested reader to [54, 104] and [33, 34], respectively. Let further $\mathbb{I} = [0,1]$ wlog.

4.1 Signature Cumulants

Many results in statistics, including Corollary 3.1.3 via (3.3), are based on the well-known facts that laws of $\mathbb{R}^d$-valued random variables are often characterised by their moments, and that statistical independence turns into simple algebraic relations when expressed in terms of cumulants. Our main object of interest are $\mathcal{C}_d$-valued random variables (stochastic processes), for which the so-called expected signature [33] provides a natural generalisation of the classical moment sequence, see Section 2.2.3. Similar to classical moments, these signature moments form multi-indexed collections of numbers that can characterize the laws of stochastic processes. Similar still, upon their 'logarithmic compression' (Sect. 2.3.2) these number collections give rise to signature cumulants that quantify the statistical dependencies within multidimensional stochastic processes (that is, between their coordinates and over time).

Denote by $[d]^\star := \bigcup_{m\geq 0}[d]^{\times m}$ the set of all multi-indices[2] with entries in $[d] = \{1, \ldots, d\}$. We restate Definition 2.2.9 for convenience.

Definition 4.1.1 (Expected Signature)**.** For $Y = (Y_t^1, \cdots, Y_t^d)_{t\in[0,1]}$ a stochastic process in $\mathbb{R}^d$ with sample-paths of bounded variation, the collection of real numbers (if it exists)

[1] A path $x = (x_t)_{t\in[0,1]} \in \mathcal{C}_d$ is called *of bounded variation* if its variation norm $\|x\|_{1\text{-var}} := |x_0| + \sup\sum|x_{t_{i+1}} - x_{t_i}|$ is finite, where the supremum is taken over all finite partitions $\{0 \leq t_1 \leq \cdots \leq t_n \leq 1\}$ $(n \in \mathbb{N})$ of $[0,1]$; cf. also Definition 2.1.4 and Section 4.7.3.

[2] We define $[d]^{\times 0} := \{\emptyset\}$ with $\emptyset$ the empty set, and let $\{k\}^\star := \bigcup_{m\geq 0}\{k\}^{\times m}$ $\big(= \{\emptyset, k, kk, kkk, \ldots\}\big)$.

$\mathfrak{S}(Y) := (\sigma_{\boldsymbol{i}}(Y))_{\boldsymbol{i}\in[d]^\star}$ defined by the expected iterated Stieltjes integrals

$$\sigma_{\boldsymbol{i}}(Y) := \mathbb{E}\Big[\int_{0\le t_1\le t_2\le\cdots\le t_m\le 1} \mathrm{d}Y_{t_1}^{i_1}\mathrm{d}Y_{t_2}^{i_2}\cdots\mathrm{d}Y_{t_m}^{i_m}\Big] \quad \text{for} \quad \boldsymbol{i} = (i_1,\ldots,i_m), \tag{4.1}$$

with $\sigma_\emptyset(Y) := 1$, is called *the expected signature of* Y.

The expected signature is to a stochastic process roughly what the sequence of moments is to a vector-valued random variable (Sect. 2.2.3), and analogous to the case of classical moments, for many statistical purposes the concept of cumulants is better suited (Sect. 2.3.2). This leads to the notion of signature cumulants [19] below.

We recall the definition (2.30) of the log transform on formal power series (and refer to Sections 2.3.2 and 4.1.3 for further motivation of the following definition.)

Definition 4.1.2 (Signature Cumulants)**.** For Y a stochastic process in $\mathbb{R}^d$ with sample-paths of bounded variation, the collection of real numbers[3]

$$(\kappa_{\boldsymbol{i}}(Y))_{\boldsymbol{i}\in[d]^\star} := \log[\mathfrak{S}(Y)] \tag{4.2}$$

is called *the signature cumulant* of Y. We further define

$$\bar\kappa_{\boldsymbol{i}}(Y) := \frac{\kappa_{\boldsymbol{i}}(Y)}{\kappa_{11}(Y)^{\eta_1(\boldsymbol{i})/2}\cdot\ldots\cdot\kappa_{dd}(Y)^{\eta_d(\boldsymbol{i})/2}} \quad \text{for} \quad \boldsymbol{i} = (i_1,\ldots,i_m)\in[d]^\star, \tag{4.3}$$

where $\eta_\nu(\boldsymbol{i})$ denotes the number of times the index-value ν appears in $\boldsymbol{i}$. We refer to $(\bar\kappa_{\boldsymbol{i}}(Y))_{\boldsymbol{i}\in[d]^\star}$ as the *standardized signature cumulant* of Y.

Remark 4.1.3. The signature cumulant of a process Y gives an efficiently computable [95], informationally condensed and hierarchically graded [cf. (2.22), (2.33)] compression of the statistical information contained in (the distribution of) Y [cf. Sects. 2.2.3 and 2.3.2], which enjoys a broad variety of excellent practical and theoretical features [34]. Note that the (classical) cumulants of a random vector Z in $\mathbb{R}^d$ read

$$\kappa^Z_{i_1\cdots i_m} \;=\; \langle \pi_{\mathrm{Sym}}(\log[\mathfrak{m}(Z)]),\, \mathtt{i}_1\cdots\mathtt{i}_m\rangle \tag{4.4}$$

(for $\pi_{\mathrm{Sym}}(\mathtt{i}_1\cdots\mathtt{i}_m) := \sum_{\tau\in S_m}\mathtt{i}_{\tau(1)}\cdots\mathtt{i}_{\tau(m)}$ the projection onto the ($[d]$-adic closure of) the subspace spanned by all symmetric polynomials in $\mathbb{R}[d]^*$, and for $\mathfrak{m}(Z)$ the moment series defined in (2.17)) and hence are identical to the signature cumulant of the linear process $Y := (Z\cdot t)_{t\in[0,1]}$. Just as for standardized classical cumulants, the normalisation (4.3) contributes the additional benefit of scale invariance which facilitates our below usage of signature cumulants as a contrast function.

[3] The log in (4.2) denotes the logarithm (2.30) on the space of formal power series.

4.1.1 On the Significance of Signature Cumulants

Let us further illuminate the concept of the signature cumulant from Definition 4.1.2 as essentially that of a natural ('moment-like') multi-indexed log-compressed coordinate vector for the law of a stochastic process.

As further detailed in Section 2.2.3, the idea behind the classical concept of 'moments' of a random variable – both for random vectors in $\mathbb{R}^d$ and stochastic processes alike – is that they provide a set of deterministic *coordinates* for the law of that random variable.

For the simplest case of a scalar random variable, i.e. a random vector in $\mathbb{R}^1$, the most established such coordinates are the sequence of its central moments, that is an ordered list of expectations over increasingly nonlinear functionals [specifically: (centered) monomials] of that random variable. Analogous coordinates for a random vector in $\mathbb{R}^d$, namely its multivariate (central) moments, are obtained if one considers the expectation over the multivariate (centered) monomials of that random vector.[4]

The expected signature is yet a further generalisation of this classical concept of coordinatisation, this time from random elements in $\mathbb{R}^d$ to random elements in $\mathcal{C}_d$, i.e. continuous-time stochastic processes: As before one considers a multiindexed family of numbers (4.1) given as expectations over increasingly nonlinear functionals of the process, but this time the functionals in question are no longer (multivariate) monomials on $\mathbb{R}^d$ but rather iterated integrals on (elements of) $\mathcal{C}_d$. Owing to the more complex (time-ordered) global structure that is innate to the realisations of a stochastic process (its sample paths), these path-space functionals are better suited [than classical multivariate monomials] to capture the (geometrical) complexity of the open sets (wrt. the uniform topology) of $\mathcal{C}_d$ whose numerical valuation constitutes the law of the process. The family of expectations (4.1) over these functionals of the process, also referred to as *noncommutative moments* due to their non-invariance under index-permutations, does hence form a more informative global statistic of the stochastic process than can be provided by the classical moments of the process at any fixed time-point of its evolution. In fact, these expected signature coefficients do provide a high-resolution description of the stochastic process which is so fine-grained that in total they are able to characterise the law [34] of the process. Another of their main advantages, leveraged below, is that these coefficients are interrelated by way of a rich algebraic and combinatorial global structure, which (akin to classical moments for the case of random

[4] The only difficulty with this multidimensional generalisation is that the resulting 'coordinate vector' is no longer a linearly ordered list (isomorphic to (the coefficients of) a formal power series in a single variable) but rather a multiindexed family of numbers (isomorphic to a formal power series in several variables); cf. Section 2.3.1.

vectors) makes core aspects of stochastic process statistics amenable to a lucid algebraic description.

These relations reveal, however, that considered in isolation the family of coefficients (4.1), that is the expected signature of the stochastic process, appears somewhat 'bloated' in that it exhibits a certain level of internal algebraic redundancy. In fact, we have seen in Section 2.3.2 that the expected signature is 'close to an exponential', see (2.26) and (2.29), which allows for the information it contains to be efficiently compressed by a logarithmic 'change of coordinates'. What results is the multi-indexed coordinate vector (4.2), an algebraically accessible reservoir of conveniently organised statistical information that characterises the law of a stochastic process.

4.2 Signature Contrasts for Nonlinear ICA

Similar to how classical cumulants are traditional in linear ICA, cf. page 36, the usage of signature cumulants in our present ICA-context is due to the following observation: Recall that a random vector Y in $\mathbb{R}^d$ has independent components if and only if all of its cross-cumulants vanish, that is iff, in the notation of (4.4) and for $*$ the concatenation of indices,

$$\kappa_{\boldsymbol{q}}^{Y} \;=\; 0 \quad \text{for all} \quad \boldsymbol{q} \in \bigsqcup_{k=2}^{d} \{\boldsymbol{i} * \boldsymbol{j} \;|\; \boldsymbol{i} \in [k-1]^\star \setminus \{\emptyset\},\; \boldsymbol{j} \in \{k\}^\star \setminus \{\emptyset\}\}. \tag{4.5}$$

Now in the same way that the expected signature generalises the classical concept of moments, cf. Sections 2.2.3 and 4.1.1, it was shown in [19] that signature cumulants generalise this classical relation (4.5) to an algebraic characterisation of statistical independence between [the components of] stochastic processes. This is particularly useful in our context as it yields a natural and explicitly computable contrast function for path-valued random variables (Proposition 4.2.2) as desired for nonlinear ICA.

Algebraically, cf. Remark 2.3.3, the (4.2)-based extension of the characterisation (4.5) to stochastic processes requires us to replace the simple operation $*$ of index concatenation by a slightly more involved combinatorial operation on $[d]^\star$. This operation is defined next.

Notation 4.2.1. For convenience, we denote by $[d]^\star_+$ the family of all finite sums of indices in $[d]^\star$, and for any such sum $\boldsymbol{i} \equiv \boldsymbol{i}_1 + \ldots + \boldsymbol{i}_\ell \in [d]^\star_+$ define $\kappa_{\boldsymbol{i}} := \kappa_{\boldsymbol{i}_1} + \ldots + \kappa_{\boldsymbol{i}_\ell}$.

Recalling (2.28) for convenience, the *shuffle product* of two multi-indices $\boldsymbol{i} = (i_1, \ldots, i_m)$ and $\boldsymbol{j} = (i_{m+1}, \ldots, i_{m+n})$ in $[d]^\star$ is defined as the element of $[d]^\star_+$ which is given by

$$\boldsymbol{i} \sqcup\!\sqcup \boldsymbol{j} \;:=\; \sum_{\tau} (i_{\tau(1)}, \ldots, i_{\tau(m+n)}) \;\in\; [d]^\star_+ \tag{4.6}$$

where the sum is taken over the family of permutations

$$\{\tau \in S_{m+n} \mid \tau(1) < \cdots < \tau(m) \ \text{ and } \ \tau(m+1) < \cdots < \tau(m+n)\}.$$

This enables us to formulate the following central observation.

Proposition 4.2.2. *For any stochastic process $Y = (Y^1, \cdots, Y^d)$ in $\mathbb{R}^d$ whose expected signature exists, the components $Y^1, \ldots, Y^d$ are mutually independent if and only if*

$$\bar{\kappa}_{\mathrm{IC}}(Y) \ := \ \sum_{k=2}^{d} \sum_{\boldsymbol{q} \in \mathcal{W}_k} \bar{\kappa}_{\boldsymbol{q}}(Y)^2 \ = \ 0 \tag{4.7}$$

where $\mathcal{W}_k \ := \ \{\boldsymbol{i} \sqcup \boldsymbol{j} \mid \boldsymbol{i} \in [k-1]^\star \setminus \{\emptyset\}, \ \boldsymbol{j} \in \{k\}^{\times m}, \ m \geq 1\} \ \subset \ [d]^\star_+$.

Proof. Observe that the coordinate processes $Y^1, \ldots, Y^d$ are mutually independent iff:

$$\text{for each } \ 2 \leq k \leq d, \qquad \text{the process } \ Y^k \ \text{ is independent of } \ (Y^1, \cdots, Y^{k-1}). \tag{4.8}$$

Indeed: Denoting by $\mathcal{C}_k \cong C_1^{\times k}$ the co-domain of the process $Y_{[k]} \equiv (Y^1, \cdots, Y^k)$ and writing $\mathcal{B}_k \equiv \mathcal{B}(\mathcal{C}_k) \cong \mathcal{B}_1^{\otimes k}$ for the Borel σ-algebra on $\mathcal{C}_k$ (Lemma 2.1.2), the characterisation (4.8) is immediate by the definition of mutual independence and the fact that

$$\mathbb{P}_{Y_{[k]}}(A_1 \times \cdots \times A_k) = \mathbb{P}_{(Y_{[k-1]}, Y^k)}((A_1 \times \cdots \times A_{k-1}) \times A_k) \quad (k \geq 2)$$

for any $A_1, \ldots, A_k \in \mathcal{B}_1$. The asserted characterisation (4.7) now follows from (4.8) and [19, Theorem 1.2 (iii)]. □

We may now combine Proposition 4.2.2 with Theorems 3.3.3 and 3.3.7 to obtain the following instance of (3.22) for the inversion '$X \mapsto S$' that is desired in (1.2) (cf. Corollary 3.1.3).

(Recall Remark 4.0.1 for well-definedness of the signature statistics featured in (4.9).)

Theorem 4.2.3. *Let the process S in (3.15) be α-, β- or γ-contrastive with sample-paths of bounded variation. Then it holds with probability one that*

$$\left[\operatorname*{arg\,min}_{h \in \Theta} \bar{\kappa}_{\mathrm{IC}}\big(h(X)\big)\right] \cdot X \ \subseteq \ \mathrm{DP}_d \cdot S \tag{4.9}$$

for any family of transformations $\Theta \subseteq C^{2,2}(D_X)$ with $\Theta\big|_{D_X} \cap \left(\mathrm{DP}_d(D_S) \cdot f^{-1}\right)\big|_{D_X} \neq \emptyset$.

This theorem states that the initial problem (1.2) of nonlinear blind source separation can be reformulated as a problem of optimisation-based function approximation. More specifically, statement (4.9) says that the desired demixing transformations of the data can be found as minimizers of the energy-like functional (4.7). We conclude with a few practical remarks.

Remark 4.2.4. (i) For $\Theta \subseteq \mathrm{GL}_d$ and under the temporally degenerate hypothesis of Theorem 3.1.1, the procedure (4.9) reduces to Comon's method (3.2) for $\phi = \phi_c$ since

$$\bar{\kappa}_{\mathrm{IC}}\big((Y \cdot t)_{t\in[0,1]}\big) \;=\; \phi_c(Y) \quad \text{if } Y \text{ is a random vector in } \mathbb{R}^d \qquad \text{(cf. Remark 4.1.3).}$$

(ii) Regarding implementations of (4.9), one may choose to realise the above domain Θ by way of an Artificial Neural Network, see e.g. Section 4.8.3. This choice is mathematically justified by the fact that neural networks can be designed as universal approximators to $C^{2,2}(D_X)$ [149] with a favourable convergence topology [124] (cf. also Remarks 4.4.2, 4.8.3).

(iii) In practice, only discrete-time observations $(X_t)_{t\in\mathcal{I}}$ of X for a finite $\mathcal{I} \subset \mathbb{I}$ are available. Our framework covers these discretised observations as well, as we can naturally identify the data $(X_t)_{t\in\mathcal{I}}$ with a continuous bounded variation process in $\mathbb{R}^d$ via piecewise-linear interpolation of the points $\{X_t \mid t \in \mathcal{I}\}$. The identifiability procedure of Theorem 4.2.3 is robust under this discretisation, see Theorem 4.4.13 and Section 3.7.1 (iii) in particular. A quick inspection of Theorems 3.3.3 & 3.3.7 further reveals that the identifiability approach of the preceding sections can be immediately extended to discrete time-series that are not necessarily generated from continuous-time processes, see Section 3.7 for details.

(iv) The contrast function $\bar{\kappa}_{\mathrm{IC}}$ can be efficiently approximated by restricting the summation in (4.7) to multindices $(i_1, \ldots, i_m)$ up to a maximal order $m \leq m_0$ and estimating these remaining summands using the unbiased minimum-variance estimators for signature cumulants introduced in [19, Section 4]. For the latter, a more naive but straightforward approach that is sufficient for our experiments is to just use the Monte-Carlo estimator, see Sections 4.4.3 & 4.4.5 for details.

(v) For a fixed and finite data set, lower-order summands in the above (capped) approximation of the contrast $\bar{\kappa}_{\mathrm{IC}}$ are typically estimated more accurately than higher-order summands. In practical applications this may be accounted for by applying weights to the estimated summands of the contrast, leading one to estimate the alternative objective

$$\textstyle\sum_{m=2}^{m_0} \sum_{\boldsymbol{q}\in\mathfrak{C}_m} w_{\boldsymbol{q}} \cdot \bar{\kappa}_{\boldsymbol{q}}(Y)^2 \quad \text{for} \quad 0 < w_{\boldsymbol{q}} \equiv w_{\boldsymbol{q}}(Y) \;\; \text{decreasing in the order of } \boldsymbol{q},$$

where we used the notation of (4.22) for convenience. Appropriate choices of weights $(w_{\boldsymbol{q}})$ will generally depend on m_0 and the respective domain of $\bar{\kappa}_{\mathrm{IC}}$, but may otherwise be arbitrary provided that the $\arg\min$ of the thus-weighted objective coincides with the $\arg\min$ of the default case $w_{\boldsymbol{q}} \equiv 1$ (as, e.g., is guaranteed by the final assumption of Theorem 4.2.3).

4.3 'Classical' Contrast Functions on Time-Series

On a theoretical note, let us remark how alternative contrast functions for multivariate time-series may be obtained by an extension of classical independence criteria for random vectors:

In principle, every classical contrast function $\phi : \mathcal{M}_1(\mathbb{R}^d) \to \mathbb{R}_+$ (cf. Cor. 3.1.3) naturally induces an independence criterion $\tilde{\phi} : \mathcal{M}_1(\mathbb{R}^{d\times n}) \to \mathbb{R}_+$ for $\mathbb{R}^d$-valued time-series

$$\mathcal{Y} \;\equiv\; (Y_t^1, \cdots, Y_t^d)_{t\in[n]} \;\cong\; \Big((Y_t^1)_{t\in[n]}, \cdots, (Y_t^d)_{t\in[n]}\Big) \tag{4.10}$$

over a finite time-horizon $n \in \mathbb{N}$, namely by composing ϕ with (the d-fold Cartesian product of) a measurable injection from $\mathbb{R}^n$ to $\mathbb{R}$.

Proposition 4.3.1. *If $\phi_{\mathrm{cl}} : \mathcal{M}_1(\mathbb{R}^d) \to \mathbb{R}_+$ is a classical contrast function and $\psi : \mathbb{R}^n \to \mathbb{R}$ is a measurable injection mapping measurable sets to measurable sets, then the composition $\tilde{\phi} := \phi_{\mathrm{cl}} \circ (\psi \times \cdots \times \psi) : \mathcal{M}_1(\mathbb{R}^{d\times n}) \to \mathbb{R}_+$, given by (for arguments $\mathcal{Y}$ as in (4.10))*

$$\tilde{\phi}(\mathcal{Y}) \;=\; \phi_{\mathrm{cl}}(\psi((Y_t^1)_{t\in[n]}), \cdots, \psi((Y_t^d)_{t\in[n]})),$$

is a contrast function on the product space $\mathbb{R}^{d\times n} \cong (\mathbb{R}^n)^{\times d}$, i.e. such that

$$\tilde{\phi}(\mathcal{Y}) = 0 \quad \textit{if and only if} \quad (Y_t^i)_{t\in[n]},\ i \in [d],\ \textit{are mutually independent.} \tag{4.11}$$

Proof. Let $\mathcal{Y}$ be as in (4.10) with spatial components $\mathcal{Y}_i \equiv (Y_t^i)_{t\in[n]}$. The 'if'-direction in (4.11) is clear as independence is preserved under measurable transformations. Suppose hence that the random variables $Z_1 := \psi(\mathcal{Y}_1), \ldots, Z_d := \psi(\mathcal{Y}_d)$ are mutually independent. As ψ is injective we have that $A = \psi^{-1}(\psi(A))$ for each $A \subseteq \mathbb{R}^n$, where by assumption $\psi(A)$ is measurable if A is measurable. Thus for any measurable $A_1, \ldots, A_d \subseteq \mathbb{R}^n$ we find that

$$\begin{aligned}\mathbb{P}\left(\bigcap_{i\in[d]} \mathcal{Y}_i^{-1}(A_i)\right) &= \mathbb{P}\left(\bigcap_{i\in[d]} (\mathcal{Y}_i^{-1} \circ \psi^{-1})(\psi(A_i))\right)\\ &= \prod_{i\in[d]} \mathbb{P}\Big(Z_i^{-1}(\psi(A_i))\Big) = \prod_{i\in[d]} \mathbb{P}\Big(\mathcal{Y}_i^{-1}(A_i)\Big)\end{aligned}$$

which shows that the vectors $\mathcal{Y}_1, \ldots, \mathcal{Y}_d$ are independent, as desired. □

We call an injection ψ as in Proposition 4.3.1 a *link function*, and introduce Cantor's interlacing map as a particular example of such a function.

Definition 4.3.2 (Cantor's n-Link)**.** Consider the map $\mathfrak{c} : (0,1)^2 \to (0,1)$ which takes the decimal digits of its arguments and interleaves them, i.e.

$$(x, y) \equiv (0.a_1a_2a_3\ldots, 0.b_1b_2b_3\ldots) \;\mapsto\; \mathfrak{c}(x, y) := (0, a_1b_1a_2b_2a_3b_3\ldots) \tag{4.12}$$

for $x = \sum_{i=1}^{\infty} a_i \cdot 10^{-i}$ the (terminating[5]) decimal representation of x (likewise for y), and for $n \in \mathbb{N}_{\geq 3}$ define $\mathfrak{c}_n : (0,1)^n \to (0,1)$ recursively via $\mathfrak{c}_2 := \mathfrak{c}$ and

$$\mathfrak{c}_n(x_1, \cdots, x_n) := \mathfrak{c}(\mathfrak{c}_{n-1}(x_1, \ldots, x_{n-1}), x_n). \tag{4.13}$$

For $\tau_1 : \mathbb{R}^n \to \mathbb{R}^n_{>0},\ (x_i) \mapsto (e^{x_i})$, and $\tau_2 : \mathbb{R}^n_{>0} \to (0,1)^n,\ (y_i) \mapsto (\frac{y_i}{y_i+1})$, the map

$$\psi_{\mathcal{C}} := \mathfrak{c}_n \circ (\tau_2 \circ \tau_1) \,:\, \mathbb{R}^n \to (0,1) \tag{4.14}$$

will be referred to as *Cantor's n-link*.

Lemma 4.3.3. *Cantor's n-link is a link function in the sense of Proposition 4.3.1.*

Proof. We need to show that for any $n \in \mathbb{N}_{\geq 2}$ fixed, the function (4.14) is a Borel-measurable injection which maps Borel-sets to Borel-sets; note that since both τ_1 and τ_2 are diffeomorphisms, it remains to verify these properties for (4.13).

We do this via induction on n. Let $n = 2$ first. The fact that $\mathfrak{c}$ is injective is clear from (4.12), so let us next prove that $\mathfrak{c}$ is Borel-measurable. For this it suffices to show that

$$\mathfrak{c}^{-1}(\mathfrak{B}) \subseteq \mathcal{B}((0,1)^2) \quad \text{ for } \mathfrak{B} \text{ any (topological) base of } (0,1). \tag{4.15}$$

To this end, we consider the particular base of $(0,1)$ which is given by

$$\mathfrak{B} := \left\{ I_{\alpha,2k} := (\alpha,\, \alpha + 10^{-2k}) \;\middle|\; \alpha = \sum_{i=1}^{\infty} a_i \cdot 10^{-i} \text{ for } (a_i) \in \mathfrak{d}_{2k} \text{ and } k \in \mathbb{N} \right\} \tag{4.16}$$

where for $k \in \mathbb{N}$ we denote by $\mathfrak{d}_k$ the space of all decimal representations $(a_i)_{i\in\mathbb{N}} \subset ([9]_0)^{\mathbb{N}}$ with $\sup_{i>k} a_i = 0$ and $a_k < 9$. (The fact that (4.16) defines a base of the Euclidean topology on $(0,1)$ is then clear by definition.) To for this base prove (4.15), take any $I_{\alpha,2k} \equiv (\alpha, \alpha + 10^{-2k}) = 0.a_1 \ldots a_{2k} + (0, 10^{-2k})$ for $\alpha = \sum_{i=1}^{\infty} a_i \cdot 10^{-i}$ with $(a_i) \in \mathfrak{d}_{2k}$, and note that then

$$\begin{aligned}
\mathfrak{c}^{-1}(I_{\alpha,2k}) &= \mathfrak{c}^{-1}\big[(0.a_1a_2a_3 \ldots a_{2k},\ 0.a_1a_2a_3 \ldots (a_{2k}+1))\big] \\
&= \big(0.a_1a_3 \ldots a_{2k-1},\ 0.a_1a_3 \ldots (a_{2k-1}+1)\big) \times \big(0.a_2a_4 \ldots a_{2k},\ 0.a_2a_4 \ldots (a_{2k}+1)\big) \\
&= I_{\alpha',k} \times I_{\alpha'',k} \qquad \text{for} \quad \alpha' = \sum_{i=1}^{\infty} a_{2i-1} \cdot 10^{-i} \text{ and } \alpha'' = \sum_{i=1}^{\infty} a_{2i} \cdot 10^{-i}.
\end{aligned}$$

To see that $\mathfrak{c}$ maps Borel sets to Borel sets, note that by the injectivity of $\mathfrak{c}$ (and the resulting fact that $\mathfrak{c}(A \cap B) = \mathfrak{c}(A) \cap \mathfrak{c}(B)$ for any $A, B \subseteq (0,1)^2$, implying that $\mathfrak{c}$ is compatible with any countable set-theoretic operation on $(0,1)^2$) this holds if

$$\mathfrak{c}(I \times (0,1)),\ \mathfrak{c}((0,1) \times I) \in \mathcal{B}(0,1) \quad \text{ for each } \quad I \in \mathfrak{B},$$

[5] Whence, e.g., $\frac{1}{2} = 0.500000\ldots$ instead of $\frac{1}{2} = 0.499999\ldots$.

and by symmetry it suffices to check that $\mathfrak{c}(I \times (0,1)) \in \mathcal{B}(0,1)$ for any given $I = I_{\alpha,k} \in \mathfrak{B}$, say $\alpha = \sum_{i=1}^{k} a_i \cdot 10^{-i}$. And indeed, the definition (4.12) yields

$$\begin{aligned}\mathfrak{c}(I_{\alpha,k} \times (0,1)) &= \left\{ \sum_{i=1}^{\infty} b_i \cdot 10^{-i} \;\middle|\; (b_{2\nu-1})_{\nu\in[k]} = (a_\nu)_{\nu\in[k]},\ b_j \in [9]_0 \text{ else} \right\} \\ &= \bigcup_{\beta\equiv(\beta_\nu)\in([9]_0)^{\times k}} \left(\sum_{i=1}^{k} (a_i \cdot 10^{-(2i-1)} + \beta_i \cdot 10^{-2i}) + (0,\ 10^{-2k}) \right) \\ &= \bigcup_{\beta\in([9]_0)^{\times k}} I_{\alpha_\beta,k} \ \in\ \mathcal{B}(0,1) \qquad \text{as desired,}\end{aligned}$$

where $\alpha_\beta \equiv \alpha_{(\beta_\nu)} := \sum_{i=1}^{k}(a_i \cdot 10^{-(2i-1)} + \beta_i \cdot 10^{-2i})$.

With (4.12) hence being a Borel-measurable injection which maps Borel sets to Borel sets, it is clear that the composition (4.13) inherits this property for each $n \in \mathbb{N}_{\geq 2}$, which concludes the proof. □

4.4 Statistical Consistency

In practice, the mixture X is usually not observed as a whole stochastic process with fully known distribution but rather as a time-discretized sample consisting of finitely many data points in $\mathbb{R}^d$, often taken over non-equally spaced time grids. Further, often only a single such (time-discretized) sample path of the process is available rather than many independent sample realisations,[6] for example in the classical cocktail party problem. In this section we demonstrate that our proposed nonlinear ICA method is stable under such discretisations. More precisely, we formalise the usual observation schemes (Section 4.4.1) and prove that if these sampling discretisations of X get 'finer' in some natural sense, then our ICA method produces a signal that gets uniformly closer to the unobserved source S underlying X (Section 4.4.6). The main steps towards this are outlined in Section 4.4.2.

We assume throughout this section that the mixture X is a continuous-time signal. As before, the (simpler) case where the underlying mixture X is discrete in time is covered analogously up to some minor modifications, as we explain in detail in Section 4.5.0.1.

4.4.1 Sampling

In practical applications, observations of the mixture X are typically not available as continuous paths, i.e. elements of $\mathcal{C}_d$, but rather as discrete, sequentially ordered collections of data points in $\mathbb{R}^d$. Formally, such data can be modelled as a discrete time series

$$\mathfrak{x} := (X_t(\omega))_{t\in\mathcal{I}} \quad \text{with} \quad \mathcal{I} := \{0 = t_0 < t_1 < \ldots < t_{n-1} = 1\} \tag{4.17}$$

[6] Note, however, that the latter can be regarded as a special case of the former by concatenating the available (independent) sample observations into a single long observation.

for $\omega \in \Omega$ any fixed elementary event in the probability space $(\Omega, \mathscr{F}, \mathbb{P})$ underlying X. (In empirical language, the above dissection $\mathcal{I}$ of $[0,1]$ can then be regarded as a 'protocol' describing a sequence of measurements of X carried out per unit time with frequency n.)

The time series data (4.17) is then typically collected over not just one but several ($\nu \in \mathbb{N}$) time intervals (per observation), and different observations of X ($k \in \mathbb{N}$) may vary in the frequency at which their fixed-time measurements are made. This gives rise to the data scheme

$$\mathfrak{x}^{(k)} \equiv (\mathfrak{x}_1^{(k)}, \mathfrak{x}_2^{(k)}, \cdots) := (X_t(\omega))_{t\in\mathcal{J}_k}, \quad \text{for} \quad \mathcal{J}_k := \mathcal{I}_1^{(k)} \sqcup \mathcal{I}_2^{(k)} \sqcup \ldots = \bigsqcup\nolimits_{\nu\in\mathbb{N}} \mathcal{I}_\nu^{(k)} \quad (4.18)$$

a dissection[7] of $[0,\infty)$ such that $\mathcal{I}_1^{(k)} < \mathcal{I}_2^{(k)} < \ldots$ and $\hat{n}_k := \sup_{\nu\in\mathbb{N}} |\mathcal{I}_\nu^{(k)}| < \infty$,[8] with $\mathcal{I}_1^{(k)}$ a dissection of $[0,1]$ and $\mathfrak{x}_\nu^{(k)} = (X_t(\omega))_{t\in\mathcal{I}_\nu^{(k)}}$ for each $\nu \in \mathbb{N}$. Any such family $(\mathcal{J}_k)_{k\in\mathbb{N}}$ will be called a *protocol*, where the index $k \in \mathbb{N}$ enumerates the different observations of X.

Adopting the ergodicity perspective common in time-series analysis and signal processing, the relevant statistical information of X (in our case: the signature-cumulant coordinates (4.2)) may be averaged from the discretely-sampled observation $\mathfrak{x}^{(k)}$ of a single, sufficiently long realisation of X. For this, we may generalise established observation schemes[9] by assuming that the data $\mathfrak{x}^{(k)}$ in (4.18) be obtained according to nothing else but the assumptions

$$X = (X_t)_{t\geq 0} \quad \text{and} \quad \lim_{k\to\infty} \|\mathcal{I}_1^{(k)}\| = 0\,, \tag{4.19}$$

that is, the requirements that the continuous observable X be 'infinitely long' (i.e., defined over $[0,\infty)$) and the frequency $|\mathcal{I}_1^{(k)}|$ of observations per initial interval be going to infinity.

A protocol $(\mathcal{J}_k)_{k\in\mathbb{N}}$ as in (4.18) & (4.19) will be called *exhaustive* with *base lengths* $n_k := |\mathcal{I}_1^{(k)}|$ and *maximal (observation) length* $\hat{n}_k$; we called it *balanced* if $|\mathcal{I}_\nu^{(k)}| = n_k$ for all $\nu \in \mathbb{N}$.

Remark 4.4.1. Note that the sampling scheme (4.18) allows for the units $\mathcal{I}_\nu^{(k)} =: \{t_0^{(k|\nu)} < \cdots < t_{n_{k,\nu}-1}^{(k|\nu)}\}$ to span observation intervals $[t_0^{(k|\nu)}, t_{n_{k,\nu}-1}^{(k|\nu)}]$ of different lengths and dissect them at different, non-constant frequencies. Notice further that while the mesh $\mathcal{J}_k$ in (4.18) is infinite, this does not restrict us to considering observations $\mathfrak{x}^{(k)}$ that contain an infinitude

[7] We call $\mathcal{J}_T := \{t_0 < t_1 < \ldots\}$ a *dissection of* $[0,\infty)$ if $t_0 = 0$ and $t_j \nearrow \infty$ as $j \to \infty$. Also, we then assume X to be defined over the full positive time-axis $[0,\infty)$, see (4.19).

[8] Notice that $|\mathcal{I}|$ denotes the cardinality of a set $\mathcal{I} \subset \mathbb{R}$ (i.e. the number of its elements) while the *mesh-size* of $\mathcal{I}$ is denoted $\|\mathcal{I}\|$, cf. (4.29). (Consequently $|\mathcal{I}| \geq (\max\mathcal{I} - \min\mathcal{I})/\|\mathcal{I}\| + 1$.)

[9] Observations schemes are typically required to be equispaced, see e.g. [52, Sect. 3] and [159] for an overview.

of information $(\mathfrak{x}_\nu^{(k)})$. In fact, everything presented in this section works as stated if we relax the above definition of a protocol by replacing the observable X in (4.18) with the cutoff

$$X \cdot \mathbb{1}_{[0,T_k]} \quad \text{for some finite observation horizon } \ T_k \geq 0 \text{ with } \lim_{k\to\infty} T_k = \infty,$$

thus extending the permissible data $(\mathfrak{x}^{(k)})$ in (4.18) to finite (eventually zero) sequences.

4.4.2 Section Overview

Let $(\mathfrak{x}^{(k)})$ be data associated to an exhaustive observation $\big(X, (\mathcal{J}_k)\big)$ via the sampling scheme (4.18). This section aims to establish conditions under which the optimisation procedure of Theorem 4.2.3, when applied to $(\mathfrak{x}^{(k)})$ in lieu of X, yields a sequence of approximations $(\hat{\theta}_k)$ of the true demixing inverse f^{-1} which is statistically consistent in the sense that

$$\lim_{k\to\infty} \operatorname{dist}\big(\hat{\theta}_k(X),\, \mathrm{DP}_d \cdot S\big) \ = \ 0 \tag{4.20}$$

almost surely or in probability, where the distance is taken with respect to the uniform norm on $\mathcal{C}_d$.[10] Since the original optimisation (4.9) is composed of three ('limiting') operations that each involve an 'infinite amount of information', namely the *infinite series* $\bar{\kappa}_{\mathrm{IC}}$ from (4.7) whose summands (4.3) are each defined by taking *expectations* of nonlinear functionals (4.1) of the *continuous-time* stochastic processes $Y = \theta(X)$, one may expect the consistency (4.20) to result as a combination of the following three sublimits:

- *Capping Limit* (Section 4.4.3). In practice, only finitely many summands of the infinite statistics $\bar{\kappa}_{\mathrm{IC}}$ from (4.7) can be computed from the data. This is to say that the series

$$\bar{\kappa}_{\mathrm{IC}}(Y) \ = \ \sum_{m=2}^{\infty} \sum_{\boldsymbol{q}\in\mathfrak{C}_m} \bar{\kappa}_{\boldsymbol{q}}(Y)^2, \tag{4.21}$$

  with $\mathfrak{C}_m \subset [d]_+^\star$ denoting the set of all cross-shuffles $\mathfrak{C} := \bigsqcup_{k=2}^d \mathcal{W}_k$ of word-length $m \in \mathbb{N}$ (see Prop. 4.2.2), needs to be capped at some index $m = m_0$. Denoting this capped series by

$$\bar{\kappa}_{\mathrm{IC}}^{[m_0]}(Y) \ := \ \sum_{m=2}^{m_0} \sum_{\boldsymbol{q}\in\mathfrak{C}_m} \bar{\kappa}_{\boldsymbol{q}}(Y)^2 \tag{4.22}$$

  we show that in the *capping limit* $m_0 \to \infty$ the minimizers of $\theta \mapsto \bar{\kappa}_{\mathrm{IC}}^{[m_0]}\big(\theta(X)\big)$ approach those of (4.21) with respect to a naturally chosen topology on Θ; this provides the first ingredient for the consistency limit (4.20).

---

[10] The trivial case $\operatorname{dist}\big(\hat{\theta}_k(X), \mathrm{DP}_d \cdot S\big) \leq \|\hat{\theta}_k(X) - \alpha_k(S)\| \leq \|\hat{\theta}_k(X)\| + \|\alpha_k(S)\| \overset{!}{\to} 0 \ \ (k \to \infty)$ is automatically excluded if (X is non-trivial and) $(\hat{\theta}_k) \subseteq \Theta$ is bounded away from zero. This is guaranteed by the below assumption, in Theorem 4.4.13, of Θ being a compact subset of $C^{1,1}(D_X)$.

- *Interpolation Limit* (Section 4.4.4). The mixture X is usually observed along a discrete set of time-points $\mathcal{I}$ rather than continuously over time (Sect. 4.4.1). Via their piecewise-linear interpolation $\hat{X}_{\mathcal{I}}$, these discrete observations $(X_t)_{t\in\mathcal{I}}$ can be reinterpreted as $\mathcal{C}_d$-valued data, allowing us to approximate the summands in (4.22) via

$$\bar{\kappa}_{\boldsymbol{q}}(\theta(X)) \approx \bar{\kappa}_{\boldsymbol{q}}(\hat{X}^{\theta}_{\mathcal{I}}), \quad \text{for } \hat{X}^{\theta}_{\mathcal{I}} \text{ the linear interpolant of } \big(\theta(X_t)\big)_{t\in\mathcal{I}}. \tag{4.23}$$

  Showing that (4.23) defines a Θ-uniform approximation as $\|\mathcal{I}\| \to 0$, we obtain that the minimizers of $\theta \mapsto \bar{\kappa}^{[m_0]}_{\mathrm{IC}}(\hat{X}^{\theta}_{\mathcal{I}})$ converge to those of (4.22); our second ingredient for (4.20).

- *Ergodicity Limit* (Section 4.4.5). Finally, as the available data (4.18) is but a single realisation of the discrete time-series $(X_t)_{t\in\mathcal{J}_k}$, we propose to compute the approximations (4.23) by estimating their constituent signature moments (4.1) via

$$\sigma_{\boldsymbol{i}}(\hat{X}^{\theta}_{\mathcal{I}^{(k)}_1}) \approx \frac{1}{T}\sum_{\nu=1}^{T}\mathfrak{sig}_{\boldsymbol{i}}(\hat{\mathfrak{x}}^{\theta|k}_{\nu}), \quad \text{for } \hat{\mathfrak{x}}^{\theta|k}_{\nu} \text{ the linear interpolant of } \theta(\mathfrak{x}^{(k)}_{\nu}) \tag{4.24}$$

  and where $\mathfrak{sig}_{\boldsymbol{i}}(Y)$ denotes the iterated integrals inside the expectation (4.1) (cf. Sect. 2.2.2, eqs. (2.18) and (2.15)). Showing that for many popular time-series models and stochastic signals the above estimation scheme (4.24) for $\bar{\kappa}_{\boldsymbol{q}}(\hat{X}^{\theta}_{\mathcal{I}})$ is Θ-uniformly consistent as $T \to \infty$, we obtain our third and final ingredient for (4.20).

In Section 4.4.6, these three sublimits are then combined to prove that our nonlinear ICA method (Theorem 4.2.3) gives rise to statistically consistent estimators of the sources underlying the data, see Theorem 4.4.13 which also includes the consistency limit (4.20) as a special case. The resulting approach is condensed into a readily implementable source estimator in Section 4.6.

The majority of the proofs for this section are deferred to Section 4.7 as they are mostly technical and independent of the argumentation developed in the main body of this work.

4.4.3 The Capping Limit

Let the subset $\mathfrak{C}_m$ of $[d]^{\star}_{+}$ denote the set of all cross-shuffles $\mathfrak{C} := \bigsqcup_{k=2}^{d}\mathcal{W}_k$ of fixed word-length m,[11] for $\mathcal{W}_2, \ldots, \mathcal{W}_d$ as in Proposition 4.2.2 and $m \in \mathbb{N}$.

[11] The *word-length* of an element $\boldsymbol{i} \in [d]^{\star}_{+}$ is defined as the maximal order of the (finitely many) indices in $[d]^{\star}_{+}$ whose formal sum is $\boldsymbol{i}$ (cf. Notation 4.2.1). Thus $\mathfrak{C}_m = V_m \cap \bigsqcup_{k=2}^{d}\mathcal{W}_k$ in the language of Section 2.3.3.

Let further Θ be a given set of nonlinearities (as specified below), and for $(\kappa_q(Y) \mid q \in [d]^\star)$ as in (4.2) and $\theta \in \Theta$, consider the (in)finite cumulant series

$$Q(\theta) := \sum_{\nu=2}^{\infty} \sum_{\boldsymbol{q} \in \mathfrak{C}_\nu} c_{\boldsymbol{q}}^{-1} \cdot \kappa_{\boldsymbol{q}}\big(\theta(X)\big)^2 \quad \text{and} \quad Q_m(\theta) := \sum_{\nu=2}^{m} \sum_{\boldsymbol{q} \in \mathfrak{C}_\nu} c_{\boldsymbol{q}}^{-1} \cdot \kappa_{\boldsymbol{q}}\big(\theta(X)\big)^2 \tag{4.25}$$

for $m \geq 2$, where $c_{\boldsymbol{q}}$ denotes the number of monomials in $\boldsymbol{q}$ (cf. Remark 4.7.1).

We make following technical compatibility assumptions on Θ and X. (For notation, 2.3.3.)

Assumption 1. Let $\Theta \subseteq C(D_X; \mathbb{R}^d)$ be equipped with the topology of compact convergence and suppose that Θ is compact, satisfies $\Theta \cdot X \subseteq \mathcal{BV}$ and is such that it holds with probability one that for every convergent sequence $(\theta_j)_{j\geq 1}$ in Θ there is some $p \in [1, 2)$ with

$$\sup\nolimits_{j\geq 1} \|\theta_j(X)\|_{p\text{-var}} < \infty \tag{4.26}$$

where $\|\cdot\|_{p\text{-var}}$ is the p-variation seminorm (2.4). On side of the signature moments (4.1), suppose that the expected signatures $\mathfrak{S}(\theta(X)) \equiv \mathbb{E}[\mathfrak{sig}(\theta(X))]$, $\theta \in \Theta$, exist and characterize the law of their arguments, that their collection $\{\mathfrak{S}(\theta(X)) \mid \theta \in \Theta\}$ is $|\!|\!|\cdot|\!|\!|_\lambda$-bounded[12] for some $\lambda > 2$, and that for each $m \geq 1$ (with $\mathfrak{sig}_m := \pi_m \circ \mathfrak{sig}$, $\mathfrak{sig}$ as in (2.13)),[13]

$$\mathbb{E}\left[\sup_{\theta\in\Theta} \big\|\mathfrak{sig}_m(\theta(X))\big\|_m\right] < \infty. \tag{4.27}$$

(To avoid potential measurability problems, we may as well replace (4.27) by the (weaker) requirement that $\big[\mathbb{E}\big[\sup_{\theta\in\mathcal{F}} \big\|\mathfrak{sig}_m(\theta(X))\big\|_m\big] < \infty, \ \forall$ countable $\mathcal{F} \subseteq \Theta\big]$ if desired.)

Remark 4.4.2. Notice that the above conditions on Θ and X are quite natural and well-established in the contexts of Artificial Neural Networks (ANNs) and Stochastic Analysis. Indeed: The topological requirement of compact convergence is typically met if Θ is given as the realisation space of an ANN, see e.g. [15, 124], while the assumptions $\Theta \cdot X \subseteq \mathcal{BV}$ resp. (4.26) hold for instance if $\Theta \subseteq C^1(D_X; \mathbb{R}^d)$ resp. if the elements of Θ are continuously differentiable with uniformly bounded Jacobians. The growth assumptions on the signature coordinates (including (4.27)), on the other hand, have been extensively studied, established and applied in the context of rough path analysis and statistics, see e.g. [33] and [19, 34].

Lemma 4.4.3. *Let X and Θ be as described in Assumption 1, and $Q, Q_m : \Theta \to \mathbb{R}$ as given in* (4.25)*. Then the following holds:*

[12] As the source S can be recovered up to (a componental permutation and) a monotone scaling only, we can and will assume wlog (cf. Lemma 2.3.5 (vii)) that the set $\{\mathfrak{S}(\theta \cdot X) - 1 \mid \theta \in \Theta\} \subseteq V_{(0)}$ is $|\!|\!|\cdot|\!|\!|_\lambda$-bounded by 1.

[13] Some of the $\sup_\theta$-related (or $\operatorname{dist}(\,\cdot\,, \mathrm{DP}_d \cdot S)$-related) expressions in the following sections may be non-measurable, in which case any probability statements involving these expressions are to be understood in terms of outer measure (cf. [152, 155]).

(i) *the functions* $Q, Q_m : \Theta \to \mathbb{R}$ *are continuous ;*

(ii) *the capped objectives* Q_m *approximate* Q *uniformly as* m *goes to infinity, in symbols:*

$$\lim_{m\to\infty} \|Q - Q_m\|_\Theta \ = \ 0 \qquad \textit{for} \quad \|q\|_\Theta := \sup_{\theta\in\Theta} |q(\theta)| \, ;$$

(iii) *if* Q *is uniquely minimized at* $\theta_\star \in \Theta$*, i.e. such that* $Q(\theta) > Q(\theta_\star)$ *if* $\theta \neq \theta_\star$*, then any ('minimising') sequence* $(\theta^\star_m)$ *in* Θ *such that* $Q_m(\theta^\star_m) \leq \inf_{\theta\in\Theta} Q_m(\theta) + \eta_m$ *for some* $\eta_m \geq 0$ *with* $\lim_{m\to\infty} \eta_m = 0$ *a.s., converges to* $\theta_\star$ *almost surely as* $m \to \infty$*.*

Proof. See Section 4.7.2. □

The following is but a reformulation of Lemma 4.4.3 in terms of the standardized signature cumulants (4.7). It also anticipates the consistency assertion in Theorem 4.4.13 below.

Proposition 4.4.4 (Capping Limit)**.** *Let* X *and* $\Theta \subseteq C^{2,2}(D_X)$ *fulfil Assumption 1, and suppose that there is a unique* $\theta_\star \in \Theta$ *such that* $\theta_\star(X)$ *is IC. Let* $\bar{\kappa}^{[m]}_{\mathrm{IC}}$ *be as in* (4.22)*. Then for any sequence of minimizers* $(\theta^\star_m)$ *in* Θ *such that*

$$\bar{\kappa}^{[m]}_{\mathrm{IC}}(\theta^\star_m(X)) \ \leq \ \min_{\theta\in\Theta} \bar{\kappa}^{[m]}_{\mathrm{IC}}(\theta(X)) \ + \ \eta_m$$

for some (η_m) *in* $\mathbb{R}_+$ *with* $\lim_{m\to\infty} \eta_m = 0$ *a.s., it holds with probability one that*

$$\lim_{m\to\infty} \mathrm{dist}_{\|\cdot\|_\infty}\big(\theta^\star_m(X),\, \mathrm{DP}_d \cdot S\big) \ = \ 0. \tag{4.28}$$

Proof. Since for any $\boldsymbol{q} \in \mathfrak{C}$ we have that $\kappa_{\boldsymbol{q}}(Y) = 0$ iff $\bar{\kappa}_{\boldsymbol{q}}(Y) = 0$ (recall (4.3) and Notation 4.2.1), it holds that $\arg\min_{\theta\in\Theta} Q_m(\theta) = \arg\min_{\theta\in\Theta} \bar{\kappa}^{[m]}_{\mathrm{IC}}(\theta(X))$ for each $m \geq 2$. The convergence (4.28) is thus immediate by Lemma 4.4.3 (iii) and Prop. 4.2.2/ Thm. 4.2.3. □

4.4.4 The Interpolation Limit

Let $\mathbb{I}$ be a compact interval; say $\mathbb{I} = [0,1]$ wlog as above. A finite subset $\mathcal{I}$ of $\mathbb{I}$ is called a *dissection of* $\mathbb{I}$ if it contains the boundary points of $\mathbb{I}$, and a sequence $(\mathcal{I}_\mu)_{\mu\geq 1}$ of dissections $\mathcal{I}_\mu \equiv \{t^{(\mu)}_0, \ldots, t^{(\mu)}_{n_\mu - 1} \mid t^{(\mu)}_0 < t^{(\mu)}_1 < \ldots < t^{(\mu)}_{n_\mu-1}\}$ of $\mathbb{I}$ is called *refined* if the maximal distance $\|\mathcal{I}_\mu\|$ between two successive points in $\mathcal{I}_\mu$, the so-called *mesh-size* of $\mathcal{I}_\mu$, goes to zero as $\mu \to \infty$; in symbols:

$$\|\mathcal{I}_\mu\| := \max_{j\in[n_\mu - 1]} \big|t^{(\mu)}_j - t^{(\mu)}_{j-1}\big| \ \longrightarrow \ 0 \qquad \text{as} \quad \mu \to \infty. \tag{4.29}$$

Writing $\hat{X}^\theta_{\mathcal{I}}$ for the piecewise linear interpolant[14] of the transformed data $X^\theta_{\mathcal{I}} := \big(\theta(X_t)\big)_{t\in\mathcal{I}}$, $\theta \in \Theta$, the next lemma shows that, as $\|\mathcal{I}\| \to 0$, the statistic (cf. (4.25))

$$\widehat{Q}_m(Y) := \sum_{\nu=2}^{m} \sum_{\boldsymbol{q}\in\mathfrak{C}_\nu} c_{\boldsymbol{q}}^{-1} \cdot \kappa_{\boldsymbol{q}}(Y)^2 \quad \text{with} \quad Y := \hat{X}^\theta_{\mathcal{I}} \qquad (m \geq 2) \tag{4.30}$$

yields a Θ-uniform approximation of the contrast Q_m from (4.25).

Lemma 4.4.5 (Interpolation Limit)**.** *Let Θ and X be as in Assumption 1, $\widehat{Q}_m$ as in* (4.30) *and Q_m as in* (4.25)*. Then for $(\mathcal{I}_n)_{n\in\mathbb{N}}$ any refined sequence of dissections of $\mathbb{I}$ and any $m \in \mathbb{N}_{\geq 2}$,*

$$Q_m(\theta) \;=\; \lim_{n\to\infty} \widehat{Q}_m(\hat{X}^\theta_{\mathcal{I}_n}) \quad \textit{uniformly on } \;\Theta. \tag{4.31}$$

Proof. See Section 4.7.4. □

4.4.5 The Ergodicity Limit

We formalise the estimation scheme (4.24) and show that it holds uniformly on Θ for a large class of time-series models and stochastic processes.

Notation 4.4.6. Let $Z := \mathbb{R}^d$. Given $z \equiv (z_j)_{j\in\mathbb{N}}$ and $J \subset \mathbb{N}$, we write $z_J := (z_j)_{j\in J}$ and $z_{(\ell_1:\,\ell_2]} := (z_{\ell_1+1}, z_{\ell_1+2}, \ldots, z_{\ell_2})$ for $\ell_1, \ell_2 \in \mathbb{N}_0$ with $\ell_1 < \ell_2$, and denote by $\hat{z}_{\mathcal{E}_n} \equiv \hat{\iota}_{\mathcal{E}_n}(z)$ the piecewise-linear interpolation of $z \equiv (z_j)_{j\in[n]} \in Z^{\times n}$ along the equidistant dissection $\mathcal{E}_n := \{(\nu-1)/(n-1) \mid \nu \in [n]\}$ of $[0,1]$ (cf. Section 4.7.3). For $X_* \equiv (X_j)_{j\in\mathbb{N}}$ a discrete time-series in $\mathbb{R}^d$, we denote by $D_{X_*} := \overline{\bigcup_{j\in\mathbb{N}} \operatorname{supp}(X_j)}^{|\cdot|}$ its *spatial support.*

Set further $\mathfrak{sig}_{[m]} := \pi_{[m]} \circ \mathfrak{sig}$ for the signature capped at level $m \geq 2$, cf. Section 2.3.3.

The signature transform (2.14), and thereby its cumulants (4.2), (4.3), are invariant under time-domain reparametrisations of X, see Lemma 2.3.5 (iii). Hence, the statistics (4.22) of an interpolant $Y \equiv \hat{X}_{\mathcal{I}}$ depend only on the time series $X_{\mathcal{I}}$ — i.e. on the random variables $Z_1 := X_{t_0}, \ldots, Z_n := X_{t_{n-1}}$ and their sequential order — and not on the dissection along which $X_{\mathcal{I}} = (Z_j)_{j\in[n]}$ is interpolated. In symbols, see Section 4.7.3 (4.64) for notation,

$$\mathfrak{sig}(\hat{X}_{\mathcal{I}}) \;=\; \mathfrak{sig}(\hat{\iota}_{\mathcal{J}}(Z_1, \ldots, Z_n)) \qquad \text{for any dissection } \mathcal{J} \tag{4.32}$$

with cardinality $|\mathcal{J}| = |\mathcal{I}|$. This justifies to abstract from the topology of the time-indices $t \in \mathcal{I}$ in (4.18), as done in the formulation of Definition 4.4.7 below.

All expectations in the following definition are assumed to exist.

[14] For a formal definition of this operation see Section 4.7.3, where a unified notation for the projection of continuous-time data to discrete time series – and, conversely, the embedding (via interpolation) of the latter type of data into $C(\mathbb{I}; \mathbb{R}^d)$ – is provided.

Definition 4.4.7 (Signature Ergodicity)**.** Let $X_* = (X_j)_{j\in\mathbb{N}}$ be a discrete time-series in $\mathbb{R}^d$ and $n, m \in \mathbb{N}$. We call X_* m^{th}*-order signature ergodic to length* n if, almost surely,

$$\mathbb{E}[\phi(X_{[n]})] \;=\; \lim_{T\to\infty} T^{-1}\sum_{j=1}^{T}\phi(X_{(n(j-1):\,nj]}) \quad \text{for} \quad \phi(z) := \mathfrak{sig}_{[m]}(\hat{z}_{\mathcal{E}_n})\,, \tag{4.33}$$

and X_* will be called *weakly* m^{th}-order signature ergodic to length n if (4.33) holds in probability. We call the process X_* [weakly] *signature ergodic* to length n if X is [weakly] m^{th}-order signature ergodic to length n for every $m \geq 1$.

Given $\Theta \subseteq C(D_{X_*};\mathbb{R}^d)$, we call X_* [weakly/ m^{th}-order] signature ergodic to length n *on* Θ if the respective property holds for each $\theta(X_*) := \big(\theta(X_j)\big)_{j\in\mathbb{N}}$, $\theta \in \Theta$.

We refer to the LHS of (4.33) as the $[m]^{\text{th}}$*-signature moment* of the batch $(X_1, \dots, X_n)$.

Remark 4.4.8. (i) In other words, the time-series $X = (X_j)_{j\in\mathbb{N}}$ is [weakly] m^{th}-order signature ergodic to length n iff the sequence of empirical measures (on $\mathcal{B}(\mathcal{C}_d)$)

$$\hat{\mu}_T \;:=\; \frac{1}{T}\sum_{j=1}^{T}\delta_{\hat{X}_j} \qquad \text{for} \qquad \hat{X}_j := \hat{\iota}_{\mathcal{E}_n}(X_{n(j-1)+1}, \dots, X_{nj})$$

yields a consistent estimator for the expected signature $\mathfrak{S}_{[m]}(\hat{X}_1)$ of $\hat{X}_1$, that is iff

$$\mathfrak{S}_k(\hat{X}_1) \;=\; \lim_{T\to\infty}\int_{\mathcal{C}_d}\mathfrak{sig}_k(x)\,\hat{\mu}_T(\mathrm{d}x) \quad \text{a.s. [in probab.]} \tag{4.34}$$

for each $1 \leq k \leq m$. Notice that due to (4.32), the equidistant dissection $\mathcal{E}_n$ in (4.33) may be replaced by any other $[0,1]$-dissection of the same cardinality.

(ii) A time-series $(X_t)_{t\in\mathcal{J}_k}$ for $\mathcal{J}_k$ as in (4.18), is called m^{th}*-order signature ergodic* if a.s.

$$\mathbb{E}[\tilde{\phi}(\hat{X}_{\mathcal{I}_1^{(k)}})] \;=\; \lim_{T\to\infty} T^{-1}\sum_{\nu=1}^{T}\tilde{\phi}(\hat{X}_{\mathcal{I}_\nu^{(k)}}) \quad \text{for} \quad \tilde{\phi} := \mathfrak{sig}_{[m]}(\,\cdot\,)\,, \tag{4.35}$$

for $\hat{X}_{\mathcal{I}} := \hat{\iota}_{\mathcal{I}}(X_{\mathcal{I}})$ the piecewise-linear interpolation of $X_{\mathcal{I}} \equiv (X_t)_{t\in\mathcal{I}}$ along $\mathcal{I} \subset \mathbb{R}$. The remaining notions of Definition 4.4.7 carry over analogously. Notice that in consequence of Lemma 2.3.5 (iii), the above notions (4.35) of signature ergodicity for protocol-indexed time-series are in fact a special case of Definition 4.4.7: see Lemma 4.7.3. Hence also for time-series of this protocol-indexed kind, the results of this section all apply as stated upon replacing their respective ergodicity assumptions by their (4.35)-type counterparts.

Let as before the space $C(D_{X_*};\mathbb{R}^d)$ be endowed with the compact-open topology. Using the universality of the signature transform (Lemma 2.3.5 (iv)), we find that the [weak] signature ergodicity of $X_* = (X_j)_{j\in\mathbb{N}}$ is passed onto $\theta(X_*) = \big(\theta(X_j)\big)_{j\in\mathbb{N}}$ for any $\theta \in C(D_{X_*};\mathbb{R}^d)$.[15]

Proposition 4.4.9. *Let $\Theta \subseteq C(D_{X_*};\mathbb{R}^d)$ and $X_* = (X_j)_{j\in\mathbb{N}}$ be a discrete time-series in $\mathbb{R}^d$ with compact spatial support and such that for each $\theta \in \Theta$ the expectations*

$$\mathbb{E}[\mathfrak{sig}_m(\hat{X}_1^\theta)] \quad \textit{exist for all } m \geq 1, \quad \textit{with } \hat{X}_1^\theta \textit{ the interpolant of } \theta(X_1), \ldots, \theta(X_n).$$

It then holds that: if X_* is [weakly] signature ergodic to length n, then X_* is [weakly] signature ergodic to length n on Θ.

Using a Glivenko-Cantelli type result yields the following observation of uniform convergence.

For the lemma below, let Θ be as in Assumption 1 and $\kappa_{\mathrm{IC}}^{[m]}$ as in (4.22), and for $n \in \mathbb{N}$ denote by $(\mathcal{I}_{n|j})_{j\in\mathbb{N}}$ any fixed sequence of $[0,1]$-dissections with $|\mathcal{I}_{n|j}| = n$ for all $j \in \mathbb{N}$.

Lemma 4.4.10 (Ergodicity Limit)**.** *Let $X_* = (X_j)_{j\in\mathbb{N}}$ be a discrete time-series which for some m is [weakly] m^{th}-order signature ergodic to some length $n \in \mathbb{N}$ on Θ, and denote*

$$\mathfrak{K}^{m|n|T}(\theta) := \log_{[m]}(\hat{\mathfrak{S}}_T^{m|n}(\theta)) \quad \textit{for} \quad \hat{\mathfrak{S}}_T^{m|n}(\theta) := \frac{1}{T}\sum_{j=1}^{T} \mathfrak{sig}_{[m]}(\hat{X}_j^\theta),$$

$$\textit{and}^{16} \quad \bar{\mathfrak{K}}_{\boldsymbol{i}}^{m|n|T}(\theta) := \frac{\mathfrak{K}_{\boldsymbol{i}}^{m|n|T}(\theta)}{\big(\mathfrak{K}_{11}^{m|n|T}(\theta)\big)^{\eta_1(\boldsymbol{i})/2} \cdot \ldots \cdot \big(\mathfrak{K}_{dd}^{m|n|T}(\theta)\big)^{\eta_d(\boldsymbol{i})/2}}, \quad \boldsymbol{i} \in [d]^\star,$$

for each $\theta \in \Theta$, where $\hat{X}_j^\theta$ is the interpolant of $\theta(X_{n(j-1)+1}), \ldots, \theta(X_{nj})$ along $\mathcal{I}_{n|j}$. For any $m \geq 2, n, T \in \mathbb{N}$ and $\theta \in \Theta$, denote further (recalling Notation 4.2.1)

$$\hat{\kappa}_T^{m|n}(\theta) := \sum_{\nu=2}^{m} \sum_{\boldsymbol{q}\in\mathfrak{C}_\nu} \bar{\mathfrak{K}}_{\boldsymbol{q}}^{m|n|T}(\theta)^2. \tag{4.36}$$

Provided that $\mathbb{E}\big[\sup_{\theta\in\Theta} \|\mathfrak{sig}_k(\hat{X}_1^\theta)\|_k\big] < \infty$ for each $k \in [m]$, it then holds that

$$\bar{\kappa}_{\mathrm{IC}}^{[m]}(\hat{X}_1^\theta) = \lim_{T\to\infty} \hat{\kappa}_T^{m|n}(\theta) \quad \textit{uniformly on} \ \Theta \quad \textit{a.s. [in probability]}. \tag{4.37}$$

A detailed study of the class of (weakly) signature-ergodic stochastic processes is beyond the scope of this thesis, but Section 4.7.7 and the examples below show that the ergodicity assumption (4.34) is met for many popular time series models and stochastic processes.

Definition 4.4.11 (Ergodic Observations)**.** For $\tilde{X} = (\tilde{X}_t)_{t\geq 0}$ a continuous stochastic process in $\mathbb{R}^d$ and $\mathcal{J} = (\mathcal{J}_k)_{k\in\mathbb{N}} \subset 2^{[0,\infty)}$ an exhaustive protocol with base lengths $(n_k)_{k\in\mathbb{N}}$,

[15] Proposition 4.4.9 and Lemma 4.4.10 are proved in the Sections 4.7.5 and 4.7.6, respectively.
[16] Where $\eta_\nu(\boldsymbol{i})$ denotes the number of times the index-entry ν appears in $\boldsymbol{i}$; cf. (4.3).

(i) the pair $(\tilde{X}, \mathcal{J})$ will be called an *ergodic observation* if for almost all $k \in \mathbb{N}$,

$$(\tilde{X}_t)_{t\in\mathcal{J}_k} \quad \text{is} \quad \text{signature ergodic to length } n_k, \tag{4.38}$$

and $(\tilde{X}, \mathcal{J})$ will be called *ergodic** if in addition the spatial support of $\tilde{X}$ is compact;

(ii) the pair $(\tilde{X}, \mathcal{J})$ will be called a *weakly* ergodic observation if for almost all $k \in \mathbb{N}$,

$$(\tilde{X}_t)_{t\in\mathcal{J}_k} \quad \text{is} \quad \text{weakly signature ergodic to length } n_k \tag{4.39}$$

and the running maximum of $|\tilde{X}|$ has finite expectation,[17] i.e. $\mathbb{E}[\sup_{t\in[0,1]}|\tilde{X}_t|] < \infty$.

Given a finite-horizon process $X = (X_t)_{t\in\mathbb{I}}$, we call a pair $(\tilde{X}, \mathcal{J})$ a [weakly] ergodic$^{[*]}$ observation *of* X if it is a [weakly] ergodic$^{[*]}$ observation and $X = (\tilde{X}_t)_{t\in\mathbb{I}}$ almost surely, and we call $(\tilde{X}, \mathcal{J})$ a [weakly] ergodic$^{[*]}$ observation of X *on* $\Theta \subseteq C(D_{\tilde{X}}; \mathbb{R}^d)$ if in addition the pair $\big(\theta(\tilde{X}), \mathcal{J}\big)$ is a [weakly] ergodic$^{[*]}$ observation for each $\theta \in \Theta$ individually.

Examples 4.4.12. Lemma 4.7.7 implies that the [strong resp. weak] ergodicity assumptions (4.33) resp. (4.35) are satisfied by a large number of time series and continuous stochastic processes $\tilde{X}$ (and with it by $\theta(\tilde{X})$ for $\theta \in \Theta$), for the latter by way of (4.38) resp. (4.39) and Lemma 4.7.3 via any protocol $\mathcal{J}$ chosen such that $\tilde{X}_{\mathcal{J}}$ is appropriately stationary. These include, adequate stationarity provided (cf. e.g. Definition 4.7.6),

1. (trivially) all q-dependent time series (e.g. all moving-average processes of finite degree);
2. various linear and related processes such as certain [MC]ARMA, ARCH and GARCH models (see, e.g., [55, 108, 127, 141]);
3. many Markov processes, diffusions and stochastic dynamical systems (e.g. [46, 109, 135, 156]);

see e.g. [21] for an overview. In practice however, infringements of the above (sufficient) conditions for signature ergodicity may typically be innocuous, cf. Section 4.8.

4.4.6 The Consistency Limit

The considerations of Subsections 4.4.2 to 4.4.5 combine to the following consistency result for our ICA-method (Theorem 4.2.3).

Theorem 4.4.13 (Consistency)**.** *Let X, S and Θ be as in Theorem 4.2.3 and Assumption 1, and let $(\tilde{X}, (\mathcal{J}_k)_{k\in\mathbb{N}})$ be an ergodic* [resp. weakly ergodic on Θ] observation of X with base lengths $(n_k)_{k\in\mathbb{N}}$. Suppose that there is $\theta_\star \in \Theta$ such that $\theta_\star(X)$ is IC. Then for any*

[17] The integrability of the running maximum of $|\tilde{X}|$ is discussed in, e.g., [20, Chapter 13] and [107]

error bound $\varepsilon > 0$ there exists a capping threshold $m_0 \geq 2$ such that for any fixed $m \geq m_0$ the following holds: There is a mesh-index $k_0 = k_0(m) \in \mathbb{N}$ such that for any $k \geq k_0$ and any sequence $(\hat{\theta}^\star_T)$ in Θ with the property that, for $\hat{\kappa}_T^{m|n_k}$ as in (4.36) *but computed from $X_* \equiv \tilde{X}_{\mathcal{J}_k}$,*

$$\hat{\kappa}_T^{m|n_k}(\hat{\theta}^\star_T) \ \leq \ \min_{\theta\in\Theta} \hat{\kappa}_T^{m|n_k}(\theta) \ + \ \eta_T \qquad (T \in \mathbb{N}) \tag{4.40}$$

for some $(\eta_T) \subset \mathbb{R}_+$ with $\lim_{T\to\infty} \eta_T = 0$ almost surely [resp. in probability], it holds that

$$\lim_{\tau\to\infty} \max\left\{\sup_{T\geq\tau}\left[\mathrm{dist}_{\|\cdot\|_\infty}(\hat{\theta}^\star_T(X),\, \mathrm{DP}_d\cdot S)\right],\, \varepsilon\right\} \ = \ \varepsilon \tag{4.41}$$

almost surely [resp. in probability]. If $(\tilde{X}, (\mathcal{J}_k)_{k\in\mathbb{N}})$ is ergodic on Θ and the spatial support of X is not necessarily compact, then (4.41) *holds almost surely with the above threshold m_0 depending on the realisation of $\tilde{X}$.*

The above theorem shows that the optimality-based inversion scheme (4.9) is provably robust under a variety of approximations arising in statistical practice.

Remark 4.4.14. In particular,[18] Theorem 4.4.13 gives the following guarantee: Provided that the hyperparameters m (capping threshold) and k (observation frequency) are chosen large enough, the minimizers (4.40) of the empirical signature contrasts $\hat{\kappa}_T^{m|n}$ from (4.36) will, in the infinite-data limit $T \to \infty$, recover each component of the source to an arbitrarily high $\mathbb{I}$-uniform precision (up to order and monotone scaling). In other words, as the grid of observational time-points gets finer ($k \to \infty$) and the length of the observed time series increases ($T \to \infty$), our method produces a signal that gets uniformly closer to the unobserved source. ◆

Proof of Theorem 4.4.13. For brevity, only the statement for $(\tilde{X}, (\mathcal{J}_k))$ ergodic* is proved here; a proof of the remaining non-compact [weakly] ergodic case is given in Section 4.7.8.

So let $(\tilde{X}, (\mathcal{J}_k))$ be an ergodic observation with D_X compact.

Making the (m, k, T)-dependence of each minimizer $\hat{\theta}^\star_T$ in (4.40) explicit by writing $\hat{\theta}^\star_T =: \theta_T^{m|k}$, fix any $\theta_\star \in \Theta$ with $\theta_\star(X)$ IC and observe that assertion (4.41) follows from the claim:

$$\begin{gathered}\forall\,\tilde{\varepsilon} > 0 \,:\, \exists\, m_0 \geq 2 \,:\, \text{for each } m \geq m_0 \text{ there is } k_0 \equiv k_0(m) \text{ such that:}\\ \lim_{\tau\to\infty} \alpha_\tau^{m|k} \vee \tilde{\varepsilon} = \tilde{\varepsilon} \quad \text{a.s.} \quad \text{with} \quad \alpha_\tau^{m|k} := \sup\nolimits_{T\geq\tau} \tilde{d}(\theta_T^{m|k}, \Theta_\star), \quad \text{for each } k \geq k_0;\end{gathered} \tag{4.42}$$

here: $a \vee b := \max\{a, b\}$ and $\tilde{d}$ denotes the topology-inducing metric (4.89) on Θ, and

$$\Theta_\star := \{\beta\circ\theta_\star \mid \beta \in \mathfrak{M}_\Theta\} \quad \text{for} \quad \mathfrak{M}_\Theta := \mathrm{DP}_d(D_{\theta_\star(X)}) \cap \left[\Theta\circ\theta_\star^{-1}\right].$$

[18] Since the maximum norm and the Euclidean norm on $\mathbb{R}^d$ are equivalent.

Indeed: Note first that for each $\theta \in \Theta$, we have that $\theta(X)$ is IC iff $\theta \in \mathfrak{M}_\Theta \cdot \theta_\star$. (For this, recall that if $\theta(X) \equiv (\theta \circ f)(S)$ is IC then $\tilde{\beta} := \theta \circ f \in \mathrm{DP}_d(D_S)$ by [the respective proofs of] Theorems 3.3.3 and 3.3.7; thus also the residual $\beta_\star := \theta_\star \circ f : D_S \to D_{\theta_\star(X)}$ (cf. Lemma 3.2.5 (i)) is in $\mathrm{DP}_d(D_S)$, whence the map $\beta := \theta \circ \theta_\star^{-1} = \tilde{\beta} \circ \beta_\star^{-1}$ is both monomial and in $\Theta \circ \theta_\star^{-1}$, i.e. $\beta \in \mathfrak{M}_\Theta$, as claimed.) In other words (recall Proposition 4.2.2), we have that

$$\Theta_\star = \arg\min_{\theta\in\Theta} \bar{\kappa}_{\mathrm{IC}}(\theta(X)). \tag{4.43}$$

(In the following, we import the setting and notation of Remark 4.7.8 for usage below.)

Provided now that (4.42) holds, we find that for every $\varepsilon > 0$ there is $m_0 \geq 2$ with the property that each capping index $m \geq m_0$ comes with a mesh-threshold $k_0 = k_0(m) \in \mathbb{N}$ which is such that for each observation at mesh-level $k \geq k_0$ we have with probability one that there is a sequence of transformations $(\theta_T)_{T\in\mathbb{N}}$ in $\Theta_\star$ such that almost surely:

$$\big\|\theta_T^{m|k}(X) - \theta_T(X)\big\|_\infty \;\leq\; \varepsilon \quad \text{ for almost all } T \in \mathbb{N}, \quad \text{where } \; \big(\theta_T(X)\big) \subset \mathrm{DP}_d \cdot S \tag{4.44}$$

with probability one due to (4.43) and Theorem 4.2.3. This readily implies (4.41) as desired.

To derive (4.44) from (4.42), let $\varepsilon > 0$ be arbitrary, assuming $\varepsilon < 1$ wlog, and note that $D_X \subseteq K_{\nu_0}$ for some $\nu_0 \in \mathbb{N}$ as D_X is compact. Observe then that (4.90) provides the inclusion $B_\varepsilon^{d_{\nu_0}}(\theta) \supseteq B_{\varepsilon/2}^{\rho_{\nu_0}}(\theta)$ for each $\theta \in \Theta$, while the definition of $\tilde{d}$ yields $\tilde{B}_{\tilde{\varepsilon}}(\theta) \subseteq B_{\varepsilon/2}^{\rho_{\nu_0}}(\theta)$ for $\tilde{\varepsilon} := 2^{-\nu_0}\varepsilon/2$, cf. (4.89). Given (4.42), this $\tilde{\varepsilon}$ comes with associated $m_0, m, k_0, k \in \mathbb{N}$ and a $\mathbb{P}$-full set $\Omega' \equiv \Omega'_{m,k} \in$ [19]$\mathscr{F}$ such that

$$\alpha^{m|k}(\omega) := \lim_{\tau\to\infty} \sup\nolimits_{T\geq\tau} \tilde{d}(\theta_T^{m|k}(\omega), \Theta_\star) \;\leq\; \tilde{\varepsilon}/2 \qquad \text{for each } \; \omega \in \Omega'$$

(note: $\alpha^{m|k}$ exists as $(\alpha_\tau^{m|k})_{\tau\in\mathbb{N}}$ is monotone and bounded). Thus, for each $\omega \in \Omega' \cap \Omega''$ (for $\Omega'' \in \mathscr{F}$ the $\mathbb{P}$-full set on which the traces of X are all contained in D_X; Lemma 3.2.5 (ii)) there is $\tau_0\,(= \tau_0(\omega)) \in \mathbb{N}$ together with a sequence $(\theta_T)_{T\in\mathbb{N}}\,(\equiv (\theta_T(\omega))_{T\in\mathbb{N}}) \subset \Theta_\star$ such that: $\theta_T^{m|k}(\omega) \subseteq \tilde{B}_{\tilde{\varepsilon}}(\theta_T)\,\big(\subseteq B_{\varepsilon/2}^{\rho_{\nu_0}}(\theta_T) \subseteq B_\varepsilon^{d_{\nu_0}}(\theta_T)\big)$ for each $T \geq \tau_0$, and hence

$$\big\|\theta_T^{m|k}(\omega)(X(\omega)) - \theta_T(X(\omega))\big\|_\infty \leq \big\|\theta_T^{m|k}(\omega) - \theta_T\big\|_{K_{\nu_0}} = d_{\nu_0}(\theta_T^{m|k}(\omega), \theta_T) \;\leq\; \varepsilon \tag{4.45}$$

for each $T \geq \tau_0$, which gives (4.44) as desired. To prove (4.42) next, let $\tilde{\varepsilon} > 0$ be arbitrary.

Then for $\kappa_{m,k}(\theta) := \widehat{Q}_m(\hat{X}^\theta_{\mathcal{I}_1^{(k)}})$ as in (4.30) with $\mathcal{J} =: (\mathcal{J}_k \equiv \sqcup_{\nu=1}^\infty \mathcal{I}_\nu^{(k)} \mid k \in \mathbb{N})$ the given protocol under consideration, there is $m_0 \geq 2$ such that for each $m \geq m_0$ it holds that

$$\exists\, k_0\,(\equiv k_0(m)) \in \mathbb{N} \;:\; \quad \arg\min_{\theta\in\Theta} \kappa_{m,k}(\theta) \;\subseteq\; \Theta_\star^{\tilde{\varepsilon}/2} := \bigcup\nolimits_{\theta\in\Theta_\star} \tilde{B}_{\tilde{\varepsilon}/2}(\theta), \quad \forall\, k \geq k_0. \tag{4.46}$$

[19] For convenience, we may assume the underlying probability space $(\Omega, \mathscr{F}, \mathbb{P})$ to be complete.

Indeed: The identity (4.43) (together with the fact that $\bar{\kappa}_{\mathrm{IC}}(\theta_\star(X)) = 0$) implies that $\zeta_{\tilde{\varepsilon}} := \min_{\theta\in\Theta\setminus C_{\tilde{\varepsilon}}} Q(\theta) > 0$ for[20] $C_{\tilde{\varepsilon}} := \Theta_\star^{\tilde{\varepsilon}/2} \cap \Theta$ and Q as in (4.25), while Lemma 4.4.3 (ii) provides an $m_0 \geq 2$ with $\sup_{m\geq m_0} \|Q - Q_m\|_\Theta < \zeta_{\tilde{\varepsilon}}/2$. Fixing any $m \geq m_0$ and using that $(\mathcal{J}_k)_{k\in\mathbb{N}}$ is exhaustive (whence the sequence $(\mathcal{I}_1^{(k)})_{k\in\mathbb{N}}$ is refined), Lemma 4.4.5 yields some $k_0 \equiv k_0(m) \in \mathbb{N}$ with $\sup_{k\geq k_0} \|Q_m - \kappa_{m,k}\|_\Theta < \zeta_{\tilde{\varepsilon}}/2$. Hence for any fixed $k \geq k_0$ and each $\tilde{\theta} \in \Theta$ with $\kappa_{m,k}(\tilde{\theta}) = \min_{\theta\in\Theta} \kappa_{m,k}(\theta)$, it holds that

$$Q(\tilde{\theta}) = \kappa_{m,k}(\tilde{\theta}) + (Q - \kappa_{m,k})(\tilde{\theta}) \ < \ \zeta_{\tilde{\varepsilon}}/2 + \zeta_{\tilde{\varepsilon}}/2 = \zeta_{\tilde{\varepsilon}} \quad \text{and hence} \quad \tilde{\theta} \in \Theta_\star^{\tilde{\varepsilon}/2},$$

where we used the fact that $0 \leq \min_{\theta\in\Theta} \kappa_{m,k}(\theta) \leq \kappa_{m,k}(\theta_\star) = 0$ (i.e., the argmins of $\kappa_{m,k}$ on Θ coincide with its roots), where this last property follows from Proposition 4.2.2 and the obvious fact that if $\theta_\star(X)$ is IC then so is $\hat{X}^{\theta_\star}_{\mathcal{I}_1^{(k)}}$. Together with (4.46) the identity $\arg\min_\Theta(\kappa_{m,k}) = \{\theta \in \Theta \mid \kappa_{m,k}(\theta) = 0\} =: \mathcal{N}(\kappa_{m,k})$ now implies that, for $\bar{\kappa}_{m,k}(\theta) := \bar{\kappa}_{\mathrm{IC}}^{[m]}(\hat{X}^\theta_{\mathcal{I}_1^{(k)}})$ as in (4.22),

$$\mathcal{M} := \arg\min_{\theta\in\Theta} \bar{\kappa}_{m,k}(\theta) \ \subseteq \ \Theta_\star^{\tilde{\varepsilon}/2} \quad \text{for each} \ \ k \geq k_0, \tag{4.47}$$

because $\arg\min_\Theta \bar{\kappa}_{m,k} = \mathcal{N}(\bar{\kappa}_{m,k})$ (as above) and the zero sets $\mathcal{N}(\bar{\kappa}_{m,k})$ and $\mathcal{N}(\kappa_{m,k})$ coincide (by Definition 4.1.2). Next we claim that, for m and k as above,

$$\lim_{\tau\to\infty} \sup\nolimits_{T\geq\tau} \tilde{d}(\theta_T^{m|k}, \mathcal{M}) \ = \ 0 \qquad \text{almost surely,} \tag{4.48}$$

which by way of (4.47) implies (4.42) as desired. To see (4.48), observe first that

$$\lim_{T\to\infty} \bar{\kappa}_{m,k}(\theta_T^\star) \ = \ 0 \quad \text{almost surely,} \quad \text{with} \ \ (\theta_T^\star) \equiv (\theta_T^{m|k}) \tag{4.49}$$

as in (4.48), which due to $(4.37)|_{X_*\equiv\tilde{X}_{\mathcal{J}_k}}$ follows by the same arguments that led to (4.60). Using proof by contradiction, assume now that (4.48) does not hold. Then, pointwise on an event of positive probability, there will be $\delta_0 > 0$ together with a subsequence $(T_j)_{j\in\mathbb{N}} \subseteq \mathbb{N}$ such that $\tilde{d}(\theta^\star_{T_j}, \mathcal{M}) \geq \delta_0$ for each $j \in \mathbb{N}$. But as Θ is compact, we (upon passing to a convergent subsequence) may assume that $(\theta^\star_{T_j})_{j\in\mathbb{N}}$ converges to some $\theta_0 \in \Theta$. Then by continuity $\bar{\kappa}_{m,k}(\theta_0) = \lim_{j\to\infty} \bar{\kappa}_{m,k}(\theta^\star_{T_j})$, whence $\bar{\kappa}_{m,k}(\theta_0) = 0$ by (4.49) and thus $\theta_0 \in \mathcal{M}$. Yet the latter contradicts that $(\theta^\star_{T_j})_{j\in\mathbb{N}}$ is bounded away from $\mathcal{M}$ (by δ_0), proving (4.48). □

4.5 Consistency in the Discrete-Time Case

Throughout Section 4.4 we assumed the data-generating signal X from (4.19) to be continuous-time on $[0,1]$. If this is not the case, then the 'refinement assumption' $\lim_{k\to\infty} \|\mathcal{I}_1^{(k)}\| = 0$ in (4.19) can be dropped and naturally replaced by the compensating assumption that

$$\text{there is} \ \ k_0 \quad \text{s.t.} \quad \text{the observation } (X_t)_{t\in\mathcal{I}_1^{(k_0)}} \text{ is } \ \bar{\alpha}\text{-, } \bar{\beta}\text{- or } \bar{\gamma}\text{-contrastive.}$$

[20] Provided that $\tilde{\varepsilon}$ is small enough such that $C_{\tilde{\varepsilon}} \subsetneq \Theta$, which can be assumed without loss of generality.

This renders Sect. 4.4.4 void and removes the necessity to, as in Sect. 4.4.1, consider k-dependent protocols $\mathcal{J}_k$ with ever growing base lengths $|\mathcal{I}_1^{(k)}|$. Sections 4.4.3, 4.4.5, however, stay applicable as stated and Theorem 4.4.13 remains valid – by the same proof – up to the following straightforward modifications[21].

4.5.0.1 Consistency in the Discrete-Time Case

In Section 4.4 we assumed that the data-generating signal X from (4.19) is continuous-time on $[0,1]$. If this is not the case, then the 'refinement assumption' $\lim_{k\to\infty}\|\mathcal{I}_1^{(k)}\| = 0$ in (4.19) can be dropped and naturally replaced by the compensating assumption that

$$\text{there is } \; k_0 \quad \text{s.t.} \quad \text{the observation } (X_t)_{t\in\mathcal{I}_1^{(k_0)}} \text{ is } \; \bar{\alpha}\text{-}, \bar{\beta}\text{- or } \bar{\gamma}\text{-contrastive.}$$

This renders Section 4.4.4 void and removes the necessity to consider k-dependent protocols $\mathcal{J}_k$ with ever growing base lengths $|\mathcal{I}_1^{(k)}|$ (Section 4.4.1), while Sections 4.4.3 and 4.4.5 stay applicable as stated and Theorem 4.4.13 remains valid – by the same proof – up to the following straightforward modifications. (In essence, we only need to remove the k-dependence from Section 4.4.)

Let X_* and S_* be as in (3.112) with $\tilde{X}_* = (\tilde{X}_j)_{j\in\mathcal{J}}$ for some $\mathcal{J}\subseteq\mathbb{Z}$ such that

$$\mathcal{J} = \bigsqcup\nolimits_{\nu\in\mathbb{N}}\mathcal{I}_\nu \quad \text{with} \quad \mathcal{I}_1 < \mathcal{I}_2 < \ldots \;\; \text{s.t.} \;\; (X_j)_{j\in\mathcal{I}_1} =: (\tilde{X}_j)_{j\in\mathcal{I}_1} \text{ is } \; \bar{\alpha}\text{-}, \bar{\beta}\text{- or } \bar{\gamma}\text{-contrastive.}$$

Let further Θ be as in Theorem 4.2.3 and such that Assumption 1 holds[22] for $(X_j)_{j\in\mathcal{I}_1}$. We then call the discrete process $\tilde{X}_*$ an *ergodic observation of* X_* if $\tilde{X}_*$ is signature ergodic to length $|\mathcal{I}_1| =: n$; the remaining notions of Definition 4.4.11 are adopted analogously.

Let finally $\|\cdot\|_{[n]}$ denote the uniform norm on $\mathbb{R}^{d\times n}$ (so that $\|(x_j)\|_{[n]} = \max_{j\in[n]}|x_j|$).

Theorem 4.5.1. *Let X_*, S_* and Θ, $\tilde{X}_*$ be as above, and suppose that there is $\theta_\star\in\Theta$ such that $\theta_\star(X_*)$ is IC. Suppose further that $\tilde{X}_*$ is an ergodic* [resp. weakly ergodic on Θ] observation of X_*. Then for any error bound $\varepsilon>0$ there exists a capping threshold $m_0\geq 2$ such that for any fixed $m\geq m_0$ the following holds: For every sequence $(\hat{\theta}_T^\star)$ in Θ such that*

$$\hat{\kappa}_T^{m|n}(\hat{\theta}_T^\star) \;\leq\; \min_{\theta\in\Theta}\hat{\kappa}_T^{m|n}(\theta) \;+\; \eta_T \quad (T\in\mathbb{N}), \quad \textit{for} \quad \hat{\kappa}_T^{m|n} \textit{ as in } (4.36)|_{X_*:=\tilde{X}_*} \textit{ and } \; n := |\mathcal{I}_1|$$

and some $(\eta_T)\subset\mathbb{R}_+$ with $\lim_{T\to\infty}\eta_T = 0$ almost surely [resp. in probability], it holds that

$$\lim_{\tau\to\infty}\max\left\{\sup_{T\geq\tau}\left[\operatorname{dist}_{\|\cdot\|_{[n]}}(\hat{\theta}_T^\star(X_*),\, \mathrm{DP}_d\cdot S_*)\right],\, \varepsilon\right\} \;=\; \varepsilon$$

[21] In essence, we only need to remove the k-dependence from Section 4.4.
[22] Following the piecewise-linear interpolation of $(X_j)_{j\in\mathcal{I}_1}$ as in point (ii) (b) above.

almost surely [resp. in probability]. If $\tilde{X}_*$ is ergodic on Θ and the spatial support of X_* is not necessarily compact, then (4.41) *holds almost surely with the above threshold m_0 depending on the realisation of $\tilde{X}_*$.*

4.6 Algorithm

The computational procedures of Section 4.4 (and Section 4.5.0.1) can be summarized into the following practical algorithm whose consistency is established by Theorem 4.4.13.

Algorithm: Nonlinear ICA via Signature Cumulants

Goal: For $X = f(S)$ with a contin.-/discrete-time process S in $\mathbb{R}^d$, invert X for S.

Hyperparameters: candidate nonlinearities Θ, capping order m_0 (as in (4.22)), base length n ($:= |\mathcal{I}_1^{(k)}|$, as in (4.18)), observation horizon T (as in (4.36), or $T \equiv \max\{j \geq 1 \mid \mathcal{I}_j^{(k)} \leq T_k\}$ for T_k as in Rem. 4.4.1).[a]

1. **Input:** sample observation $\mathfrak{x} \equiv (\mathfrak{x}_j)$ of X (as in (4.18), with index (k) omitted).
2. Compute the estimated contrast $\hat{\phi} := \hat{\kappa}_T^{m_0|n}$ as in (4.36), that is compute

$$\hat{\phi}(\theta) = \sum_{\nu=2}^{m_0} \sum_{\boldsymbol{q} \in \mathfrak{C}_\nu} \bar{\varphi}_{\boldsymbol{q}}(\theta)^2 \quad \text{with}^{23} \quad \bar{\varphi}_{\boldsymbol{i}}(\theta) := \frac{\varphi_{\boldsymbol{i}}(\theta)}{\varphi_{11}(\theta)^{\frac{\eta_1(\boldsymbol{i})}{2}} \cdots \varphi_{dd}(\theta)^{\frac{\eta_d(\boldsymbol{i})}{2}}}$$

$$\text{where} \qquad \varphi_{\boldsymbol{i}}(\theta) := \langle \hat{\mathscr{C}}(\theta), \boldsymbol{i} \rangle \quad (\boldsymbol{i} \in [d]^\star) \tag{4.50a}$$

$$\text{and} \qquad \hat{\mathscr{C}}(\theta) := \log_{[m_0]}\left[T^{-1} \sum_{j=1}^{T} \mathfrak{sig}_{[m_0]}\big(\hat{\mathfrak{x}}_j^\theta\big)\right], \tag{4.50b}$$

for $\hat{\mathfrak{x}}_j^\theta$ the piecewise-linear interpolation of the data $\theta(\mathfrak{x}_j) \equiv \big(\theta(X_t(\omega)) \mid t \in \mathcal{I}_j\big)$.

3. Compute a minimiser $\theta_\star$ of $\hat{\phi}$ over Θ, that is find

$$\theta_\star \in \underset{\theta \in \Theta}{\arg\min}\ \hat{\phi}(\theta). \tag{4.51}$$

4. Compute the estimated source realisation $\hat{\mathfrak{s}} := \theta_\star(\mathfrak{x}) \equiv \big(\theta_\star(X_t(\omega)) \mid t \in \mathcal{I}_1\big)$.
5. **Output:** $\hat{\mathfrak{s}}$ and $\theta_\star$.

[a] Optional: weights $(w_{\boldsymbol{q}})$ as in Remark 4.2.4 (v).

The above algorithm involves two independent subroutines, namely the computation of the free logarithm of averages of signatures of piecewise-linearly interpolated data batches (4.50b) followed by the subsequent extracion of its relevant [cross-shuffle-indexed] coeffi-

[23] Recall Notation 4.2.1 and that $\eta_\nu(\boldsymbol{i})$ equals the number of times the index-value ν appears in $\boldsymbol{i}$, for example $\eta_3(123433235) = 4$. Recall further that $\langle \hat{\mathscr{C}}(\theta), \boldsymbol{i} \rangle$ denotes the $\boldsymbol{i}^{\text{th}}$ entry of the multiindexed list $\hat{\mathscr{C}}(\theta)$ (cf. (2.3.2)).

cients (4.50a), and the optimisation (4.51) of the contrast $\hat{\phi}$ over a given set Θ of candidate demixing transformations. The first of these routines can be conveniently implemented by use of the functionality provided with specialised signature libraries such as [95], while the second task can be performed with great flexibility by choosing Θ as an artificial neural network that has $\hat{\phi}$ as its loss function. The minimiser $\theta_\star$ can then be learnt as an optimal network configuration reached by training $(\Theta, \hat{\phi})$ via backpropagation, cf. e.g. Remark 4.2.4 (ii) and Section 4.8.3.

Several example applications of the above method are detailed in Section 4.8 and on the public repository [140] where the above algorithm is also implemented, including a differentiable (i.e. backpropagatable) implementation of the contrast function $\hat{\phi}$.

4.7 Lemmas and Proofs for Section 4.4

The following contains auxiliary considerations and proofs for Section 4.7. We make tacit use of the notation from Section 2.3.3.

4.7.1 Basic Cross-Shuffle Combinatorics

Using the notation of Sections 4, 4.4.3 and 2.3.3 throughout, this subsection considers the family of *cross-shuffles* $\mathfrak{C} \equiv \bigsqcup_{k=2}^{d} \mathcal{W}_k = \bigsqcup_{\nu=2}^{\infty} \mathfrak{C}_\nu$.

By definition (4.6) of the shuffle product, each element $\boldsymbol{q} \in \mathfrak{C}_\nu$ (seen as an element of (2.3.2)) is a homogeneous polynomial of degree ν whose monomial coefficients are all 1, i.e. there is $c_{\boldsymbol{q}} \in \mathbb{N}$ such that $\boldsymbol{q} = \boldsymbol{q}_1 + \cdots + \boldsymbol{q}_{c_{\boldsymbol{q}}}$ with $\boldsymbol{q}_j \in [d]^*$ for each $j \in [c_{\boldsymbol{q}}]$. Partitioning

$$\mathfrak{C}_\nu = \bigsqcup_{k=2}^{d} \mathfrak{C}_{\nu|k} \quad \text{for} \quad \mathfrak{C}_{\nu|k} := \mathfrak{C}_\nu \cap \mathcal{W}_k,$$

the definition of $\mathcal{W}_k$ yields that for any given $\boldsymbol{q} \in \mathfrak{C}_\nu$ the pair $\vartheta_{\boldsymbol{q}} \equiv (k_{\boldsymbol{q}}, \mu_{\boldsymbol{q}}) \in [d]_{\geq 2} \times [m-1]$, with $k_{\boldsymbol{q}} := \max\{i \in [d] \mid i \in \boldsymbol{q}_1\}$ being the largest letter contained in $\boldsymbol{q}$ and $\mu_{\boldsymbol{q}} := \sum_{i \in \boldsymbol{q}_1} \delta_{i,k_{\boldsymbol{q}}}$ denoting the number of times this largest letter appears in one (and hence any) of the monomials of $\boldsymbol{q}$, determines $\boldsymbol{q}$ uniquely (in $\mathfrak{C}_\nu$) up to a $\sqcup\!\sqcup$-left-factor of word length $\nu - \mu_{\boldsymbol{q}}$. (Indeed: Given $\boldsymbol{q} \in \mathfrak{C}_\nu$, the number $k_{\boldsymbol{q}} \in \mathbb{N}$ is the (unique) index s.t. $\boldsymbol{q} \in \mathfrak{C}_{\nu|k_{\boldsymbol{q}}} \subset \mathcal{W}_{k_{\boldsymbol{q}}}$, whence $\boldsymbol{q} = \boldsymbol{w} \sqcup\!\sqcup (k_{\boldsymbol{q}})^{*\mu_{\boldsymbol{q}}}$ for some $\boldsymbol{w} \in [k_{\boldsymbol{q}} - 1]^*$ with $|\boldsymbol{w}| = \nu - \mu_{\boldsymbol{q}}$.)

Since the shuffle product (4.6) of two words $\boldsymbol{i}, \boldsymbol{j} \in [d]^*$ is precisely the sum over the $c_{\boldsymbol{i} \sqcup\!\sqcup \boldsymbol{j}}$ $\left(= \frac{(|\boldsymbol{i}|+|\boldsymbol{j}|)!}{|\boldsymbol{i}|!|\boldsymbol{j}|!}\right)$ ways of *interleaving* $\boldsymbol{i}$ and $\boldsymbol{j}$, any two monomials in $\boldsymbol{i} \sqcup\!\sqcup \boldsymbol{j}$ are composed of exactly the same letters and differ only in the order in which their letters appear. Consequently,

(a) any two $\boldsymbol{q}, \boldsymbol{q}' \in \mathfrak{C}_\nu$ have a monomial in common iff $\boldsymbol{q} = \boldsymbol{q}'$;

(b) given any $\boldsymbol{q} \in \mathfrak{C}_\nu$ with its unique (up to the order of summands) decomposition $\boldsymbol{q} = \boldsymbol{q}_1 + \ldots + \boldsymbol{q}_{c_{\boldsymbol{q}}}$ into monic monomials $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_{c_{\boldsymbol{q}}} \in [d]^*$, these monomials $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_{c_{\boldsymbol{q}}}$ are pairwise distinct.

(Note that point (b) follows inductively: Let $\boldsymbol{q} \equiv \boldsymbol{w} ⧢ (k_{\boldsymbol{q}})^{*\mu_{\boldsymbol{q}}} \in \mathfrak{C}_\nu$ with $(k_{\boldsymbol{q}}, \mu_{\boldsymbol{q}}) = \vartheta_{\boldsymbol{q}}$. The assertion clearly holds if $\mu_{\boldsymbol{q}} = 1$ or (by symmetry) $|\boldsymbol{w}| = 1$. Fixing any $\boldsymbol{q}$ as above, assume that (b) holds for any $\boldsymbol{r} \equiv \boldsymbol{w}' ⧢ (k_{\boldsymbol{q}})^{*\mu_{\boldsymbol{r}}} \in \mathfrak{C}$ with $\mu_{\boldsymbol{r}} = \mu_{\boldsymbol{q}} - 1$ or $|\boldsymbol{w}'| = |\boldsymbol{w}| - 1$. Note that

$$\boldsymbol{q} = (\boldsymbol{w}' * \mathtt{i}) ⧢ (k_{\boldsymbol{q}})^{*(\mu'_{\boldsymbol{q}}+1)} = \boldsymbol{r}_0 * \mathtt{i} + \boldsymbol{r}_1 * k_{\boldsymbol{q}} \qquad (\boldsymbol{w} =: \boldsymbol{w}' * \mathtt{i}, \ \mathtt{i} \in [k_{\boldsymbol{q}} - 1] \setminus \{\epsilon\})$$

for the polynomials $\boldsymbol{r}_0 := \boldsymbol{w}' ⧢ (k_{\boldsymbol{q}})^{*\mu_{\boldsymbol{q}}}$ and $\boldsymbol{r}_1 := \boldsymbol{w} ⧢ (k_{\boldsymbol{q}})^{*(\mu_{\boldsymbol{q}}-1)}$ (by the recursive formulation of the shuffle product, e.g. [132, p. 25 f.]). Now since the monomials of $\boldsymbol{r}_0$ and $\boldsymbol{r}_1$ are all monic and pairwise distinct by induction hypothesis, the same applies to $\boldsymbol{r}_0 * \mathtt{i}$ and $\boldsymbol{r}_1 * k_{\boldsymbol{q}}$ and, hence (as $\mathtt{i} \neq k_{\boldsymbol{q}}$), to $\boldsymbol{q}$. Thus by induction, assertion (b) holds for $\boldsymbol{q}$ as desired.)

4.7.2 Proof of Lemma 4.4.3

Proof of Lemma 4.4.3. (i): For $\nu \geq 2$ fixed, consider the function $q_\nu : \Theta \to V_\nu$ given by

$$q_\nu(\theta) := \sum_{\boldsymbol{q} \in \mathfrak{C}_\nu} \kappa_{\boldsymbol{q}}(\theta(X)) \cdot \frac{\boldsymbol{q}}{\sqrt{c_{\boldsymbol{q}}}} \tag{4.52}$$

with $c_{\boldsymbol{q}}$ the number of monomials in $\boldsymbol{q}$ (cf. Remark 4.7.1). As the sets $\{c_{\boldsymbol{q}}^{-1/2} \cdot \boldsymbol{q} \mid \boldsymbol{q} \in \mathfrak{C}_\nu\}$ are each finite and orthonormal wrt. the Euclidean structure on V_ν (cf. Sect. 2.3.3), we have that $Q_m = \sum_{\nu=2}^m \|q_\nu\|_\nu^2$ for each $m \geq 2$ and thus obtain the continuity of Q_m from the continuity of (4.52). To convince ourselves of the latter, fix any index $\boldsymbol{q} \in \mathfrak{C}_\nu$ and let $(\theta_j)_{j \in \mathbb{N}}$ be an arbitrary convergent sequence in Θ, with limit $\lim_{j\to\infty} \theta_j =: \tilde{\theta}$. By a classical interpolation inequality (see e.g. [54, Proposition 5.5. (i)]) we for any $2 > p' > p$ have that

$$\|\tilde{\theta}(X) - \theta_j(X)\|_{p'\text{-var}} \leq C \cdot \|\tilde{\theta}(X) - \theta_j(X)\|_\infty^{1-p/p'} \longrightarrow 0 \quad \text{a.s.} \quad (\text{as } j \to \infty) \tag{4.53}$$

for the a.s. finite (by (4.26)) random variable $C := 2^{1-p/p'} \sup_{j\geq 1} [\|\tilde{\theta}(X) - \theta_j(X)\|_{p\text{-var}}]^{p/p'}$, where the convergence in (4.53) then follows by the compact convergence $\theta_j \to \tilde{\theta}$ and the fact that almost every realisation of X has a compact trace in D_X (Lemma 3.2.5 (ii)).

Hence by the p'-variation continuity of $\mathfrak{sig}$ (Lemma 2.3.5 (ii)) followed by dominated convergence (cf. (4.27)) and the fact that $\log_{[\nu]} \equiv \pi_{[\nu]} \circ \log = \log_{[\nu]} \circ \pi_{[\nu]}$ is continuous (Lemma 2.3.5 (v)), we see that the convergence (4.53) implies that, for any $\boldsymbol{q} \in \mathfrak{C}_\nu$,

$$\begin{aligned} \kappa_{\boldsymbol{q}}(\theta_j(X)) &\overset{\text{def}}{=} \langle \log[\mathbb{E}[\mathfrak{sig}(\theta_j(X))]], \boldsymbol{q} \rangle = \langle \log_{[\nu]}[\mathbb{E}[(\pi_{[\nu]} \circ \mathfrak{sig})(\theta_j(X))]], \boldsymbol{q} \rangle \\ &\longrightarrow \langle \log_{[\nu]}[\mathbb{E}[(\pi_{[\nu]} \circ \mathfrak{sig})(\tilde{\theta}(X))]], \boldsymbol{q} \rangle = \kappa_{\boldsymbol{q}}(\tilde{\theta}(X)) \quad \text{as } j \to \infty. \end{aligned} \tag{4.54}$$

Since our topology on Θ is metrizable, the sequential convergence (4.54) characterizes the continuity of $\Theta \ni \theta \mapsto \kappa_{\boldsymbol{q}}(\theta(X))$, which yields that (4.52) (hence Q_m) is continuous as desired. The continuity of Q thus follows from assertion (ii) of this lemma, i.e. from the claim

$$Q_m \to Q \quad \text{uniformly on } \Theta \qquad \text{as } m \to \infty. \tag{4.55}$$

(ii) : To see that (4.55) holds, observe that since $\sup_{\theta\in\Theta} |||\mathfrak{S}(\theta(X)) - 1|||_\lambda \leq 1$ for some $\lambda > 2$ by assumption, Lemma 2.3.5 (v) yields that the set $\mathcal{L} := \log(\{\mathfrak{S}(\theta(X)) \mid \theta \in \Theta\}) \equiv \{\ell(\theta) \mid \theta \in \Theta\}$ is $|||\cdot|||_\rho$-bounded for some $\rho > 1$. Hence by Lemma 2.3.5 (vi),

$$\varsigma_m := \sup_{\boldsymbol{\ell}\in\mathcal{L}} \sum_{\nu>m} \|\pi_\nu(\boldsymbol{\ell})\|_\nu \longrightarrow 0 \qquad \text{as } m \to \infty. \tag{4.56}$$

Writing now $\boldsymbol{q} = \boldsymbol{w}_1(\boldsymbol{q}) + \ldots + \boldsymbol{w}_{c_{\boldsymbol{q}}}(\boldsymbol{q})$ for the decomposition of a polynomial $\boldsymbol{q} \in \mathfrak{C}_\nu$ into its (monic) monomials $\boldsymbol{w}_j(\boldsymbol{q}) \in [d]^\star_\nu$ (cf. Remark 4.7.1 (b)), we have that for each $\boldsymbol{q} \in \mathfrak{C}_\nu$ the monomials $\boldsymbol{w}_1(\boldsymbol{q}), \ldots, \boldsymbol{w}_{c_{\boldsymbol{q}}}(\boldsymbol{q})$ are pairwise distinct (Rem. 4.7.1 (b)), and further that the union $\bigcup_{\boldsymbol{q}\in\mathfrak{C}_\nu}\{\boldsymbol{w}_1(\boldsymbol{q}), \ldots, \boldsymbol{w}_{c_{\boldsymbol{q}}}(\boldsymbol{q})\} \subset [d]^*_\nu$ is disjoint (Rem. 4.7.1 (a)). Hence and since[24]

$$\kappa_{\boldsymbol{q}}^2 \overset{\text{def}}{=} \Big(\kappa_{\boldsymbol{w}_1(\boldsymbol{q})} + \ldots + \kappa_{\boldsymbol{w}_{c_{\boldsymbol{q}}}(\boldsymbol{q})}\Big)^2 \leq c_{\boldsymbol{q}} \cdot \sum_{j=1}^{c_{\boldsymbol{q}}} \kappa^2_{\boldsymbol{w}_j(\boldsymbol{q})} \tag{4.57}$$

by the Cauchy-Schwarz inequality, we for each $m \geq 2$ obtain the estimate

$$\begin{aligned} \|Q - Q_m\|_\Theta &\leq \sup_{\theta\in\Theta} \sum_{\nu>m} \sum_{\boldsymbol{q}\in\mathfrak{C}_\nu} c_{\boldsymbol{q}}^{-1} \cdot \kappa_{\boldsymbol{q}}\big(\theta(X)\big)^2 \\ &\leq \sup_{\theta\in\Theta} \sum_{\nu>m} \sum_{\boldsymbol{w}\in[d]^*_\nu} \kappa_{\boldsymbol{w}}\big(\theta(X)\big)^2 = \sup_{\boldsymbol{\ell}\in\mathcal{L}} \sum_{\nu>m} \|\pi_\nu(\boldsymbol{\ell})\|^2_\nu. \end{aligned} \tag{4.58}$$

Hence, and since $\lim_{\nu\to\infty} \|\pi_\nu(\boldsymbol{\ell})\|_\nu = 0$ uniformly on $\mathcal{L}$ by (4.56), there will be an $m_0 \geq 2$ such that $\sup_{\boldsymbol{\ell}\in\mathcal{L},\, \nu\geq m_0} \|\pi_\nu(\boldsymbol{\ell})\|_\nu < 1$ and therefore, by (4.58), $\|Q - Q_m\|_\Theta \leq \varsigma_m$ for all $m \geq m_0$, implying (4.55) as claimed.

(iii) : Let $\theta_\star \in \Theta$ be as above, and $\varepsilon > 0$ be arbitrary. Since Θ is compact so is its closed subset[25] $C_\varepsilon := \Theta \setminus B_\varepsilon(\theta_\star)$, and for $\zeta_\varepsilon := \inf_{\theta\in C_\varepsilon} Q(\theta) = Q(\theta_\epsilon) > Q(\theta_\star)$ (for some $\theta_\varepsilon \in C_\varepsilon$; recall that Q is continuous) and any $\theta \in \Theta$ we have the obvious implication that:

$$\text{if} \quad Q(\theta) < \zeta_\varepsilon, \qquad \text{then} \quad \theta \in B_\varepsilon(\theta_\star). \tag{4.59}$$

Let now $(\theta^\star_m) \subset \Theta$ be a minimising sequence of the required kind. As then $Q(\theta_\star) \leq Q(\theta^\star_m)$ and $Q_m(\theta^\star_m) \leq Q_m(\theta_\star) + \eta_m$ for each $m \geq 2$, we find that $Q(\theta_\star) \leq Q_m(\theta^\star_m) + \big(Q(\theta^\star_m) -$

[24] To ease notation, we in (4.57) drop the argument of the cumulants, i.e. denote $\kappa_q \equiv \kappa_q(\theta(X))$.

[25] The topology (of compact convergence) on Θ is metrizable (cf. Remark 4.7.8), and $B_\varepsilon(\theta_\star)$ denotes the open ball of radius ε defined wrt. any applicable metric on Θ.

$Q_m(\theta^\star_m)) \le Q_m(\theta_\star) + (Q(\theta^\star_m) - Q_m(\theta^\star_m) + \eta_m) = Q(\theta_\star) + r_m$ for $r_m := (Q_m(\theta_\star) - Q(\theta_\star) + Q(\theta^\star_m) - Q_m(\theta^\star_m) + \eta_m)$. Hence $Q(\theta_\star) \le Q(\theta^\star_m) \le Q(\theta_\star) + r_m$ and therefore

$$\lim_{m\to\infty} |Q(\theta_\star) - Q(\theta^\star_m)| \;\le\; \lim_{m\to\infty} r_m \;=\; 0 \qquad \text{(a.s.)}, \tag{4.60}$$

where the last identity is due to the uniform convergence (ii) (and our assumption on (η_m)). Consequently $Q(\theta^\star_m) < \zeta_\varepsilon$ for almost all m, which in light of (4.59) implies that $\theta^\star_m \in B_\varepsilon(\theta_\star)$ for almost all m (a.s.), as desired. □

4.7.3 Linear Interpolation of Discrete-Time Data

Let $\mathbb{I}$ be a compact interval; say $\mathbb{I} = [0,1]$ wlog. Any dissection $\mathcal{I} \equiv \{t_0, \ldots, t_{n-1} \mid t_0 < \ldots < t_{n-1}\}$ of $\mathbb{I}$ can be uniquely assigned the family of *$\mathcal{I}$-centered hat functions* $\tau_0, \ldots, \tau_{n-1} \in C(\mathbb{I};\mathbb{R})$ characterised by:

$$\tau_j \;\text{ is }\; \mathcal{I}\text{-piecewise affine} \qquad \text{and} \qquad \tau_j(t_\nu) \;=\; \delta_{j\nu} \;\text{ for each } \nu \in [n-1]_0 \tag{4.61}$$

for all $j \in [n-1]_0$. A path in $\mathcal{C} \equiv C(\mathbb{I};\mathbb{R}^d)$ will be called *$\mathcal{I}$-piecewise linear* if it lies in

$$\mathcal{C}_\mathcal{I} \;:=\; \{v_0 \cdot \tau_0 + \ldots + v_{n-1} \cdot \tau_{n-1} \;\big|\; v_0, \ldots, v_{n-1} \in \mathbb{R}^d\} \tag{4.62}$$

(the 'vectorial span' of (4.61)). Clearly, the set (4.62) is a closed linear subspace of $(\mathcal{C}, \|\cdot\|_\infty)$, in fact of $(\mathcal{BV}, |\!|\!|\cdot|\!|\!|_{p\text{-var}})$ (cf. below), and each element $\hat{x} = (\hat{x}_t) \in \mathcal{C}_\mathcal{I}$ is of the form

$$\hat{x}_t \;=\; \hat{x}_{t_{j-1}} + \frac{t - t_{j-1}}{t_j - t_{j-1}} \cdot (\hat{x}_{t_j} - \hat{x}_{t_{j-1}}) \quad \text{for} \quad t \in [t_{j-1}, t_j] \qquad (j \in [n-1]).$$

The space $\mathcal{C}_\mathcal{I}$ is the co-domain of two natural operators, namely the linear projection

$$\hat{\pi}_\mathcal{I} \;:\; \mathcal{C} \twoheadrightarrow \mathcal{C}_\mathcal{I}, \quad \hat{\pi}_\mathcal{I}(x) := x_{t_0} \cdot \tau_0 + \ldots + x_{t_{n-1}} \cdot \tau_{n-1} \equiv \hat{x}_\mathcal{I} \tag{4.63}$$

as well as the (continuous wrt. both $\|\cdot\|_\infty$ and $|\!|\!|\cdot|\!|\!|_p$; set $Z := \mathbb{R}^d$) linear injection

$$\hat{\iota}_\mathcal{I} \;:\; Z^{\times n} \hookrightarrow \mathcal{C}_\mathcal{I}, \quad \hat{\iota}_\mathcal{I}(v_0, \ldots, v_{n-1}) := v_0 \cdot \tau_0 + \ldots + v_{n-1} \cdot \tau_{n-1}. \tag{4.64}$$

It is clear that the linear operator $\hat{\pi}_\mathcal{I}$ is bounded on $\mathcal{C}$ with operator norm $\|\hat{\pi}_\mathcal{I}\| = 1$.

Remark 4.7.1. (i) As any two points in $\mathbb{R}^d$ uniquely determine the affine path-segment that joins them, the $\mathcal{I}$-piecewise linear projection $\hat{x}_\mathcal{I}$ of a path x can be seen as the 'unbiased continuous-time approximation' of x given the observations $(x_t \mid t \in \mathcal{I})$.[26]

[26] Likewise, the injection (4.64) can be seen as the 'unbiased $\mathcal{I}$-centered *continuous-time localisation*' of a sequence $(v_1, \ldots, v_n) \in Z^{\times n}$.

(ii) For any $(z_j)_{j\in[n]} \in Z^{\times n}$ and any $\mathbb{I}$-dissection of cardinality $|\mathcal{I}| = n$,

$$\|\hat{\iota}_{\mathcal{I}}(z_1,\ldots,z_n)\|_{1\text{-var}} = \sum_{j=1}^{n-1} |z_{j+1} - z_j| \;\leq\; 2\|(z_1,\ldots,z_n)\|_1.$$

Denote by $\mathcal{BV}_p := \{x \in \mathcal{C} \mid (2.4) \text{ is finite}\}$ the space of all continuous paths of bounded p-variation ($p \geq 1$), and remark that each $\mathcal{BV}_p$ is a Banach space wrt. the norm $|||x|||_{p\text{-var}} := \|x\|_{p\text{-var}} + |x(0)|$ (e.g. [54, Thm. 5.25 (i)]).

Lemma 4.7.2. *For $(\mathcal{I}_n)_{n\in\mathbb{N}}$ a refined sequence of dissections of a compact interval $\mathbb{I}$,*

$$\lim_{n\to\infty} \hat{\pi}_{\mathcal{I}_n} \;=\; \mathrm{id}_{\mathcal{C}} \quad \textit{pointwise on } \; \mathcal{C}(\mathbb{I};\mathbb{R}^d) \tag{4.65}$$

where for each argument the above convergence is understood to take place in $(\mathcal{C}, \|\cdot\|_\infty)$. In addition, the family of operators $(\hat{\pi}_{\mathcal{I}_n} \mid n \in \mathbb{N})$ is equicontinuous, whence in particular the convergence (4.65) *is uniform on compact subsets of $\mathcal{C}$. The family of operators $(\hat{\pi}_{\mathcal{I}_n} \mid n \in \mathbb{N})$ remains equicontinuous if $(\mathcal{C}, \|\cdot\|_\infty)$ is replaced by $(\mathcal{BV}_p, |||\cdot|||_{p\text{-var}})$ for any $p \geq 1$.*

Proof. The pointwise convergence (4.65) is an easy consequence of definition (4.63) and the fact that every element of $\mathcal{C}$ is uniformly continuous on $\mathbb{I}$. The equicontinuity of the family of linear operators $(\hat{\pi}_{\mathcal{I}_n} \mid n \in \mathbb{N})$ is immediate by the fact that this family is uniformly bounded (by 1) in the operator norm. That a pointwise convergent sequence of equicontinuous functions on a metric space (with values in a complete metric space) converges uniformly on compact subsets of its domain is a well-known fact from real analysis (e.g. [137, Exercise 7.16]). As shown in [54, Prop. 5.20], the operator family $(\hat{\pi}_{\mathcal{I}} : \mathcal{BV}_p \to \mathcal{BV}_p \mid n \in \mathbb{N})$ remains uniformly bounded in the operator norm (and hence is equicontinuous) if the latter is defined wrt. the p-variation norm $|||\cdot|||_{p\text{-var}}$ on the Banach space $\mathcal{BV}_p$. ☐

This subsection concludes with a proof that for protocol-indexed discrete time-series (4.18), the ergodicity notions of Remark 4.4.8 (ii) and Definition 4.4.7 coincide.

Note to this end that for $n \leq \hat{n}$, any $\mathcal{I} \equiv \{t_0 < \cdots < t_{n-1}\}$ embeds monotonically into $\mathcal{E}_{\hat{n}}$,

$$t_j \;\mapsto\; \hat{t}_j := q_j/(\hat{n}-1) \in \mathcal{E}_{\hat{n}} \quad \text{for} \quad q_j := \left\lceil \tfrac{t_j - t_0}{t_{n-1}-t_0} \cdot (\hat{n}-1) \right\rceil, \tag{4.66}$$

where $\lceil\cdot\rceil$ is the ceiling function; lifting (4.66) to a map $\varphi_{\mathcal{I}}^{\hat{n}} : [t_0, t_{n-1}] \to [0,1]$ via piecewise-linear extension defines a strictly monotonous continuous injection of intervals. The embedding (4.66) of $\mathcal{I}$ will be denoted $\mathcal{I}_{\mathcal{E}_{\hat{n}}}$ $(= \{\hat{t}_j \mid j \in [n-1]_0\} = \varphi_{\mathcal{I}}^{\hat{n}}(\mathcal{I}))$.

Given $(X, \mathcal{J}_k)$ as in (4.18) with maximal observation length $\hat{n}_k$, define the *equidistant augmentation of* $(X, \mathcal{J}_k)$ as the time-series $\bar{X}^*_{\mathcal{J}_k} := (\hat{X}^*_t)_{t \in \overline{\mathcal{J}}_k}$ given by (for some fixed $q > 1$)

$$\overline{\mathcal{J}}_k := \bigsqcup_{\nu \in \mathbb{N}} \overline{\mathcal{I}}^{(k)}_\nu \quad \text{with} \quad \overline{\mathcal{I}}^{(k)}_\nu := q(\nu - 1) + \mathcal{E}_{\hat{n}_k} \equiv \left\{ \bar{t}^{(k|\nu)}_0 < \cdots < \bar{t}^{(k|\nu)}_{\hat{n}_k - 1} \right\},$$

$$\text{and} \qquad \hat{X}^* := \sum_{\nu=1}^{\infty} \hat{\iota}_{q(\nu-1)+\left[\mathcal{I}^{(k)}_\nu\right]_{\mathcal{E}_{\hat{n}_k}}} (X_s \mid s \in \mathcal{I}^{(k)}_\nu) \cdot \mathbb{1}_{\left[\bar{t}^{(k|\nu)}_0, \bar{t}^{(k|\nu)}_{\hat{n}_k - 1}\right]}$$

(where $\hat{\iota}_{q(\nu-1)+\left[\mathcal{I}^{(k)}_\nu\right]_{\mathcal{E}_{\hat{n}_k}}} (X_s \mid s \in \mathcal{I}^{(k)}_\nu)$ is the piecewise-linear interpolation (4.64) of the observation $(X_t)_{t \in \mathcal{I}^{(k)}_\nu}$ along the $\overline{\mathcal{I}}^{(k)}_\nu$-embedded (via (4.66)) equidistant dissection $q(\nu - 1) + [\mathcal{I}^{(k)}_\nu]_{\mathcal{E}_{\hat{n}_k}}$ of $[\bar{t}^{(k|\nu)}_0, \bar{t}^{(k|\nu)}_{\hat{n}_k - 1}]$). Then the following holds.

Lemma 4.7.3. *A time-series* $(X_t)_{t \in \mathcal{J}_k}$ *for* $\mathcal{J}_k$ *as in* (4.18)*, is [weakly]* m^{th}*-order signature ergodic in the sense of* (4.35) *iff its equidistant augmentation* $\bar{X}^*_{\mathcal{J}_k}$ *is [weakly]* m^{th}*-order signature ergodic to length* $\hat{n}_k$ *in the sense of Definition 4.4.7.*

Proof. This follows from Lemma 2.3.5 (iii). Indeed: Fix any ν. Denoting $\bar{\mathcal{J}} := q(\nu - 1) + [\mathcal{I}^{(k)}_\nu]_{\mathcal{E}_{\hat{n}_k}}$ $(= \varphi(\mathcal{I}^{(k)}_\nu)$ for $\varphi \equiv q(\nu - 1) + \varphi^{\hat{n}_k}_{\mathcal{I}^{(k)}_\nu}$, with $\varphi^{\hat{n}_k}_{\mathcal{I}^{(k)}_\nu}$ defined as in (4.66) above), note that $\bar{\mathcal{J}} \subseteq \overline{\mathcal{I}}^{(k)}_\nu =: \bar{\mathcal{I}}$ and hence $\hat{\pi}_{\bar{\mathcal{J}}} = \hat{\pi}_{\bar{\mathcal{I}}} \circ \hat{\pi}_{\bar{\mathcal{J}}}$ (directly by (4.63)). Consequently,

$$\hat{X}^*_{(\nu)} := \hat{X}^*\Big|_{\left[\bar{t}^{(k|\nu)}_0, \bar{t}^{(k|\nu)}_{\hat{n}_k - 1}\right]} = \hat{\iota}_{\bar{\mathcal{J}}}(X_s \mid s \in \mathcal{I}^{(k)}_\nu) = \hat{\pi}_{\bar{\mathcal{J}}}(X_{\varphi^{-1}}) = \hat{\pi}_{\bar{\mathcal{I}}}(\hat{X}^*_{(\nu)}) \tag{4.67}$$

and therefore, for ϕ and $\tilde{\phi}$ as in (4.33) and (4.35) respectively, and $Y := \bar{X}^*_{\mathcal{J}_k}$ and $\hat{n} := \hat{n}_k$,

$$\phi(Y_{(\hat{n}(\nu-1), \hat{n}\nu]}) = \mathfrak{sig}_{[m]}(\hat{\pi}_{\bar{\mathcal{I}}}(\hat{X}^*_{(\nu)})) = \tilde{\phi}(\hat{X}^*_{(\nu)}). \tag{4.68}$$

Now since $\hat{X}^*_{(\nu)} \overset{(4.67)}{=} \hat{\pi}_{\bar{\mathcal{J}}}(X_{\varphi^{-1}}) \overset{(4.63)}{=} \sum_{t \in \varphi(\mathcal{I})} X_{\varphi^{-1}(t)} \cdot \tau^{[\bar{\mathcal{J}}]}_t = \sum_{s \in \mathcal{I}} X_{\varphi^{-1}(\varphi(s))} \cdot \tau^{[\bar{\mathcal{J}}]}_{\varphi(s)} = \sum_{s \in \mathcal{I}} X_s \cdot (\tau^{[\mathcal{I}]}_s \circ \varphi^{-1}) = \hat{\pi}_{\mathcal{I}}(X) \circ \varphi^{-1} = \hat{\pi}_{\mathcal{I}}(\hat{X}_{\mathcal{I}}) \circ \tilde{\varphi}$ for $\mathcal{I} := \mathcal{I}^{(k)}_\nu$ and $\tilde{\varphi} := \varphi^{-1}$, we have

$$\tilde{\phi}(\hat{X}^*_{(\nu)}) = \tilde{\phi}(\hat{\pi}_{\mathcal{I}}(\hat{X}_{\mathcal{I}})) = \tilde{\phi}(\hat{X}_{\mathcal{I}^{(k)}_\nu}) \tag{4.69}$$

where the first of these equalities is due to Lemma 2.3.5 (iii). As the above the choice of ν was arbitrary, we by combination of (4.68) and (4.69) obtain that the identities (4.33) and (4.35) are equivalent. This concludes the proof. □

4.7.4 Proof of Lemma 4.4.5

Proof of Lemma 4.4.5. Recalling that $\mathrm{tr}(X) \subseteq D_X$ almost surely (Lemma 3.2.5 (ii)), notice

$$\lim_{n \to \infty} \sup_{\theta \in \Theta} \left\| \hat{X}^\theta_{\mathcal{I}_n} - \theta(X) \right\|_{\tilde{p}\text{-var}} = 0 \qquad \text{almost surely} \tag{4.70}$$

for any $\tilde{p} > p$ with $p \geq 1$ as in (4.26). Indeed: Denoting by $x := (X_t(\omega))_{t\in\mathbb{I}} \subseteq D_X$ (inclusion a.s.) a given realisation of X, we remark first that (cf. Section 4.7.3 for notation)

$$\Theta_x \;:=\; \{x^\theta \equiv (\theta(x_t))_{t\in\mathbb{I}} \mid \theta \in \Theta\} \quad \text{is a compact subset of} \quad (\mathcal{BV}_{\tilde{p}}, |||\cdot|||_{\tilde{p}\text{-var}}) \tag{4.71}$$

for any $\tilde{p} > p$. To see that (4.71) holds, observe that [54, Lemma 5.27 (i)] (together with (4.26)) implies that, for any $\tilde{p} > p$, each of the functions

$$\alpha_n,\, \alpha \;:\; \Theta \to \mathcal{BV}_{\tilde{p}}, \quad \alpha_n(\theta) := \hat{\pi}_{\mathcal{I}_n}(\theta(x)) \;\text{ and }\; \alpha(\theta) := \theta(x) \qquad (n \in \mathbb{N}) \tag{4.72}$$

are continuous. Hence $\Theta_x = \alpha(\Theta)$ is compact (as continuous image of a compact set).

In addition, [54, Lemma 5.27 (i)] (by virtue of Lemma 4.7.2 (4.65) and [54, Proposition 5.20 (5.13)]) implies that $\lim_{n\to\infty} \alpha_n = \alpha$ pointwise on Θ. Hence by (4.71) and the last assertion of Lemma 4.7.2 (which implies that $(\hat{\pi}_{\mathcal{I}_n})$ converges uniformly on Θ_x; see the proof of Lemma 4.7.2 for details), we obtain that $\lim_{n\to\infty} \alpha_n = \alpha$ uniformly on Θ, which in turn yields (4.70) by the fact that $\hat{X}^\theta_{\mathcal{I}_n} = \hat{\pi}_{\mathcal{I}_n}(\theta(X))$ for each $\theta \in \Theta$.

Given (4.70) for any fixed $\tilde{p} \in (p, 2)$, the $\tilde{p}$-variation continuity of $\mathfrak{sig}$ (Lemma 2.3.5 (ii)) together with the equicontinuity of $(\hat{\pi}_{\mathcal{I}_n} : \mathcal{BV}_{\tilde{p}} \to \mathcal{BV}_{\tilde{p}} \mid n \in \mathbb{N})$ (Lemma 4.7.2) yields that

$$\lim_{n\to\infty} \sup_{\theta\in\Theta} \left\| \mathfrak{sig}_m(\hat{X}^\theta_{\mathcal{I}_n}) - \mathfrak{sig}_m(\theta(X)) \right\|_m \;=\; 0 \quad \text{almost surely} \qquad (m \in \mathbb{N}).$$

Indeed, the above holds path-wise, with probability one, by Lemma 4.7.4 (applied to Θ as above, $B = \mathcal{BV}_{\tilde{p}}$, $V = V_{[m]}$, $\Psi = \mathfrak{sig}_m$, α and α_n as in (4.72) and $\tau = \alpha$).

Thus for $\mathfrak{S}_{m|n}(\theta) := \mathbb{E}[\mathfrak{sig}_m(\hat{X}^\theta_{\mathcal{I}_n})]$ and $\mathfrak{S}_m(\theta) := \mathbb{E}[\mathfrak{sig}_m(\theta(X))]$ we have that

$$\lim_{n\to\infty} \mathfrak{S}_{m|n}(\theta) \;=\; \mathfrak{S}_m(\theta) \qquad \text{uniformly on } \Theta \tag{4.73}$$

due to [58, Theorem 22 (p. 241)] (note that the hypothesis in loc.cit. of $(\mathfrak{S}_{m|n})_n$ to be "absolutely continuous uniformly" is met in light of [58, Thm. 11 (p. 192)] and assumption (4.27)).

Finally, the fact that $\log_{[m]} \equiv \pi_{[m]} \circ \log$ is continuous (Lemma 2.3.5 (v)) together with the uniform convergence (4.73) of $\mathfrak{S}_{[m]|n} := \sum_{\nu=0}^m \mathfrak{S}_{\nu|n}$ towards $\mathfrak{S}_{[m]} := \sum_{\nu=0}^m \mathfrak{S}_\nu$, yields that

$$\kappa_n^{[m]} := \log_{[m]} \circ \mathfrak{S}_{[m]|n} \;\overset{n\to\infty}{\longrightarrow}\; \log_{[m]} \circ \mathfrak{S}_{[m]} =: \kappa^{[m]} \qquad \text{uniformly on } \Theta.$$

In particular, $\kappa_{\boldsymbol{q}}(\hat{X}^\theta_{\mathcal{I}_n}) = \langle \kappa_n^{[m]}(\theta), \boldsymbol{q}\rangle \to \langle \kappa^{[m]}(\theta), \boldsymbol{q}\rangle = \kappa_{\boldsymbol{q}}(\theta(X))$ uniformly on Θ for each $\boldsymbol{q} \in V_{[m]}$, which by definitions (4.30) (of $\widehat{Q}_m$) and (4.25) (of Q_m) yields (4.31) as desired. □

Lemma 4.7.4. *Let Θ be a compact metric space, B and V be Banach spaces, $\Psi : B \to V$ be a continuous map, and α, α_n, $\tau : \Theta \to B$, $n \in \mathbb{N}$, be continuous functions such that*

$$\alpha_n = p_n \circ \tau, \quad n \in \mathbb{N}, \quad \textit{with} \quad (p_n : \tau(\Theta) \to B \mid n \in \mathbb{N}) \;\; \textit{equicontinuous}$$

and $\lim_{n\to\infty} \alpha_n = \alpha$ *pointwise on* Θ*. Then* $\lim_{n\to\infty} \Psi \circ \alpha_n = \Psi \circ \alpha$ *uniformly on* Θ*.*

Proof of Lemma 4.7.4. Let $(\theta_n)_{n\in\mathbb{N}}$ be any convergent sequence in Θ, say $\theta_n \to \theta$ for some $\theta \in \Theta$. Then

$$\lim_{n\to\infty} \Phi_n(\theta_n) \;=\; \Phi(\theta) \qquad \text{for} \qquad \Phi_n := \Psi \circ \alpha_n \;\text{ and }\; \Phi := \Psi \circ \alpha, \tag{4.74}$$

since $\|\Phi(\theta) - \Phi_n(\theta_n)\|_V \le \|\Phi(\theta) - \Phi_n(\theta)\|_V + \|\Phi_n(\theta) - \Phi_n(\theta_n)\|_V$ with $\lim_{n\to\infty} \|\Phi(\theta) - \Phi_n(\theta)\|_V = 0$ and $\lim_{n\to\infty} \|\Phi_n(\theta) - \Phi_n(\theta_n)\|_V = 0$. For the latter convergence, take any $\varepsilon > 0$ and let $\delta_1 > 0$ be such that $\sup_{b\in B_{\delta_1}(\alpha(\theta))} \|\Psi(\alpha(\theta)) - \Psi(b)\|_V \le \varepsilon$, and $\delta_2 > 0$ be such that $\rho_n := \sup_{b\in B_{\delta_2}(\tau(\theta))\cap\tau(\Theta)} \|p_n(\tau(\theta)) - p_n(b)\|_B \le \delta_1$ for all $n \in \mathbb{N}$. Taking $n_0 \ge 1$ such that $\sup_{n\ge n_0} \|\tau(\theta) - \tau(\theta_n)\|_B \le \delta_2$ then implies that

$$\sup_{n\ge n_0} \|\alpha_n(\theta) - \alpha_n(\theta_n)\|_B = \sup_{n\ge n_0} \|p_n(\tau(\theta)) - p_n(\tau(\theta_n))\|_B \;\le\; \sup_{n\ge n_0} \rho_n \le \delta_1$$

and therefore $\sup_{n\ge n_0} \|\Phi_n(\theta) - \Phi_n(\theta_n)\|_V \le \varepsilon$, as required.

Conclude by observing that (4.74) implies $\Phi_n \to \Phi$ uniformly on Θ, as desired.

Indeed, assume otherwise that $\Phi_n \nrightarrow \Phi$ uniformly, i.e. that there is $\tilde{\varepsilon} > 0$ such that

$$\forall\, k \in \mathbb{N} \;:\; \exists\, n_k \in \mathbb{N}_{\ge n} \;\text{ with }\; \sup_{\theta\in\Theta} \|\Phi(\theta) - \Phi_{n_k}(\theta)\|_V > \tilde{\varepsilon}. \tag{4.75}$$

Then (4.75) informs the choice of a subsequence $(\theta_{n_k})_k \subseteq \Theta$, with $(n_k)_k \subseteq \mathbb{N}$ increasing, s.t.

$$\|\Phi(\theta_{n_k}) - \Phi_{n_k}(\theta_{n_k})\|_V \;>\; \tilde{\varepsilon} \qquad \text{for each} \quad k \in \mathbb{N}. \tag{4.76}$$

As Θ is compact, we may assume, by passing to a further subsequence if necessary, that this subsequence converges, say to $\tilde{\theta} \in \Theta$. The continuity of Φ then implies $\lim_{k\to\infty} \Phi(\theta_{n_k}) = \Phi(\tilde{\theta})$, while property (4.74) combined with a doubling argument (as in [131, Sect. 3.5*: remark on p. 98]) yields $\lim_{k\to\infty} \Phi_{n_k}(\theta_{n_k}) = \Phi(\tilde{\theta})$. Hence $\lim_{k\to\infty} \|\Phi(\theta_{n_k}) - \Phi_{n_k}(\theta_{n_k})\|_V = 0$, in contradiction to (4.76). $\square$

4.7.5 Proof of Proposition 4.4.9

Proof of Proposition 4.4.9. For $\tilde{m} \ge 1$ and $\theta \in \Theta$ and $w \in V_{\tilde{m}}$ all arbitrary but fixed, let $\phi = \phi_{\tilde{m}}$ be as in (4.33) and set $\xi := \langle \phi\circ\theta^{\times n},\, w\rangle$. Set further $\hat{\xi}_T(z) := \frac{1}{T}\sum_{j=1}^{T} \xi(z_{n(j-1)+1}, \ldots, z_{nj})$ for any given sequence $z = (z_\nu)_{\nu\in\mathbb{N}}$ in $\mathbb{R}^d$. The lemma then asserts that, under the given integrability and ergodicity conditions,

$$\mathbb{E}[\xi(X_1, \ldots, X_n)] \;=\; \lim_{T\to\infty} \hat{\xi}_T(X_*) \quad \text{a.s.} \quad [\text{resp.}^{27} \text{ in probab.}]. \tag{4.77}$$

To see that (4.77) holds, note first that for $\hat{X}_1 := \hat{\iota}_{\mathcal{E}_n}(X_1, \ldots, X_n)$ (cf. Def. 4.4.7 and (4.64)),

$$\xi(X_{[n]}) \;=\; \varphi(\hat{X}_1) \qquad \text{for} \qquad \varphi(x) := \left\langle \mathfrak{sig}_{[\tilde{m}]}(\hat{\pi}_{\mathcal{E}_n}(\tilde{\theta}(x))),\, w\right\rangle \tag{4.78}$$

[27] For simplicity of exposition, we present the case of almost sure convergence first and give the changes necessary for the case of convergence in probability at the end of this proof.

where $\tilde{\theta}$ is any fixed continuous extension of θ to $\widehat{D} := \operatorname{conv}(D_{X_*})$, the convex hull of D_{X_*}. (Recall that such a $\tilde{\theta}$ exists by Tietze's extension theorem.) Since the function $\varphi : \widehat{\mathcal{BV}}_n \to \mathbb{R}$ defined by (4.78) on the compact[28] subset

$$\widehat{\mathcal{BV}}_n := \left\{ x \in \mathcal{C}_{\mathcal{E}_n} \;\middle|\; x_t \in \widehat{D} \;\text{ for each }\; t \in \mathcal{E}_n \right\} \subset \mathcal{BV} \qquad \text{(cf. (4.62))}$$

is continuous (by [54, Prop. 5.20] and Lemma 2.3.5 (ii)), the universality property of the signature (e.g. Lemma 2.3.5 (iv)) implies that there is a sequence $(\boldsymbol{\ell}_j)_{j\in\mathbb{N}}$ in V° such that

$$\varphi \;=\; \lim_{j\to\infty} \langle \mathfrak{sig}(\cdot), \boldsymbol{\ell}_j \rangle \qquad \text{in} \;\; (C(\widehat{\mathcal{BV}}_n), \|\cdot\|_\infty). \tag{4.79}$$

For $\big(\hat{\mathbb{E}}_T^{(m)}(X_*)\big)_{T\in\mathbb{N}} := \big(\frac{1}{T}\sum_{\nu=1}^T \mathfrak{sig}_{[m]}(\hat{X}_\nu)\big)_{T\in\mathbb{N}}$ with $\hat{X}_\nu := \iota_{\mathcal{E}_n}(X_{n(\nu-1)+1}, \ldots, X_{n\nu})$, our assumption on X_* gives that, for each $m \in \mathbb{N}$,

$$\mathbb{E}[\mathfrak{sig}_{[m]}(\hat{X}_1)] = \lim_{T\to\infty} \hat{\mathbb{E}}_T^{(m)}(X_*) \quad \text{a.s.} \quad \text{in} \;\; \operatorname{conv}(\mathfrak{sig}_{[m]}(\widehat{\mathcal{BV}}_n)). \tag{4.80}$$

Hence upon combining (4.78) and (4.79), and using that dominated convergence applies as both sides of (4.79) are bounded (cf. Lemma 2.3.5 (ii)), we find that with probability one,

$$\begin{aligned} \mathbb{E}[\xi(X_1, \ldots, X_n)] &= \lim_{j\to\infty} \langle \mathbb{E}[\mathfrak{sig}(\hat{X}_1)], \boldsymbol{\ell}_j \rangle = \lim_{j\to\infty} \lim_{T\to\infty} \langle \hat{\mathbb{E}}_T^{(d_{\boldsymbol{\ell}_j})}(X_*), \boldsymbol{\ell}_j \rangle \\ &= \lim_{T\to\infty} \lim_{j\to\infty} \frac{1}{T} \sum_{\nu=1}^T \langle \mathfrak{sig}_{[d_{\boldsymbol{\ell}_j}]}(\hat{X}_\nu), \boldsymbol{\ell}_j \rangle \\ &= \lim_{T\to\infty} \frac{1}{T} \sum_{\nu=1}^T \lim_{j\to\infty} \langle \mathfrak{sig}(\hat{X}_\nu), \boldsymbol{\ell}_j \rangle \overset{(4.79)}{=} \lim_{T\to\infty} \frac{1}{T} \sum_{\nu=1}^T \xi(\hat{X}_\nu), \end{aligned} \tag{4.81}$$

where we denoted $d_{\boldsymbol{\ell}}$ for the degree of the index-polynomial $\boldsymbol{\ell} \in V^\circ$. Notice that the interchange of limits in the second line of (4.81) is permissible as the convergence in (4.79) is uniform (see, e.g., [136, Theorem 7.11]). This shows the almost-sure case of (4.77).

To prove that (4.77) holds in probability if (4.80) holds in probability for each $m \in \mathbb{N}$ (which is true by assumption if X_* is weakly signature ergodic), we resort to a subsequence argument, recalling that (as the topology of weak convergence is metrizable) a sequence converges in probability iff each of its subsequences admits yet another subsequence that converges almost surely. To this end, abbreviate $\mu_{m,T} := \hat{\mathbb{E}}_T^{(m)}(X_*)$ and assume that

$$\mu_m := \mathbb{E}[\mathfrak{sig}_{[m]}(\hat{X}_1)] \;=\; \lim_{T\to\infty} \mu_{m,T} \quad \text{in probability} \qquad \text{for each} \;\; m \in \mathbb{N}. \tag{4.82}$$

Then for any fixed subsequence $(T_k)_{k\in\mathbb{N}} \subset \mathbb{N}$, there is a subsequence $T_k^{(1)} < T_{k+1}^{(1)}$ of (T_k) such that $\lim_{k\to\infty} \mu_{1,T_k^{(1)}} = \mu_1$ almost surely. But since, by (4.82), $\lim_{k\to\infty} \mu_{2,T_k^{(1)}} = \mu_2$ in

[28] By [104, Prop. 1.7] and the facts that: (a) the convex hull operator on $\mathbb{R}^d$ preserves compactness, and (b) the Cartesian product of compact sets is compact (noting that $\widehat{\mathcal{BV}}_n \cong \widehat{D}^{\times n}$).

probability, there will be a subsequence $T_k^{(2)} < T_{k+1}^{(2)}$ of $(T_k^{(1)})$ such that $\lim_{k\to\infty} \mu_{2,T_k^{(2)}} = \mu_2$ almost surely (thus $\lim_{k\to\infty} \mu_{1,T_k^{(2)}} = \mu_1$ a.s. in particular). Iterating this procedure, Cantor's diagonal trick (e.g. [129, Proof of Theorem I.24]) thus allows for the choice of a subsequence $T_k^{(\infty)} < T_{k+1}^{(\infty)}$ of (T_k) such that $\lim_{k\to\infty} \mu_{m,T_k^{(\infty)}} = \mu_m$ almost surely for each $m \in \mathbb{N}$.

Repeating the above calculation (4.81) then shows that the subsequence $(\hat{\xi}_{T_k^{(\infty)}}(X_*))_{k\in\mathbb{N}}$ of $(\hat{\xi}_{T_k}(X_*))_{k\in\mathbb{N}}$ converges almost surely to $\mathbb{E}[\xi(X_{[n]})]$, as desired. □

4.7.6 Proof of Lemma 4.4.10

Proof of Lemma 4.4.10. Let $Z := \mathbb{R}^d$ and $\mathcal{E}_n$ be as in Def. 4.4.7, and for every $\theta \in \Theta$ denote

$$\xi_\theta \;:=\; \mathfrak{sig}_{[m]} \circ \hat{\iota}_{\mathcal{E}_n} \circ \theta^{\times n} \;:\; Z^{\times n} \longrightarrow V_{[m]} \cap V_{(1)}. \tag{4.83}$$

The parametrisation-invariance of $\mathfrak{sig}$ (Lemma 2.3.5 (iii)) gives that

$$\hat{\mathfrak{S}}_T^{m|n}(\theta) = \frac{1}{T}\sum_{j=1}^{T} \xi_\theta(\bar{X}_j) \;=:\; \hat{\mathbb{E}}_T[\xi_\theta(X_*)] \quad \text{for} \quad \bar{X}_j := (X_{n(j-1)+1}, \ldots, X_{nj}),$$

and the continuity of $\log_{[m]}$ (Lemma 2.3.5 (v)) yields that (4.37) follows from the convergence

$$\lim_{T\to\infty} \sup_{\theta\in\Theta} \left\|\mathbb{E}[\xi_\theta(\bar{X}_1)] - \hat{\mathbb{E}}_T[\xi_\theta(X_*)]\right\|_{[m]} = 0 \qquad \text{a.s.} \quad \text{[in probab.]} \tag{4.84}$$

for the norm $\|\cdot\|_{[m]} := \sum_{\nu=1}^m \|\cdot\|_\nu$, followed by an application of the continuous mapping theorem (e.g. [152, Theorem 2.3]). As (4.84) is equivalent to the coordinatewise convergences

$$\lim_{T\to\infty} \sup_{w\in[d]_k^*} \sup_{\theta\in\Theta} \left|\mathbb{E}[\langle \xi_\theta(\bar{X}_1), w\rangle] - \langle \hat{\mathbb{E}}_T[\xi_\theta(X_*)], w\rangle\right| = 0 \qquad \text{for} \quad k \in [m] \tag{4.85}$$

almost surely (resp. in prob.), we can see that (4.85) holds by fixing any $w \in [d]_k^*$ and showing

$$\lim_{T\to\infty} \sup_{\theta\in\Theta} \left|\mathbb{E}[\tilde{\xi}_\theta(\bar{X}_1)] - \hat{\mathbb{E}}_T[\tilde{\xi}_\theta(X_*)]\right| = 0 \quad \text{[a.s./in prob.]} \quad \text{for} \quad \tilde{\xi}_\theta := \langle \xi_\theta, w\rangle \tag{4.86}$$

and $\hat{\mathbb{E}}_T[\tilde{\xi}_\theta(X_*)] := \langle \hat{\mathbb{E}}_T[\xi_\theta(X_*)], w\rangle = T^{-1}\sum_{j=1}^T \tilde{\xi}_\theta(\bar{X}_j)$. To this end, note that the function

$$\tilde{\xi} \;:\; Z^{\times n} \times \Theta \longrightarrow \mathbb{R}, \quad (z, \theta) \;\mapsto\; \tilde{\xi}_\theta(z),$$

is continuous in θ for every $z \in Z^{\times n}$, as is seen directly from (4.83) (recalling the continuity of $\hat{\iota}_{\mathcal{E}_n} : Z^{\times n} \to \mathcal{BV}$ (Rem. 4.7.1 (ii)) and Lemma 2.3.5 (ii)). Also, by assumption, Θ is compact with $\mathbb{E}[\sup_{\theta\in\Theta} |\tilde{\xi}_\theta(\bar{X}_1)|] < \infty$ and $\hat{\mathbb{E}}_T[\xi_\theta(X_*)] \to \mathbb{E}[\xi_\theta(\bar{X}_1)]$ pointwise, which altogether implies that $\mathcal{F} := \{\tilde{\xi}(\cdot, \theta) \mid \theta \in \Theta\}$ is Glivenko-Cantelli via [155, Lem. 6.1, Thm. 6.1], i.e. that

$$\lim_{T\to\infty} \sup_{\varphi\in\mathcal{F}} \left| \mathbb{E}[\varphi(\bar{X}_1)] - \frac{1}{T}\sum_{j=1}^{T} \varphi(\bar{X}_j) \right| \; = \; 0, \tag{4.87}$$

where the mode of the convergence in (4.87) (almost surely or in probability) coincides with the mode of the pointwise convergence $\hat{\mathbb{E}}_T[\xi_\theta(X_*)] \to \mathbb{E}[\xi_\theta(\bar{X}_1)]$ on Θ (cf. [155, (Proof of) Theorem 6.1]). As (4.87) is identical to (4.86), the proof is finished. □

4.7.7 Sufficent Conditions for Signature-Ergodicity

Let $X_* \equiv (X_j)_{j\in\mathbb{N}}$ be a sequence of $\mathbb{R}^d$-valued random variables.

Definition 4.7.5. The sequence X_* is called *α-mixing* if for the sub-σ-algebras $\mathcal{X}_k^\ell := \sigma\big(X_\nu \mid k \leq \nu \leq \ell\big)$ it holds that $\lim_{\nu\to\infty} \alpha_\nu(X_*) = 0$ for the sequence

$$\alpha_\nu(X_*) \; := \sup_{A\in\mathscr{X}_1^j,\, B\in\mathscr{X}_{j+\nu}^\infty,\, j\in\mathbb{N}} \big|\mathbb{P}(A\cap B) - \mathbb{P}(A)\mathbb{P}(B)\big|\,,$$

and X_* is called *ϕ-mixing* if it holds that $\lim_{\nu\to\infty} \phi_\nu(X_*) = 0$ for the sequence

$$\phi_\nu(X_*) \; := \sup_{A\in\mathscr{X}_1^j,\, B\in\mathscr{X}_{j+\nu}^\infty,\, \mathbb{P}(A)>0,\, j\in\mathbb{N}} \big|\mathbb{P}(B\,|\,A) - \mathbb{P}(B)\big|.$$

Note that ϕ-mixing implies α-mixing, and see e.g. [21] for further information.

Definition 4.7.6. The sequence X_* will be said to have *n-seasonal increments*, $n \in \mathbb{N}$, if the sequence $\Delta(X_*) := (X_{j+1} - X_j)_{j\in\mathbb{N}}$ of its increments is suff. integrable and such that

$$\Delta(X_*)_{[n]} \;\overset{\mathrm{d}}{=}\; \Delta(X_*)_{(n(j-1),\, nj]} \qquad \text{for each } j \in \mathbb{N}.$$

We further say that a time series $(X_j)_{j\in\mathbb{N}}$ has *(m,n)-stationary sigmoments* if the $[m]^{\text{th}}$-signature moments of the batches $(X_1, \ldots, X_n), (X_{n+1}, \ldots, X_{2n}), \ldots$ exist and are equal, i.e. if for $\phi = \phi_m$ as in (4.33) we have: $\mathbb{E}[\phi_m(X_{[n]})] = \mathbb{E}[\phi_m(X_{(n(j-1):nj]})]$ for each $j \in \mathbb{N}$.

Lemma 4.7.7. *For $X_* \equiv (X_j)_{j\in\mathbb{N}}$ uniformly integrable and $n \in \mathbb{N}$, the following holds.*

(i) *If X_* is α-mixing and has (m,n)-stationary sigmoments ($m \in \mathbb{N}$), then X_* is m^{th}-order weakly signature-ergodic to length n;*

(ii) *if X_* is ϕ-mixing with $\sum_{\nu=1}^\infty \phi_{1+(\nu-1)n}^{1/2}(X_*)\frac{\log\nu}{\nu} < \infty$ and has n-seasonal increments, then X_* is signature-ergodic to length n.*

The assertions (i) and (ii) persist if X_* is replaced by $\theta(X_*) = (\theta(X_j))_{j\in\mathbb{N}}$ for any measurable $\theta : D_{X_*} \to \mathbb{R}^d$.

Proof. Starting from definition (2.14), a direct calculation yields that for any $\ell_1 < \ell_2$,

$$\mathfrak{sig}_m(\hat{\iota}_{\mathcal{E}}(X_{\ell_1}, \dots, X_{\ell_2})) \;=\; \sum_{(i_1,\dots,i_m)\in(\ell_1:\,\ell_2]^{\times m}} c_{i_1\cdots i_m} \cdot \Delta_{i_1} \otimes \cdots \otimes \Delta_{i_m} \tag{4.88}$$

for certain $c_{i_1\cdots i_m} \in \mathbb{R}$ and increments $\Delta_j := X_j - X_{j-1}$, where $\mathcal{E} \equiv \mathcal{E}_{\ell_1,\ell_2}$ is the equidistant (or any other) $\mathbb{I}$-dissection of cardinality $\ell_2 - \ell_1 + 1$. Let now $n \in \mathbb{N}$ be fixed. If we introduce the shift-map $\vartheta(i) := i + n$ (with $\vartheta^0 := \mathrm{id}$ and $\vartheta^j := \vartheta \circ \vartheta^{j-1}$) for convenience and denote

$$Y_j \;:=\; \mathfrak{sig}_m(\hat{\iota}_{\mathcal{E}}(\theta \cdot X_{n(j-1)+1}, \dots, \theta \cdot X_{nj})) \qquad (j \in \mathbb{N})$$

for brevity, then the above shows that each Y_j is a measurable function of the arguments $X_{\vartheta^{j-1}(1)}, \dots, X_{\vartheta^{j-1}(n)}$. This in turn implies the inclusion of σ-algebras

$$\mathscr{Y}_p^q := \sigma(Y_p, \dots, Y_q) \;\;\subseteq\;\; \sigma(X_{\vartheta^{p-1}(1)}, \dots, X_{\vartheta^{p-1}(n)}, \dots, X_{\vartheta^{q-1}(1)}, \dots, X_{\vartheta^{q-1}(n)})$$

for any $p \leq q$, whence in particular $\mathscr{Y}_1^j \subseteq \mathscr{X}_1^{\vartheta^{j-1}(n)}$ and $\mathscr{Y}_{j+\nu}^\infty \subseteq \mathscr{X}_{\vartheta^{j+\nu-1}(1)}^\infty$ for all $\nu \in \mathbb{N}$. Since $\vartheta^{j-1}(n) = jn$ and $\vartheta^{j+\nu-1}(1) = jn + \vartheta^{\nu-1}(1)$, we can use Definition 4.7.5 to for $Y_* := (Y_j)_{j\in\mathbb{N}}$ and $\gamma \in \{\alpha, \phi\}$ conclude that

$$\gamma_\nu(Y_*) \;\leq\; \gamma_{\vartheta^{\nu-1}(1)}(X_*) = \gamma_{1+(\nu-1)n}(X_*) \quad \text{ for each } \nu \in \mathbb{N},$$

which shows that if X_* is α-mixing (ϕ-mixing) then so is Y_*.

The proof of statement (i) is finished by a coordinatewise application of the weak law of large numbers for non-stationary α-mixing time series given in [153, Theorem 7.15].

As to (ii), we note similarly that if X_* has n-seasonal increments and is ϕ-mixing at the assumed rate, then Y_* is stationary (by (4.88)) and ϕ-mixing with

$$\sum_{\nu=1}^\infty \phi_\nu^{1/2}(Y_*)\frac{\log \nu}{\nu} \;\leq\; \sum_{\nu=1}^\infty \phi_{1+(\nu-1)n}^{1/2}(X_*)\frac{\log \nu}{\nu} \;<\; \infty\,,$$

whence assertion (ii) follows from a coordinatewise application of [97, Corollary 1].

This proof of the statements (i) and (ii) goes through without changes if the sequence $(X_j)_{j\in\mathbb{N}}$ is replaced by $(\theta \cdot X_j)_{j\in\mathbb{N}}$ for any (Borel-)measurable map $\theta : D_{X_*} \to \mathbb{R}^d$. □

4.7.8 Complementary Proofs for Theorem 4.4.13

Throughout this subsection, the setting and notation from the proof of Theorem 4.4.13 (pp. 102) applies.

Remark 4.7.8 (The Compact-Open Topology on Θ is Metrizable)**.** Since D_X is a closed subset of $\mathbb{R}^d$, there are $\{K_\nu\} \subseteq D_X$ compact with $K_\nu \subseteq K_{\nu+1}$ and $D_X = \bigcup_{\nu\in\mathbb{N}_0} K_\nu$, and the topology of compact convergence on $C(D_X; \mathbb{R}^d)$ coincides with the compact-open topology on $C(D_X;\mathbb{R}^d)$, e.g. [113, Theorem 46.8]. Defining $\|\theta\|_K := \sup_{u\in K}|\theta(u)|$, this topology is induced by the metric (see, e.g., [39, Proposition VII.1.6])

$$\tilde{d}(\theta,\tilde{\theta}) := \sum_{\nu=0}^{\infty} 2^{-\nu}\rho_\nu(\theta,\tilde{\theta}) \quad \text{with} \quad \rho_\nu(\theta,\tilde{\theta}) := \frac{\|\theta-\tilde{\theta}\|_{K_\nu}}{1+\|\theta-\tilde{\theta}\|_{K_\nu}}\,; \tag{4.89}$$

we choose $K_\nu := \overline{B_{r_\nu}(0)}$ for any $r_\nu \uparrow \infty$ monotonously with $r_0 := 0$ for convenience.

Note that the metrics (on $C(K_\nu;\mathbb{R}^d)$) ρ_ν and $d_\nu(\theta,\tilde{\theta}) := \|\theta-\tilde{\theta}\|_{K_\nu}$ are equivalent for all $\nu\in\mathbb{N}_0$. Specifically, for each $\nu\in\mathbb{N}_0$ we have $d_\nu(\theta,\tilde{\theta}) \le d_{\nu+1}(\theta,\tilde{\theta})$ for any $\theta,\tilde{\theta}\in\Theta$, and

$$d_\nu(\theta,\tilde{\theta}) \le 2\rho_\nu(\theta,\tilde{\theta}) \qquad \text{if} \quad \rho_\nu(\theta,\tilde{\theta}) \le \tfrac{1}{2}. \tag{4.90}$$

For $\eta = d_\nu, \rho_\nu, \tilde{d}$, denote $B_r^\eta(\theta_*) := \{\theta\in C(D_X;\mathbb{R}^d) \mid \eta(\theta,\theta_*) < r\}$ and $\tilde{B}_r(\theta_*) := B_r^{\tilde{d}}(\theta_*)$.

Below are the proofs of Theorem 4.4.13 for the cases $(X,\mathcal{J})$ ergodic resp. weakly ergodic.

4.7.8.1 Proof of Theorem 4.4.13 for Ergodic Observations

Let $(\tilde{X},\mathcal{J})$ be an ergodic observation such that D_X is not necessarily compact. In this case, Theorem 4.4.13 asserts that each $\varepsilon > 0$ comes with a $\mathbb{P}$-full set $\tilde{\Omega}_\varepsilon \in \mathscr{F}$ such that for each $\omega\in\tilde{\Omega}_\varepsilon$ the following holds:

$$\begin{gathered}\exists\, m_0 \equiv m_0(\omega) \ge 2 \;:\; \forall\, m \ge m_0 \;\text{ there is }\; k_0 \equiv k_0(m)\in\mathbb{N} \;\text{ s.t. }\; \forall\, k\ge k_0 : \\ \lim_{\tau\to\infty} \max\left\{ \sup_{T\ge\tau}\Big[\mathrm{dist}_{\|\cdot\|_\infty}(\hat{\theta}^\star_T(X(\omega)),\, \mathrm{DP}_d\cdot S(\omega))\Big],\, \varepsilon\right\} \;=\; \varepsilon \\ \text{for any }\; (\theta^\star_T)_{T\in\mathbb{N}} \equiv (\theta^\star_T(m,k,\omega))_{T\in\mathbb{N}} \equiv (\theta^{m|k}_T(\omega))_{T\in\mathbb{N}} \subset \Theta \;\text{ as in (4.40).}\end{gathered} \tag{4.91}$$

The above proof of (4.42), which did not involve any compactness assumption on D_X, remains valid without any changes, so that (4.91) holds if it can be derived from (4.42).

To this end, let $\Omega'' \in \mathscr{F}$ be the $\mathbb{P}$-full set on which the traces of X are all contained in D_X and (4.9) holds, and for each $n\in\mathbb{N}$ denote by Ω_n the $\mathbb{P}$-full set on which (4.42) holds for $\tilde{\varepsilon} = \frac{1}{n}$. Set $\tilde{\Omega} := \Omega''\cap\bigcap_{n\in\mathbb{N}}\Omega_n$ (another $\mathbb{P}$-full set) and let $\varepsilon > 0$ be arbitrary. Take any $\omega\in\tilde{\Omega}$. Then $\mathrm{tr}(X(\omega)) \subset K_{\nu_0}$ for some $\nu_0\in\mathbb{N}$, whence for any $n_0\in\mathbb{N}$ with $n_0^{-1} \le 2^{-\nu_0}\varepsilon/4$ we have that

$$\begin{gathered}\exists\, m_0 \equiv m_0(n_0) \ge 2 \;:\; \forall\, m\ge m_0 \;\text{ there is }\; k_0\equiv k_0(m)\in\mathbb{N} \;\text{ s.t. }\; \forall\, k\ge k_0 : \\ \alpha^{m|k}(\omega) \le n_0^{-1} \quad \text{and hence} \quad \sup\nolimits_{T\ge\tau_0} \big\|\theta^{m|k}_T(\omega)(X(\omega)) - \theta_T(X(\omega))\big\|_\infty \le \varepsilon\end{gathered}$$

(for some $\tau_0\,(\equiv \tau_0(\omega)) \in\mathbb{N}$ and some $(\theta_T)_{T\in\mathbb{N}}\,(\equiv (\theta_T(\omega))_{T\in\mathbb{N}}) \subset \Theta_\star$) by the exact same argumentation that led us to (4.45). □

4.7.8.2 Proof of Theorem 4.4.13 for Weakly Ergodic Observations

Let $(\tilde{X}, \mathcal{J})$ be a weakly ergodic observation. Adopting the setting and notation from pp. 101, suppose now that

$$\begin{aligned} &\forall\, \tilde{\varepsilon} > 0 \,:\, \exists\, m_0 \geq 2 \,:\, \text{for each } m \geq m_0 \text{ there is } k_0 \equiv k_0(m) \text{ such that:} \\ &\qquad \lim_{\tau\to\infty} \alpha_\tau^{m|k} \vee \tilde{\varepsilon} = \tilde{\varepsilon} \quad \text{in probability,}^{29} \quad \text{for each } k \geq k_0. \end{aligned} \tag{4.92}$$

Spelled out, (4.92) implies that for any given $(\varepsilon', \delta') \in (0,\infty)^2$ and $(m,k)\,(\equiv (m,k)_{\tilde{\varepsilon}}$ as in (4.92) for $\tilde{\varepsilon} := \varepsilon'/2)$

$$\text{there is} \quad \tau_* \equiv \tau_*(\varepsilon', \delta') \in \mathbb{N} \quad \text{such that} \quad \sup_{\tau\geq\tau_*} \mathbb{P}(\alpha_\tau^{m|k} \geq \varepsilon') \,\leq\, \delta'. \tag{4.93}$$

(Indeed: for any m, k as in (4.92) with $\tilde{\varepsilon} := \varepsilon'/2$, it holds $\mathbb{P}(\alpha_\tau^{m|k} \geq \varepsilon') \leq \mathbb{P}((\alpha_\tau^{m|k} \vee \frac{\varepsilon'}{2}) \geq \varepsilon') = \mathbb{P}(|(\alpha_\tau^{m|k} \vee \frac{\varepsilon'}{2}) - \frac{\varepsilon'}{2}| \geq \frac{\varepsilon'}{2}) \to 0$ as $\tau \to \infty$.) In particular, for any given $\mathfrak{p} \equiv (\varepsilon, \delta) \in (0,\infty)^2$ there will be $m_{\mathfrak{p}} \in \mathbb{N}$ such that for every $m' \geq m_{\mathfrak{p}}$ there is $k_{\mathfrak{p}} \equiv k_{\mathfrak{p}}(m')$ with the property that: for any $k' \geq k_{\mathfrak{p}}$ there is $\tau'_{\mathfrak{p}} \equiv \tau'_{\mathfrak{p}}(k') \in \mathbb{N}$ with

$$\varrho_{\mathfrak{p}} \,:=\, \sup_{\tau\geq\tau'_{\mathfrak{p}}} \mathbb{P}\big(\sup_{T\geq\tau} \mathrm{dist}_{\|\cdot\|_\infty}(\theta_T^{m'|k'}(X), \Theta_\star \cdot X) \geq \varepsilon\big) \,\leq\, \delta, \tag{4.94}$$

which due to $\sup_{\tau\geq\tau_{\mathfrak{p}'}} \mathbb{P}\big(\sup_{T\geq\tau} \mathrm{dist}_{\|\cdot\|_\infty}(\theta_T^{m'|k'}(X), \mathrm{DP}_d \cdot S) \,\geq\, \varepsilon\big) \,\leq\, \varrho_{\mathfrak{p}}$ implies that the asserted convergence (4.41) holds in probability. To see that (4.94) holds, fix $\varepsilon, \delta > 0$ and note

$$\Big\{\sup_{T\geq\tau} \mathrm{dist}_{\|\cdot\|_\infty}(\theta_T(X), \Theta_\star \cdot X) \geq \varepsilon\Big\} \cap \Omega' \quad\subseteq\quad \bigcup\nolimits_{\nu\in\mathbb{N}} A_\nu^{\hat{\theta}_\tau} \cap B_\nu \tag{4.95}$$

for any given sequence $\hat{\theta} \equiv (\theta_T)$ of Θ-valued random variables, $\tau \in \mathbb{N}$, and for the events[30] $A_\nu^{\hat{\theta}_\tau} := \{\sup_{T\geq\tau} d_\nu(\theta_T, \Theta_\star) \,\geq\, \varepsilon\}$ and $B_\nu := \{\sup_{t\in\mathbb{I}} |X_t| \,\geq\, r_{\nu-1}\}$, where r_ν denotes the radius of the 0-centered closed ball K_ν. Noting that $A_\nu^{\hat{\theta}_\tau} \subseteq A_{\nu+1}^{\hat{\theta}_\tau}$ and $B_{\nu+1} \subseteq B_\nu$ for all $\nu \in \mathbb{N}$, we from (4.95) obtain that

$$\mathbb{P}\big(\sup_{T\geq\tau} \mathrm{dist}_{\|\cdot\|_\infty}(\theta_T(X), \Theta_\star \cdot X) \geq \varepsilon\big) \;\leq\; \mathbb{P}(A_{\nu_0}^{\hat{\theta}_\tau}) \;+\; \mathbb{P}(B_{\nu_0+1}) \tag{4.96}$$

for any fixed $\nu_0 \in \mathbb{N}$. Denoting $\mu_X := \mathbb{E}[\sup_{t\in\mathbb{I}} |X_t|]$, Markov's inequality implies that

$$\mathbb{P}(B_{\nu_0+1}) \;\leq\; \frac{\mu_X}{r_{\nu_0}} \quad\longrightarrow\quad 0 \qquad (\nu_0 \to \infty), \tag{4.97}$$

while (4.90) implies $B_\varepsilon^{d_\nu}(\theta) \supseteq B_{\varepsilon/2}^{\rho_\nu}(\theta)$ for each $\theta \in \Theta$ (if $\varepsilon < 1$, assumable wlog) and hence

$$\mathbb{P}(A_{\nu_0}^{\hat{\theta}_\tau}) \;\leq\; \mathbb{P}(\sup_{T\geq\tau} \rho_{\nu_0}(\theta_T, \Theta_\star) \geq \varepsilon/2) \;\leq\; \mathbb{P}(\sup_{T\geq\tau} \tilde{d}(\theta_T, \Theta_*) \geq 2^{-\nu_0}\varepsilon/2). \tag{4.98}$$

[29] Remark that the (usual) notion of convergence in probability is well-defined for Θ-valued random variables since the topology of compact convergence on Θ is metrizable, second-countable (e.g. [110]) and, hence, separable.

[30] As the functions $\varphi_\theta : \tilde{\theta} \mapsto d_\nu(\tilde{\theta}, \theta)$ are continuous, their infimum $\varphi(\tilde{\theta}) := \inf_{\theta\in\Theta_\star} \varphi_\theta(\tilde{\theta}) = d_\nu(\tilde{\theta}, \Theta_*)$ is upper semicontinuous and hence Borel-measurable, whence the sets $A_\nu^{\hat{\theta}_\tau} = \{\sup_{T\geq\tau} \varphi(\theta_T) \,\geq\, \varepsilon\}$ are measurable. As X has continuous realisations, we have $\sup_{t\in\mathbb{I}} |X_t| = \sup_{t\in\mathbb{I}\cap\mathbb{Q}} |X_t|$ so that B_ν^θ is measurable.

Given (4.97) and (4.98), we may now fix an $\nu_0 \in \mathbb{N}$ large enough such that $\mathbb{P}(B_{\nu_0+1}) \leq \delta/2$, and for this choice of ν_0 obtain an $m_{\mathfrak{p}} \in \mathbb{N}$, as guaranteed by (4.92) for $\tilde{\varepsilon} = \tilde{\varepsilon}_\star$ with $\tilde{\varepsilon}_\star := 2^{-\nu_0}\varepsilon/4$, such that for every $m \geq m_{\mathfrak{p}}$ there is $k_{\mathfrak{p}}$ ($\equiv k_{\mathfrak{p}}(m)$) with the property that: for any $k \geq k_{\mathfrak{p}}$ there is $\tau_* \equiv \tau_*(2\tilde{\varepsilon}_\star, \delta/2) \in \mathbb{N}$, as guaranteed by (4.93), such that

$$\sup\nolimits_{\tau\geq\tau_*}\mathbb{P}(A_{\nu_0}^{\hat{\theta}_\tau}) \overset{(4.98)}{\leq} \sup\nolimits_{\tau\geq\tau_*}\mathbb{P}(\alpha_\tau^{m|k} \geq \tilde{\varepsilon}_\star) \;\leq\; \delta/2 \qquad \text{for } \hat{\theta} = (\theta_T^{m|k}).$$

Taken altogether, the estimate (4.96) then allows us to conclude that

$$\sup\nolimits_{\tau\geq\tau_*}\mathbb{P}(\sup\nolimits_{T\geq\tau}\mathrm{dist}_{\|\cdot\|_\infty}(\theta_T(X), \Theta_\star\cdot X) \geq \varepsilon) \;\leq\; \delta/2+\delta/2 \;\leq\; \delta,$$

which (via (4.43) and Thm. 4.2.3) yields the desired conclusion (4.41) for the weakly ergodic case.

It hence remains to prove (4.92), for which we may follow the previous lines of pp. 103 with only slight adaptations. Indeed: Since in the weakly ergodic case the Θ-uniform estimator convergence (4.37) holds in probability, we obtain – by way of the very same argumentation as for (4.49) – that

$$\lim_{T\to\infty}\bar{\kappa}_{m,k}(\theta_T^\star) \;=\; 0 \quad \text{in probability,} \quad \text{with } \; (\theta_T^\star) \equiv (\theta_T^{m|k}) \tag{4.99}$$

as in (4.40) for m, k as in (4.46) for some (arbitrary but) fixed $\tilde{\varepsilon} > 0$. From this we obtain that in the present context, the convergence (4.48) holds in probability. Indeed: Assuming otherwise implies the existence of $\varepsilon_0, \delta_0 > 0$ such that

$$\mathbb{P}(\mathrm{dist}(\theta_{T_j}^\star, \mathcal{M}) \geq \varepsilon_0) \;\geq\; \delta_0 \qquad \text{for each } j \in \mathbb{N}, \tag{4.100}$$

for some sequence $(T_j)_{j\in\mathbb{N}} \subset \mathbb{N}$. As a subsequence of $(\bar{\kappa}_{m,k}(\theta_T^\star))_{T\in\mathbb{N}}$, we by way of (4.99) find that $(\bar{\kappa}_{m,k}(\theta_{T_j}^\star))_{j\in\mathbb{N}}$ converges to 0 in probability, whence there is yet another subsequence $(T_{j_\ell})_{\ell\in\mathbb{N}}$ of $(T_j)_{j\in\mathbb{N}}$ such that $\lim_{\ell\to\infty}\bar{\kappa}_{m,k}(\theta_{T_{j_\ell}}^\star) = 0$ almost surely. Applying the (essentially) same argument which brought '(4.49) $\Rightarrow$ (4.48)' now yields that $\lim_{\ell\to\infty}\mathrm{dist}(\theta_{T_{j_\ell}}^\star, \mathcal{M}) = 0$ almost surely and hence in probability, contradicting (4.100).
As this proves $\lim_{\tau\to\infty}\sup_{T\geq\tau}\mathrm{dist}(\theta_T^{m|k}, \mathcal{M}) = 0$ in probability, we for any $\epsilon > 0$ obtain

$$\mathbb{P}\big(|\alpha_\tau^{m|k}\vee\tilde{\varepsilon}-\tilde{\varepsilon}| \geq \epsilon\big) \;\leq\; \mathbb{P}\big(\alpha_\tau^{m|k} \geq \tilde{\varepsilon}\big) \;\leq\; \mathbb{P}\big(\sup\nolimits_{T\geq\tau}\mathrm{dist}(\theta_T^{m|k}, \mathcal{M}) \geq \tilde{\varepsilon}/2\big) \;\to\; 0$$

as $\tau \to \infty$, where the last inequality is due to (4.47). This shows (4.92) as required. □

4.8 Numerical Experiments

We present a series of numerical examples to illustrate the practical applicability of our ICA method on discrete- and continuous-time signals. A complete account of the following experiments and results, including their full parameter settings and all relevant implementations and estimates, is provided on the public repository [140].

4.8.1 A Performance Index for Nonlinear ICA

As before, we consider processes X and S in continuous or discrete[31] time such that

$$X = f(S) \quad \text{for some} \quad f \in C^{2,2}(D_S). \tag{4.101}$$

In order to assess how close an estimate $\hat{S} \equiv \hat{S}(X)$ of S is to the true source S in (4.101), we propose to quantify the distance between $\hat{S}$ and the orbit[32] $\mathrm{DP}_d \cdot S$ by way of the following intuitive[33] performance statistic (cf. Remark 4.2.4 (iii) for applicability).

Definition 4.8.1 (Monomial Discordance)**.**Given two time series $\mathcal{X} := (X_t^1, \cdots, X_t^d)_{t\in\mathcal{I}}$ and $\mathcal{Y} := (Y_t^1, \cdots, Y_t^d)_{t\in\mathcal{I}}$ in $\mathbb{R}^d$ for $\mathcal{I}$ finite, define the *concordance matrix of* $(\mathcal{X}, \mathcal{Y})$ as

$$\mathcal{C}(\mathcal{X}, \mathcal{Y}) \;:=\; \left(\frac{1}{|\mathcal{I}|}\sum_{t\in\mathcal{I}} |\rho_{\mathrm{K}}(X_t^i, Y_t^j)|\right)_{(i,j)\in[d]^2} \;\in\; [0,1]^{d\times d}$$

where ρ_{K} is the Kendall[34] rank correlation coefficient. Furthermore, we define

$$\varrho(\mathcal{X}, \mathcal{Y}) \;:=\; \frac{1}{\sqrt{d(d-1)}} \min_{P\in\mathrm{P}_d} \|C(\mathcal{X}, \mathcal{Y}) - P\|_2 \;\in\; [0,1] \tag{4.102}$$

and call this quantity the *monomial discordance* of $\mathcal{X}$ and $\mathcal{Y}$.

Proposition 4.8.2. *Let X and S be as in* (4.101) *with S IC, and h be C^1-invertible on some open superset of D_X. Then for $\mathcal{I} \subset \mathbb{I}$ finite and ϱ as in* (4.102)*, we have that:*

$$\big(h(X_t)\big)_{t\in\mathcal{I}} \;\in\; \mathrm{DP}_d \cdot (S_t)_{t\in\mathcal{I}} \quad \textit{iff} \quad \varrho\big((h(X_t))_{t\in\mathcal{I}},\, (S_t)_{t\in\mathcal{I}}\big) = 0. \tag{4.103}$$

Proof. This is a direct consequence of the fact that Kendall's (and Spearman's) rank correlation coefficient ρ_{K} attains its extreme values ± 1 iff one of its arguments is a monotone transformation of the other (cf. e.g. [49, Theorem 3 (7.), (8.)]), combined with the fact that $\rho_{\mathrm{K}}(U, V) = 0$ if U and V are independent (cf. e.g. [49, Theorem 3 (2.)]).

Indeed, note for the 'if'-direction in (4.103) that $\varrho\big((h(X_t))_{t\in\mathcal{I}},\, (S_t)_{t\in\mathcal{I}}\big) = 0$ implies that there is $\sigma \in S_d$ with $|\rho_{\mathrm{K}}(h_i(X_t), S_t^j)| = \delta_{j,\sigma(i)}$ for each $t \in \mathcal{I}$ and $i \in [d]$, which by the above-mentioned property of ρ_{K} yields that $(h_i \circ f)(S_t) = \alpha_{i|t}(S_t^{\sigma(i)})$, and hence $(h_i \circ f)\big|_{\mathrm{supp}(S_t)} = \alpha_{i|t}\Big|_{\mathrm{supp}(S_t)}$, for some function $\alpha_{i|t} : \mathrm{supp}(S_t) \to \mathbb{R}$ with $\mathrm{supp}(S_t^{\sigma(i)}) \ni x_{\sigma(i)} \mapsto \alpha_{i|t}(x_{\sigma(i)}) \equiv \alpha_{i|t}(x)$ monotone. But since for each $i \in [d]$ we have $h_i \circ f \in C^1(\mathcal{O}_S)$ for some $\mathcal{O}_S \supset D_S$ open, the classical pasting lemma (see, for instance, [98, Corollary 2.8]) guarantees that the functions $(\alpha_{i|t} \mid t \in \mathcal{I})$ can be 'glued together' to an injective C^1-map $\alpha_i : \bigcup_{t\in\mathcal{I}} \mathrm{supp}(S_t^i) \to \mathbb{R}$ with $\alpha_i\big|_{\mathrm{supp}(S_t^i)} = \alpha_{i|t}\Big|_{\mathrm{supp}(S_t^i)}$ for each $t \in \mathcal{I}$, implying that $(h(X_t))_{t\in\mathcal{I}} = \big(P \cdot (\alpha_{\sigma^{-1}(1)} \times \cdots \times \alpha_{\sigma^{-1}(d)})(S_t)\big)_{t\in\mathcal{I}}$ for $P = (\delta_{\sigma(i),j})_{ij} \in \mathrm{P}_d$, as claimed. □

[31] See Section 3.7 and 4.5.0.1 for an explicated treatment of the latter.

[32] See (3.24) for notation, and recall that the elements of $\mathrm{DP}_d \cdot S$ are in a minimal distance from S.

[33] Recall the classical facts (e.g. [49]) that Kendall's (and Spearman's) rank correlation coefficient ρ_{K} attains its extreme values ± 1 iff one of its arguments is a monotone transformation of the other, with $\rho_{\mathrm{K}}(U, V) = 0$ if its arguments U and V are independent.

[34] If preferred, ρ_{K} might alternatively be chosen as Spearman's rank correlation coefficient.

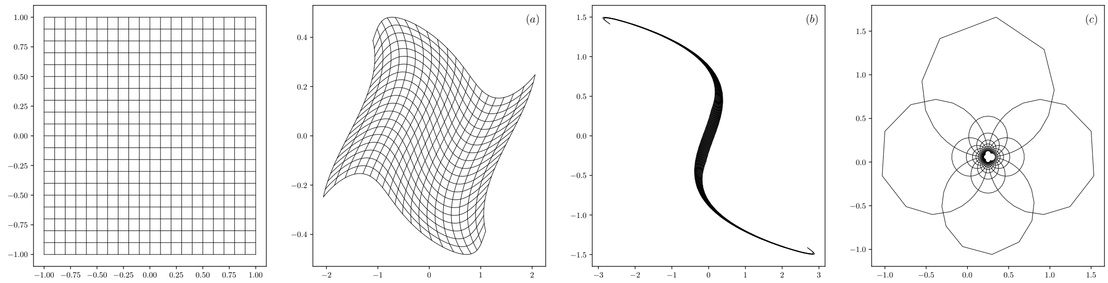

Figure 4.1: *The image of the square* $[-1,1]^2$ *(leftmost) under three increasingly nonlinear mixing transformations* f_1, f_2, f_3*, namely conjugates of the Hénon map (*f_1 *and* f_2*; panels (a) and (b), respectively) and of the Möbius transformation (*f_3*; panel (c)).*

Hence the smaller the monomial discordance between S and a transformation $h(X)$ of its observable, the closer to optimal will be the deviation between $h(X)$ and S.

Below we provide a brief synopsis of our experiments and the results that we obtained. For brevity, the truncated approximations (4.22) of the above contrast $\bar{\kappa}_{\text{IC}}$ will be denoted ϕ_{m_0}.

4.8.2 Nonlinear Mixings With Explicitly Parametrized Inverses

First we consider three families of C^2-diffeomorphisms on the plane whose inverses are explicitly parametrized.

More specifically: We sample two types of source processes in $\mathbb{R}^2$, namely: an IC Ornstein-Uhlenbeck process $S_{\text{ou}} = (S^1_{\text{ou}}, S^2_{\text{ou}})$, and an IC copula-based time-series $S_{\text{cy}} = (S^1_{\text{cy}}, S^2_{\text{cy}})$ that follows the dependence model (3.85).[35] Both S_{ou} and S_{cy} are contrastive by Prop. 3.5.5 (ii) and Cor. 3.5.2 (i), respectively.[36] These sources are first mapped to the square $[-1,1]^2$ upon centering and scaling them to unit amplitude, and then transformed by one of three mixing maps $f_j : \mathbb{R}^2 \to \mathbb{R}^2$ $(j = 1,2,3)$ with increasing degree of 'nonlinearity', see Figure 4.1. (For an explicit definition of the f_j, see [140].) Figure 4.2 shows the spatial trace of a sample realisation of S_{ou} and S_{cy} (panels (a) and (b)) next to an excerpt of the time-parametrised components of these realisations, together with their nonlinear mixtures $X^{(j)}_\eta := f_j(S_\eta)$ for $j = 1,2,3$ and $\eta =$ 'ou' (panels (c), (e), (g)) and $\eta =$ 'cy' (panels (d), (f), (h)).

[35] With $F_t^{S^i_{\text{cy}}}$ chosen as the cdf of $\mathcal{N}(0,1)$ and c chosen as the Clayton-density (cf. Proposition 3.5.1 (i)).

[36] Note further that the Ornstein-Uhlenbeck processes S_{ou}, while continuous-time by nature, are processed as discrete-time observations according to their classical Euler-Maruyama approximation. The copula-based time-series S_{cy}, on the other hand, are simulated at their observation frequency and thus showcase the applicability of our method to discrete-time signals (in accordance with Sections 3.7) and 4.5.0.1.

Each of the 'true' inverses $g^j := f_j^{-1}$ $(j = 1, 2, 3)$ are contained in an (injectively parametrized) family $\Theta_j \equiv \{g_\theta^j \in C^2(\mathbb{R}^2) \mid \theta \in \tilde{\Theta}_j\}$ of candidate de-mixing transformations g_θ^j, where $\tilde{\Theta}_j \subseteq \mathbb{R}^2$ is some open parameter set. On these parameter sets, we consider the data-based objective functions

$$\Phi_\eta^j : \tilde{\Theta}_j \to \mathbb{R}, \qquad \theta \mapsto \phi_{m_j}(g_\theta^j(X_\eta^{(j)})), \tag{4.104}$$

with $\phi_m := \bar{\kappa}_{\mathrm{IC}}^{[m]}$ as in (4.22) and capped at the cumulant orders $m_1 = m_2 = m_3 = 6$, and compare the topography of the functions (4.104) to that of the monotone discordances

$$\delta_\eta^j : \tilde{\Theta}_j \to \mathbb{R}, \qquad \theta \mapsto \varrho(g_\theta^j(X_\eta^{(j)}), S_\eta) \qquad \text{(cf. (4.102))}. \tag{4.105}$$

Recall that the latter are 'distance functions' that quantify how much a candidate source estimate $\hat{S}_\eta^\theta := g_\theta^j(X_\eta^{(j)})$ deviates [from the monomial orbit $\mathrm{DP}_d \cdot S_\eta$ of S_η, that is] from the true source S_η up to order and monotone scaling of its components.

The results are displayed in the first three columns of Figure 4.3, with the 'estimator's view' $\Phi_{\mathrm{ou}}^{1|2|3}$ of the demixing performance shown in the top-row panels and the 'true view' $\delta_{\mathrm{ou}}^{1|2|3}$ of the demixing performance shown in the bottom-row panels.[37] This shows clearly that within the given families Θ_j of candidate transformations, those candidate nonlinearities which map the data $X_\eta^{(j)}$ to a best-approximation of its source S_η are precisely those that minimise the contrast (4.104), as asserted by Theorem 4.2.3.

An analogous experiment $(j = 4)$ is performed for a mixing transformation $f_4 : \mathbb{R}^3 \to \mathbb{R}^3$, see Figures 4.4, 4.5. The results, obtained for a contrast capped at cumulant order $m_4 = 7$ and shown as the rightmost column of Figure 4.3, are again affirmative of Theorem 4.2.3.

4.8.3 Nonlinear Mixings With Inverses Approximated By Neural Networks

The practical applicability of our ICA-method is illustrated by running the optimisation (4.9) over (approximate) demixing-transformations which are modelled by an artificial neural network.

More specifically: We subject two Ornstein-Uhlenbeck sources $S^{(1)}$ and $S^{(2)}$ with two resp. four independent components to a two- resp. four-dimensional nonlinear mixing transform

[37] For brevity, Figure 4.3 shows the case $\eta = \mathrm{ou}$ only; the results for the case $\eta = \mathrm{cy}$ can be found in [140].

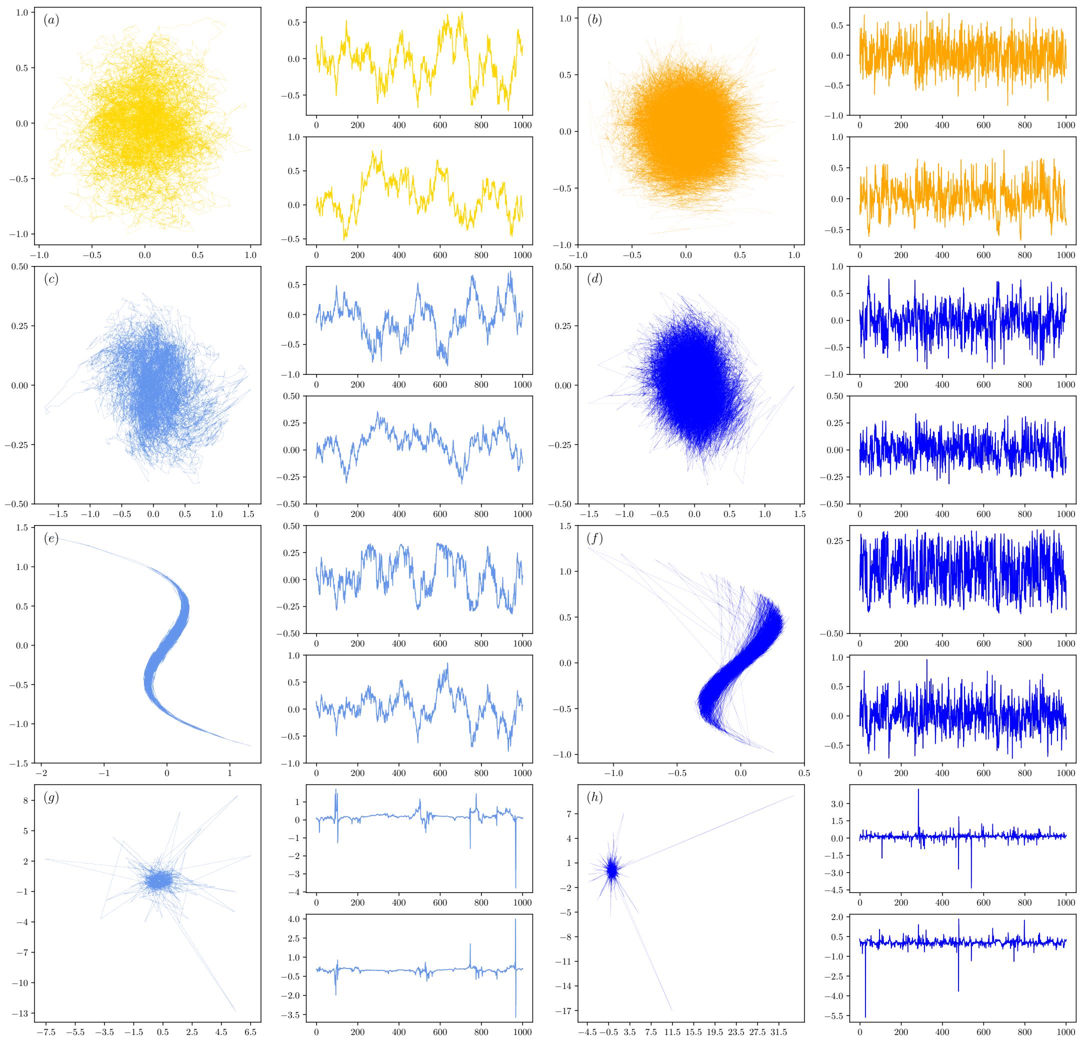

Figure 4.2: *Spatial traces and sampled components of three nonlinear mixtures of the sources* S_{ou} *and* S_{cy} *(panels (a) and (b), respectively). Depicted are the mixtures* $X^{(1)}_{\mathrm{ou}}$ *and* $X^{(1)}_{\mathrm{cy}}$ *((c) and (d)),* $X^{(2)}_{\mathrm{ou}}$ *and* $X^{(2)}_{\mathrm{cy}}$ *((e) and (f)), and* $X^{(3)}_{\mathrm{ou}}$ *and* $X^{(3)}_{\mathrm{cy}}$ *((g) and (h)). The components of the mixtures, excerpted over 1000 data points each, are shown to the right of each panel.*

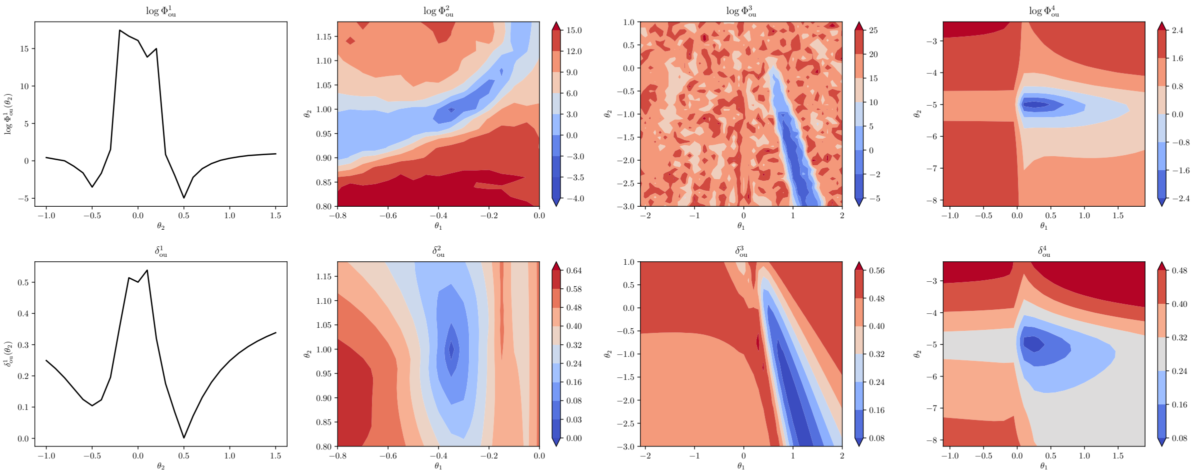

Figure 4.3: *Contour plot (leftmost column) and heatmaps of the log-transformed contrast functions* (4.104) *(top row) and of the associated discordance functions* (4.105) *(bottom row) for the mixings* $X_{\text{ou}}^{(j)} = f_j(S_{\text{ou}})$, $j = 1, \ldots, 4$. *The parameters* $\theta_\star^{(j)} \equiv (\theta_1^{(j)}, \theta_2^{(j)})$ *of the true inverses* $f_j^{-1} \equiv g^j_{\theta_\star^{(j)}} \in \Theta_j$ *are* $\theta_\star^{(1)} = 0.5$,[38] $\theta_\star^{(2)} = (-0.35, 1)$, $\theta_\star^{(3)} = (1, -2)$, *and* $\theta_\star^{(4)} = (0.2, -5)$.

(see [140] for details). The resulting mixtures $X^{(1)}$ and $X^{(2)}$ are then passed on to candidate demixing-nonlinearities $g^\nu_\theta \in \Theta_\nu$ which are given as elements of the parametrized families

$$\Theta_\nu \ := \ \left\{ g^\nu_\theta \,:\, \mathbb{R}^{2\nu} \to \mathbb{R}^{2\nu} \;\middle|\; g^\nu_\theta \ \text{ is an ANN with weights } \ \theta \in \tilde{\Theta}_\nu \right\} \qquad (\nu = 1, 2). \quad (4.106)$$

Here, the families of transformations Θ_ν are spanned by the various configurations of some artificial neural network (ANN) instantiated over weight-vectors θ which are chosen from a given parameter set $\tilde{\Theta}_\nu$ in $\mathbb{R}^{m_\nu}$, where the number of weights m_ν is part of the pre-defined architecture of the ANN. Given these candidate-inverses, the optimisations (4.9) are run by

$$\text{minimizing} \qquad \tilde{\Theta}_\nu \ \ni \ \theta \quad \longmapsto \quad \phi_{m_\nu}(g^\nu_\theta(X^{(\nu)})), \qquad (4.107)$$

i.e. by training each constituent ANN (4.106) with the truncated contrast $\phi_{m_\nu} = \bar{\kappa}_{\text{IC}}^{[m_\nu]}$ (cf. (4.22)) as its loss function, where the optimization steps are computed via backpropagation along the weights of the ANN. Technical details for the respective setups of (4.106) and (4.107) are reported in ([140] and) Remark 4.8.4.

For the case $\nu = 1$ we applied the mixing transformation depicted in Figure 4.6 (leftmost panel), and for the case $\nu = 2$ we followed the simulations of [82, 83] in using as a mixing transformation an invertible feedforward-neural network with four-nodal in- and output

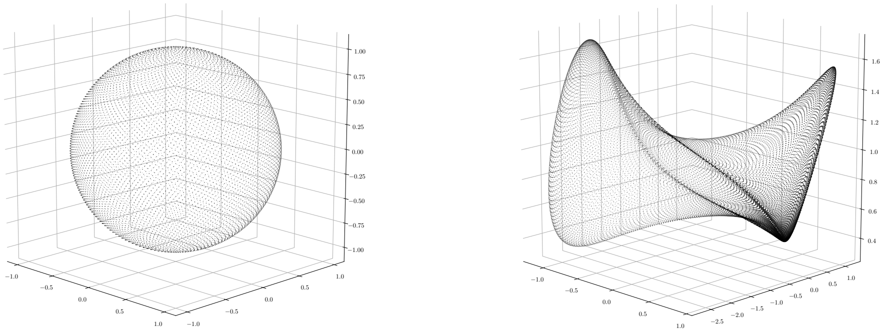

Figure 4.4: *Illustration of the three-dimensional mixing transform* $f_4 : \mathbb{R}^3 \to \mathbb{R}^3$ *via its action* $f_4(S^2)$ *(right panel) on the 2-sphere* $S^2 \equiv \{x \in \mathbb{R}^3 \mid |x| = 1\}$ *(left panel).*

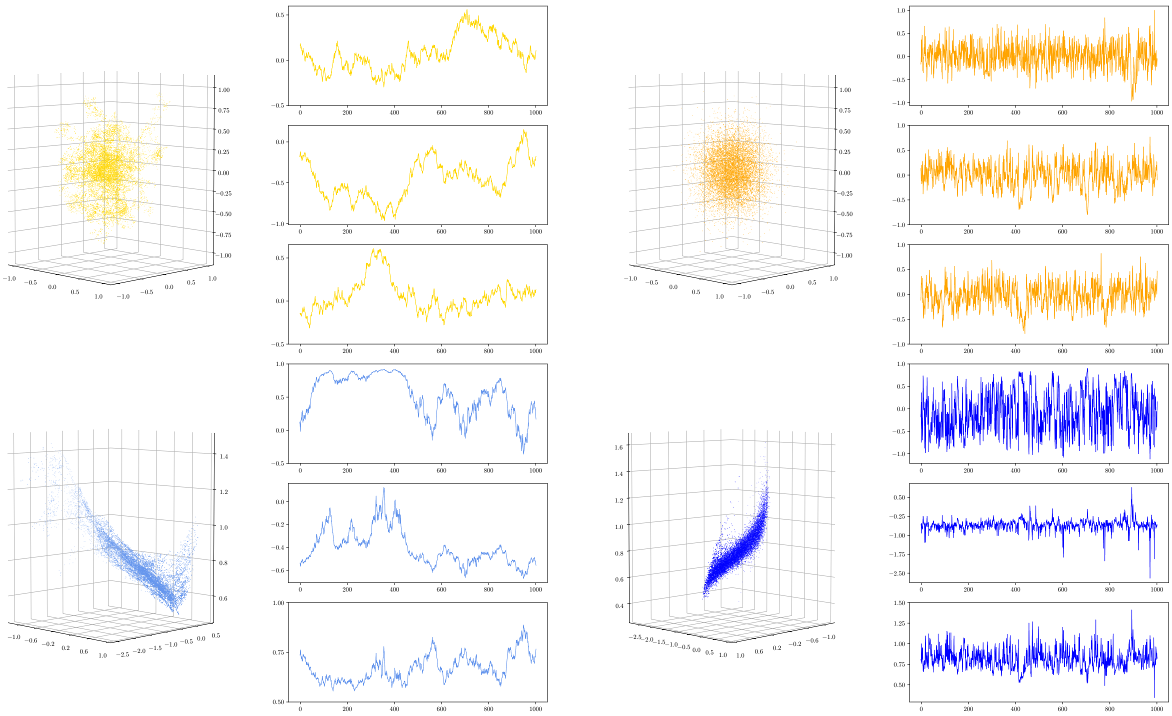

Figure 4.5: *Spatial trace and sampled components of a three-dimensional IC Ornstein-Uhlenbeck process* $\tilde{S}_{\mathrm{ou}}$ *(top left) and an IC copula-based time-series model* $\tilde{S}_{\mathrm{cy}}$ *(top right) and their respective nonlinear mixtures* $f_4(\tilde{S}_{\mathrm{ou}})$ *(bottom left) and* $f_4(\tilde{S}_{\mathrm{cy}})$ *(bottom right).*

layers and two four-nodal hidden layers with tanh activation each.

Denoting by $\theta_\nu^* \in \tilde{\Theta}_\nu$ the (local) optimum obtained by the minimisation of the objective (4.107) and setting $\hat{S}^{(\nu)} := g^\nu_{\theta_\nu^*}(X^{(\nu)})$ for the associated estimate of the source $S^{(\nu)}$ (cf. (4.9)), we as results to these experiments obtained the concordance matrices (cf. Def. 4.8.1)

$$\mathcal{C}(\hat{S}^{(1)}, S^{(1)}) \doteq \begin{pmatrix} \mathbf{0.853} & 0.065 \\ 0.079 & \mathbf{0.930} \end{pmatrix} \quad \text{and} \tag{4.108}$$

$$\mathcal{C}(\hat{S}^{(2)}, S^{(2)}) \doteq \begin{pmatrix} \mathbf{0.834} & 0.003 & 0.037 & 0.016 \\ 0.148 & \mathbf{0.725} & 0.109 & 0.069 \\ 0.037 & 0.034 & \mathbf{0.803} & 0.265 \\ 0.077 & 0.131 & 0.072 & \mathbf{0.787} \end{pmatrix}, \tag{4.109}$$

where we corrected for the permutation ambiguity between $\hat{S}$ and S to simplify comparison.

Both (4.108) and (4.109) indicate a good fit between $\hat{S}^{(\nu)}$ and $S^{(\nu)}$ in the sense that, to a good approximation, $\hat{S}^{(\nu)}$ and $S^{(\nu)}$ differ only up to (an inevitable permutation and) monotone scaling of their components,[39] as stated by Theorem 4.2.3. A visual comparison of the original samples $S^{(1)}, S^{(2)}$ and their estimates $\hat{S}^{(1)}, \hat{S}^{(2)}$, see Figures 4.6 and 4.7, confirms these results.

To reaffirm that the above results of finding good approximations to the source are not simply due to chance, we ran our experiments repeatedly with randomly chosen realisations and initial configurations for the data and the learning process (4.106) & (4.107), see [140] and Figure 4.8. The obtained discordances have mean 0.21 and standard deviation $45 \cdot 10^{-3}$ for the Ornstein-Uhlenbeck mixture (Fig. 4.8 (a)), and mean 0.18 and standard deviation $55 \cdot 10^{-3}$ for the copula-based time series mixture (Fig. 4.8 (b)). The associated average concordance matrix for this first (Ornstein-Uhlenbeck) type of mixtures is $\left(\begin{smallmatrix} 0.78 & 0.10 \\ 0.11 & 0.89 \end{smallmatrix}\right)$, and for the latter (copula) type of mixtures it is $\left(\begin{smallmatrix} 0.83 & 0.11 \\ 0.11 & 0.88 \end{smallmatrix}\right)$. The average discordance between the respective sources and their initial guesses $g_{\theta_0}(X)$ prior to applying (4.107) are at 0.59 (a) and 0.56 (b), respectively.

These experiments underline the practical applicability of our proposed ICA-method.

To conclude, we note the following empirical findings.

[38] Notice that: (a) by definition of Θ_1, the function Φ^1_{ou} depends on the one-dimensional parameter θ_2 only; (b) as the concordance matrix of $\hat{S}_{-0.5} := g^{(1)}_{-0.5}(X^{(1)}_{\mathrm{ou}})$ and S_{ou} is $\left(\begin{smallmatrix} 0.053 & 0.929 \\ 0.834 & 0.099 \end{smallmatrix}\right)$ (indicating a close proximity between $\hat{S}_{-0.5}$ and $\mathrm{DP}_d \cdot S_{\mathrm{ou}}$, cf. Prop. 4.8.2), the observation of Φ^1_{ou} attaining a low local minimum at -0.5 is in accordance with Theorem 4.2.3.

[39] Recall that the optimal deviation $\hat{S}^{(\nu)} \in \mathrm{DP}_d \cdot S^{(\nu)}$ between $\hat{S}^{(\nu)}$ and $S^{(\nu)}$ is achieved iff (4.108) and (4.109) are permutation matrices (Proposition 4.8.2).

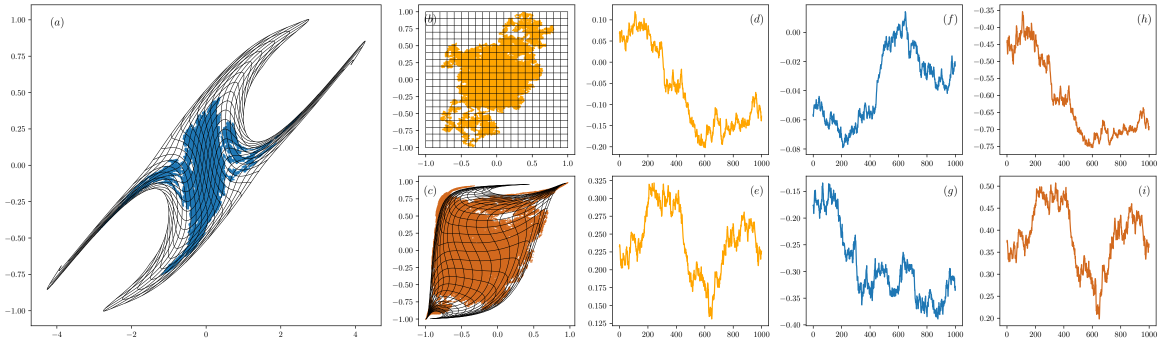

Figure 4.6: *Nonlinear mixture X (sampled trace (a) and components (f), (g)) of an IC Ornstein-Uhlenbeck source S ((b) and (d), (e)). Further shown is the residual $g \circ f|_{[-1,1]^2}$ ((c); cf.* (4.110)*) for an estimate g of $f^{-1}|_{D_X}$. The function g is found by optimising an artificial neural network (g_θ) via the loss function* (4.107)*, and the resulting estimate $\hat{S} := g(X)$ of S is shown in brown ((c) and (h), (i)). To a good approximation, the source S and its estimate coincide up to (a transposition and) a monotone scaling of their components, as quantified by* (4.108)*.*

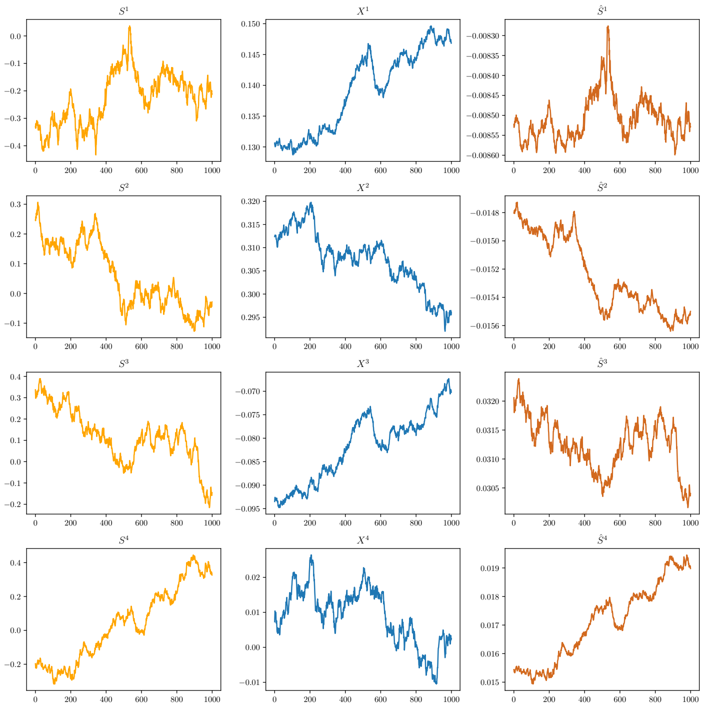

Figure 4.7: *Component processes, excerpted over 1000 data points each, of an IC Ornstein-Uhlenbeck process $S = (S^1, S^2, S^3, S^4)$ (orange), a nonlinear mixture $X = (X^1, X^2, X^3, X^4)$ (blue) of S, and an estimate $\hat{S} = (\hat{S}^1, \hat{S}^2, \hat{S}^3, \hat{S}^4)$ of S (brown). The estimate $\hat{S}$ is obtained as $\hat{S} = g_{\theta_\star}(X)$ where (g_θ) is an ANN and $\theta_\star$ is a (local) minimum of the associated objective* (4.107)*. To a good approximation, the processes $\hat{S}$ and S coincide up to (a permutation and) a monotone scaling of their components. This is in accordance with Theorem 4.2.3 and as quantified by* (4.109)*.*

Remark 4.8.3 (Empirical Comments). (i) Given an observable $X = f(S)$ together with a family Θ of candidate transformations on $\mathbb{R}^d$, the technical compatibility condition $(\mathrm{DP}_d(D_S) \cdot f^{-1})\big|_{D_X} \cap \Theta\big|_{D_X} \neq \emptyset$ of Theorem 4.2.3 can in practice typically not be guaranteed a priori. However, as indicated by the above findings (4.108) and (4.109), infringements of this (sufficient) technical condition might typically be innocuous, provided that at least

$$(\mathrm{DP}_d \cdot g)\big|_{D_X} \cap \Theta\big|_{D_X} \neq \emptyset \quad \text{for some } g \text{ with } g\big|_{D_X} \text{ `close enough' to } f^{-1}\Big|_{D_X}, \tag{4.110}$$

which will be satisfied if Θ is chosen large enough, say as a suitable ANN or another universal approximator. In a similar vein, our experiments indicate that the regularity condition $\Theta \subseteq C^2(D_X)$ may in practice be softened by merely requiring that the 'approximate inverse' g in (4.110) be 'C^2-invertible on most of D_X' (cf. e.g. Figure 4.6, panel (c)) and the parametrization of Θ be 'continuous' at (some) point $\tilde{g} \in \Theta$ with $\tilde{g}\big|_{D_X} \in \mathrm{DP}_d \cdot g\big|_{D_X}$, though this a priori reduces the optimisation (4.9) to the search for a (low) local minimum.

(ii) We emphasize that the configurations of the neural networks and their backpropagation that we used in our experiments were ad hoc and not tuned for approximational optimality. Since the loss functions (4.107) are typically non-convex with their topography crucially depending on the choice of (4.106) (cf. e.g. Figure 4.3), we expect that the accuracy and efficiency of our estimates may be significantly improved by applying our ICA-method to ANN-based approximation schemes (4.106), (4.107) which are more carefully designed.

Remark 4.8.4 (Implementation Details). (The following enumeration ($\nu = 1, 2$) refers to the mixtures $X^{(2)} = f_\nu(S^{(2)})$ considered in Section 4.8.3.) For the case $\nu = 1$, we applied as f_1 the mixing transformation depicted in Figure 4.6 (leftmost panel), and as the parametrising ANN Θ_1 we chose a feedforward neural network with a two-nodal in- and output layer and two hidden layers consisting of 4 resp. 32 neurons with tanh activation each; the cumulant series (4.7) was capped at maximal cumulant order $m_1 = 6$. For the case $\nu = 2$, we followed the simulations of [82, 83] in using as a mixing transformation f_2 an invertible feedforward-neural network with four-nodal in- and output layers and two four-nodal hidden layers with tanh activation each, and as the parametrising ANN chose a feedforward network with a four-nodal in- and output layer and one hidden layer of 1024 neurons and uniformly-weighted Leaky ReLU activations; the contrast function (4.7) was capped at the maximal cumulant order $m_2 = 5$. For both $\nu = 1, 2$, the resulting

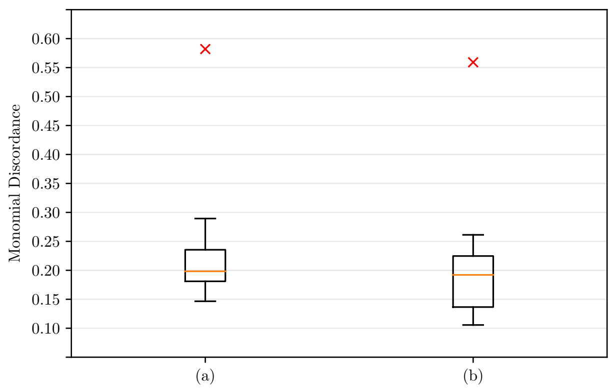

Figure 4.8: *Discrepancies (quantified by* (4.102)*, boxplotted) between the true source S and its* (4.107)*-based estimates $\hat{S}$. The latter are computed from the nonlinear mixture X shown in Figure 4.6. Displayed are the cases where: (a) the source S an IC Ornstein-Uhlenbeck process, and (b) the source S is a copula-based IC time-series (b), respectively. Shown in addition (red crosses) is the average discordance between S and its initial estimates at the start of the optimisation* (4.107)*. (The optimised discordances and their averages pre-learning are computed by applying our method to ten realisations of S and ten initial configurations g_{θ_0} of the ANNs* (4.106) *respectively, all drawn independently at random.)*

loss functions (4.107) were optimised using stochastic gradient descent (Adam) with non-vanishing ℓ_2-penalty.

4.9 Conclusion

Chapters 3 and 4 have addressed the problem of Blind Source Separation via the classical approach of Independent Component Analysis. As our main contribution, we have formulated and proved a statistical method to recover multidimensional stochastic processes (in both continuous- as well as discrete-time) from observations of their nonlinear mixtures. Conceptually, our method assumes a source process with independent component processes and, by exploiting the temporal structure of this source, characterises its nonlinear transformations by the degree of intercomponental statistical dependence that they inflict on the source. Quantifying the latter by way of an efficiently computable contrast function derived from the signature cumulants of a stochastic process, the initial source separation problem may then be reformulated as a provably robust problem of optimisation-based function approximation which in practice can be conveniently implemented by, e.g., contemporary neural network-based learning schemes. A comprehensive consistency analysis ensures that the resulting method is usable in real-world situations (discretized time, one sample trajectory), which is further illustrated by a number of theoretical and numerical examples.

The mathematics of the identifiability theory established in this work appears flexible enough to allow for extensions in various further directions. For instance, by considering third-order in place of second-order finite-dimensional distributions it may be adapted to infer the identifiability of stochastic processes from their time-dependent nonlinear mixing transformations ('invertible flows'). By adapting the ideas of our approach further, it does now also seem within reach to prove the identifiability of stochastic sources from more general nonlinear relations, such as for instance in the setting of controlled differential equations where one may be interested to recover an (independent-component) stochastic control from its nonlinear response. As with most methods involving an optimisation over flexibly parametrisable nonlinearities, however, a significant practical caveat of our approach is the occurence of spurious local minima in the approximation of the demixing transformation. This leaves room for improvement that future research might explore: In addition to practical deliberations such as spanning the optimisation domain by more carefully designed learning architectures, or amplifying the contrast function by the addition of tunable hyperparameters such as weights attached to its summands, one may attempt to tame the critical optimisation task by adjusting it to (localised) polynomial approximations of the mixing nonlinearity and harvesting the additional algebraic structure that then results from the fact (Proposition 5.9.9) that the signature transform 'dualises' the action of polynomial transformations on its arguments.

Chapter 5

Robust Blind Source Separation

5.1 Introduction

The Problem of Blind Source Separation concerns the (local) inversion of an unknown function given only the image under this function of some likewise unobserved argument:

$$\text{Provided only } X = f(S) \text{ for } f \text{ and } S \text{ unknown, invert } X \text{ for } S. \tag{5.1}$$

This problem is central to a great variety of applications ranging from uses in medical imaging, neuroscience and biology [8, 45, 99, 106] over finance and astronomy [10, 118] to physics and engineering [3, 41, 147] to name but a few, witness also the large number of practical and theoretical approaches that have been proposed to address it; see Chapter 1 and the references therein for an overview. With the generic formulation (5.1) of the BSS-problem being highly underdetermined, all of these methods have to operate under some additional assumptions on f and S, so-called 'identifiability assumptions', to guarantee that a meaningful blind inversion $X \mapsto S$ is possible at all.

The by far most successful ansatz to date to interpret and solve (5.1) assumes the hidden relation f to be linear and models the source S as a multidimensional random variable with statistically independent components. Commonly referred to as Independent Component Analysis (ICA), this framework is based on probabilistic insights obtained independently by Reiersøl, Skitovich, and Darmois [130, 145, 43] and was influentially formulated as an optimisation problem by Comon [37], see Section 3.1. Since its (surprisingly recent [81]) conception, ICA has seen various adaptations and extensions, the majority of those concerning modifications and analyses of its underling optimisation procedures or variations of the source model in terms of different (identifiability-sufficient) 'non-degeneracy' assumptions on its component distributions, cf. Chapter 1 or [38, 81, 111] and the references therein.

A comparatively under-explored yet both obvious and fundamental question, especially in light of how extensively ICA-methods are used across their considerable spectrum of applications, is if and to what extent the solutions to (5.1) persist when their underlying identifiability assumptions on f and S are violated. The scope and relevance of this question and its implications are illustrated by the classic Example 1.1.2 on p. 12.

Formally speaking, each of the infringements listed in Example 1.1.2 constitutes a potential violation of the structural prior assumptions on source and mixing that underlie the vast majority of ICA procedures, including all of the most popular such methods to date, cf. [38, 81, 111]. Referring to the pair (S, f) in (5.1) as 'the cause' of X, the central question then becomes how and to what degree such infringements of the prior assumptions on the cause of an observation affect the algorithmic consistency of an ICA procedure, and whether this consistency can be meaningfully quantified against general distortions of its supporting identifiability assumptions.

Although according investigations into the robustness of BSS and ICA solutions have been repeatedly called for [80], questions of this sort have received surprisingly little rigorous attention in the literature. To the best of our knowledge, there currently is no robustness theory for BSS or ICA to analyse the above questions comprehensively, and available results are limited to either proofs of statistical consistency[1], e.g. [29, 30, 138], or partial sensitivity analyses of numerical subroutines [1, 24, 144] or are asymptotical [29] or confined to highly specific BSS-models [11, 92] or to very particular or (semi-)parametric noise or dependence models only [12, 89, 125, 139], and even those few results are mostly restricted to the 'time-independent' special case of X being modelled as an iid sequence of random vectors.

In this chapter, we aim to close these gaps by providing a flexible and comprehensive robustness theory for the Blind Source Separation of (iid and) time-dependent observables. Modelling the source S in (5.1) as a discrete- or continuous-time stochastic process in $\mathbb{R}^d$ and working under the general paradigm of ICA, the following are our main contributions.

Formalised Robustness for BSS. Based on a concise analysis of the abstract BSS problem (5.1), we propose a general and flexible topological notion of statistical robustness[2] for Blind Source Separation and Independent Component Analysis that covers a broad range of relevant infringement scenarios (Section 5.3).

[1] Consistency results could be seen as 'asymptotic (partial) robustness with respect to empirical approximation' as these establish convergence of solutions along certain sequences of sample-based estimators (but not necessarily along all convergent sequences, which would characterize robustness under a sequential topology).

[2] Our usage of the word 'robust' is in the sense of Huber [75], as is made precise in Section 5.3.

ICA-Tailored Topology on Causes. We obtain a natural and coarse premetric topology on the space of laws of stochastic processes by considering a standardized difference between elements of a graded family of path-space functionals pertaining to the signature map from rough paths theory, and extend this to an informative ICA-tailored topology on the space of hidden causes underlying a BSS-observable. In conjunction with our formalisation of robustness, this topology provides a general and model-free way to capture violations of idealised identifiability assumptions and relate them, in a finite (i.e. 'non-asymptotic') and conveniently quantifiable manner, to the resulting effect that they have on the accuracy of an ICA-procedure (Sections 5.4 and 5.5).

General Stability for Blind Inversion. As a first vehicle on which to demonstrate our robustness theory, we present a generic blind-inversion map that performs (5.1) under ICA-typical identifiability assumptions but is the first to achieve this with unified stability guarantees for very general deviations from its underlying structural assumptions. Conceptionally, this inversion map has the property of being continuous with respect to the aforementioned ICA-topology on the space of causes, and its robustness guarantees are due to informative and entirely explicit moduli of continuity with regards to this topology (Section 5.6). The significance for applications of our approach is illustrated in a series of statistical use cases (Section 5.7), including the demonstration of provable robustness for the practical infringements listed in Example 1.1.2.

In the interest of a compact exposition, most proofs and auxiliary results are relegated to the end of the chapter. This chapter further assumes familiarity with Section 5.2.1 and its notation, and we also adopt the following convention.

Convention 5.1.1 (Identifying Signals With Their Laws)**.** From Section 5.3.1 onwards, we will tacitly identify a given signal Y in $\mathbb{R}^d$ (a random variable on $\mathcal{C}_d$) with its law $\mathbb{P}_Y := \mathbb{P} \circ Y^{-1}$ (a probability measure on $\mathcal{C}_d$). This justifies the prevalent abuse of notation $Y\,[\cong \mathbb{P}_Y] \in \mathcal{M}_1(\mathcal{C}_d)$.

(Notice that the identification in Convention 5.1.1 is equivalent to the perspective of considering two [$\mathcal{C}_d$-valued] random variables 'equal' if they are equal in distribution.)

Notice finally that unlike in the rest of this thesis (cf. Definition 2.1.4), in this chapter the symbol $\|\cdot\|_{p\text{-var}}$ denotes the p-variation *norm* (5.32).

5.2 Preliminaries

5.2.1 Discrete-Time and Continuous-Time Signals

Throughout this chapter, the term *signal* refers to a stochastic processes in $\mathbb{R}^d$ (Definition 3.2.1), that is to an ordered family

$$Y = (Y_t \mid t \in \mathbb{I}) \quad \text{of (Borel) random vectors} \quad Y_t = (Y_t^1, \cdots, Y_t^d) \text{ in } \mathbb{R}^d, \tag{5.2}$$

with $\mathbb{I} \subset \mathbb{R}$, typically either $\mathbb{I} = \mathbb{N}$ ('discrete-time') or $\mathbb{I} = [0,1]$ ('continuous-time'). The scalar processes $Y^i = (Y_t^i \mid t \in \mathbb{I}) \equiv (Y_t^i)$, $i \in [d]$, are called the components of Y.

By rescaling $\mathbb{I}$ and due to Rem. 5.2.1, we may assume $\mathbb{I} = [0,1]$ and that the vectors of any signal (5.2) are ordered continuously without loss of generality (cf. Rem. 5.4.4), i.e. that

$$\mathbb{I} \ni t \longmapsto Y_t \quad \text{is continuous} \tag{5.3}$$

with probability one. This continuity assumption (5.3) allows us to

- use the term 'stochastic process' synonymously to 'continuous-time stochastic process', as we will do unless stated otherwise, and to
- view a signal (5.2) as a random path, that is as a $\mathcal{C}_d$-valued (Borel) random variable over some probability space $(\Omega, \mathscr{F}, \mathbb{P})$.

Remark 5.2.1. Any signal $Y = (Y_t)_{t\in\mathbb{I}}$ with $\mathbb{I}$ discrete can be viewed as a continuous-time signal $\hat{Y}$ via piecewise-linear interpolation of its points $(Y_t \mid t \in \mathbb{I}) \equiv (Y_{t_j} \mid j \in \mathbb{N})$, namely

$$\hat{Y} = (\hat{Y}_t := Y_{t_j} + \tfrac{t-t_j}{t_{j+1}-t_j}(Y_{t_{j+1}} - Y_{t_j}) \text{ for } t \in [t_j, t_{j+1})).$$

(Note that this interpolation commutes with linear transformations: $A\cdot\hat{Y} \equiv (A\hat{Y}_t) = \widehat{(AY_{t_j})}$ for each $A \in \mathbb{R}^{d\times d}$.) Upon scaling $\mathbb{I}$, we may further assume $\hat{Y}$ to be $\mathcal{C}_d$-valued. ◆

By identifying a signal with its law (Convention 5.1.1) we can embed[3] the class of all signals into $\mathcal{M}_1 \equiv \mathcal{M}_1(\mathcal{C}_d)$, the space of all (Borel) probability measures on $\mathcal{C}_d$. This is of significant mathematical convenience as it allows us to describe the BSS-problem (5.1) in terms of maps on a single pre-structured domain ($\mathcal{M}_1$), from where the notion of robustness can then be flexibly formulated and analysed from the angle of general topology (see Section 5.3).
We will at times make use of the spatial support (3.7) of a signal (Definition 3.2.4), and call a signal Y 'degenerate' if its spatial support D_Y is contained in a hyperplane of $\mathbb{R}^d$.

The support (3.7) and other notions and operations on signals (5.2) can be readily generalised to measures in $\mathcal{M}_1$, as per the following subsection.

[3] ... up to equality in distribution, which for our purposes here is entirely sufficient ...

5.2.2 Operations on Measures

The Problem of Blind Source Separation naturally involves a range of (non)linear transformations φ applied to a signal. Since the space $\mathcal{M}_1(\mathcal{C}_d) \equiv \mathcal{M}_1$ of (Borel) probability measures on $\mathcal{C}_d$ is convex but not a vector space, these transformations need to be 'induced' on $\mathcal{M}_1(\mathcal{C}_d)$ in order to be well-defined.

Given some target vector space V, a canonical way to do this is via the pushforward

$$\varphi_* \,:\, \mathcal{M}_1 \ni \mu \longmapsto \varphi_*(\mu) \equiv \mu \circ \varphi^{-1} \in \mathcal{M}_1(V) \quad \text{of a (Borel) map } \ \varphi : \mathcal{C}_d \to V. \tag{5.4}$$

Owing to their ubiquity, a few special signal actions (5.4) will be given their own name: For $x = (x_t^1, \cdots, x_t^d)_{t\in[0,1]} \in \mathcal{C}_d$ any fixed path, consider the maps (for $i \in [d],\ t \in [0,1]$)

$$\pi_i \,:\, x \mapsto x^i \equiv (x_t^i)_{t\in[0,1]} \quad \text{and} \quad \pi_t \,:\, x \mapsto x_t \quad \text{and} \quad \pi_t^i \,:\, x \mapsto x_t^i \tag{5.5}$$

and denote their respectively induced pushforward measures from a signal $\mu \in \mathcal{M}_1$ by $\mu^i := \mu\circ\pi_i^{-1}$ and $\mu_t := \mu\circ\pi_t^{-1}$ and $\mu_t^i := \mu\circ(\pi_t^i)^{-1}$; we may then also write $\mu = (\mu^1, \cdots, \mu^d)$ and $\mu = (\mu_t)$ to denote that μ has the spatial marginals μ^i and the fixed-time marginals μ_t.

In accordance with (3.7), the spatial support of $\mu \in \mathcal{M}_1$ is defined as $D_\mu := \overline{\bigcup_{t\in\mathbb{I}} \operatorname{supp}(\mu_t)}$. Any $g : \mathbb{R}^d \to \mathbb{R}^d$ induces a map $\tilde{g} : \mathcal{C}_d \to \mathcal{C}_d$ via $\tilde{g}(x) := (g(x_t))_{t\in[0,1]}$, and we set

$$g(\mu) \equiv g_*\mu \,:=\, \mu \circ \tilde{g}^{-1}. \tag{5.6}$$

Extending (5.5), we may further declare in-time increments $\mu_{s,t} \equiv \mu_t - \mu_s := (\pi_t - \pi_s)_*(\mu)$ and products $\mu_{s,t}^i{\cdot}\mu_{s,t}^j := \left[\mathfrak{m}\circ(\pi_{s,t}^i \times \pi_{s,t}^j)\right]_*(\mu)$ for the maps $\pi_{s,t}^k := \pi_t^k - \pi_s^k$ and $\mathfrak{m}(a,b) := a{\cdot}b$ and $s,t \in \mathbb{I}$. Further, $\mathbb{E}\,\nu := \int_{\mathbb{R}^n} u\,\nu(\mathrm{d}u)$ is the expectation of a (possibly signed) Borel measure ν on $\mathbb{R}^n$, and a measure $\mu = (\mu_t) \in \mathcal{M}_1$ is called mean-stationary if $\mathbb{E}[\mu_t] = \mathbb{E}[\mu_0]$ for each $t \in \mathbb{I}$.

Finally, a signal $\mu \in \mathcal{M}_1$ is said to have (mutually) *independent components*, or to be *IC*, if $\mu = \mu^1 \otimes \cdots \otimes \mu^d$ (up to the identification $\mathcal{M}_1(\mathcal{C}_d) \cong \mathcal{M}_1(\mathcal{C}_1^{\times d})$; Lemma 2.1.2). The set of all IC signals in $\mathcal{M}_1$ is denoted $\mathscr{S}_\perp$.

Example 5.2.2. To illustrate, suppose that $\mu = \mathbb{P}_Y \in \mathcal{M}_1$ for a some signal $Y \equiv (Y^i) = (Y_t)$ in $\mathbb{R}^d$. Then $\mu^i = Y^i$ and $\mu_t = Y_t$ and $g(\mu) = g(Y) \equiv \mathbb{P}_{g(Y)}$ for any measurable $g : \mathbb{R}^d \to \mathbb{R}^d$. Further $\mathbb{E}\,\mu_{s,t} = \mathbb{E}Y_t - \mathbb{E}Y_s$ and $\mathbb{E}[\mu_{s,t}^i \mu_{s,t}^j] = \mathbb{E}[Y_{s,t}^i Y_{s,t}^j]$ for any $i,j \in [d]$ and $s,t \in \mathbb{I}$, and μ is mean-stationary iff the process Y is mean-stationary, and μ is IC iff the components $Y^1, \ldots, Y^d$ of Y are mutually statistically independent (as stochastic processes).

5.3 The Problem of Blind Source Separation

A well-founded stability analysis of Blind Source Separation requires a conceptionally precise formulation of the general problem statement (5.1) that we gave initially. Taking $S = (S_t)_{t\in\mathbb{I}}$ as a stochastic process in $\mathbb{R}^d$, $\mathbb{I} \subset \mathbb{R}$ compact, this section provides an according such formulation: Following a rigorous definition of the inverse problem (5.1) and its associated concept of solution (Section 5.3.1), a general notion for the stability of blind inversion is informally anticipated and turned into a natural and exact definition of robustness in the context of ICA (Section 5.3.2).

The terminology of this section, which at first reading may seem somewhat abstract to the general reader, is illustrated in Figure 5.1 and further motivated in Section 5.3.3.

5.3.1 Blind Source Separation

The problem of blind source separation is an inverse problem with the objective to recover an unobserved signal S from its mixture $X = f(S)$ under the condition that the effecting transformation f is invertible but unknown. The underlying relation between X and S is expressed by the identity

$$X_t \;=\; f(S_t) \quad \text{for all} \;\; t \in \mathbb{I}, \tag{5.7}$$

which is well-defined iff the domain of f includes the spatial support D_S of S. Assuming that f and its inverse are both continuous, we have $f(D_S) = D_X$ and call such

$$(X, S, f) \quad \text{a } \textit{BSS-triple} \tag{5.8}$$

with *observable* X*, source* S, and *mixing transformation* f. In partial summary of the above, note that the 'causal part' (S, f) of any such triple is an element of the space

$$\mathfrak{C} := \Big\{ (\tilde{S}, \tilde{f}) \in \mathcal{M}_1 \times Z^Z \;\mid\; \tilde{f} \in C^{0,0}(D_{\tilde{S}}) \Big\} \tag{5.9}$$

where we set $Z := \mathbb{R}^d$ and employed Convention 5.1.1.[4] A few more definitions will be useful.

5.3.1.1 The Problem of Blind Source Separation

Our blindness to the right-hand side of (5.7) imposes an asymmetry between the observable X and its cause (S, f): If we are only given X, bar any information on S or f, then we cannot distinguish the 'original' cause (S, f) underlying X from any other pair in the set

[4] Recall that by this convention we identify a signal $\tilde{S}$ with its law $\mathbb{P}_{\tilde{S}} \in \mathcal{M}_1$. Note further that in (5.9) we made the slight abuse of notation: $\tilde{f} \in C^{0,0}(D_{\tilde{S}}) \;:\Leftrightarrow\; \tilde{f}\big|_{D_{\tilde{S}}} \in C^{0,0}(D_{\tilde{S}})$, to not clutter our exposition.

$\{(\tilde{S},\tilde{f}) \mid (X,\tilde{S},\tilde{f})$ BSS-triple$\} \subseteq \mathfrak{C}$ of alternative causes. In this case, the original S cannot be recovered from X any better than finding an element of

$$[S]_X := \{\tilde{S} \in \mathcal{M}_1 \mid \exists\, \tilde{f} \in C^{0,0}(D_{\tilde{S}}) \,:\, (X,\tilde{S},\tilde{f}) \text{ is BSS-triple}\} = \mathrm{ev}_X(C^{0,0}(D_X)).$$

The BSS paradigm ameliorates this by upgrading our information from (nothing but) X to:

$$X \quad \text{and} \quad \mathscr{I} \quad \text{and} \quad \text{the } \textit{inclusion} \ \ (S,f) \overset{!}{\in} \mathscr{I} \tag{5.10}$$

for some a-priori given non-empty subset $\mathscr{I}$ of $\mathfrak{C}$, called an *identifiability assumption* (IA), which describes our prior knowledge on the original cause (S,f) of X. (The larger $\mathscr{I}$, the less we know about (S,f): if $\mathscr{I} = \{(S,f)\}$ then we know the original cause exactly, if $\mathscr{I} = \mathfrak{C}$ then we are 'completely blind' about it.)

The upgrade of the data from X to $(X,\mathscr{I})$ reduces the asymmetry between X and (S,f): Under (5.10), the source of a triple (X,S,f) can now be recovered up to within the smaller class

$$\langle X\rangle_{\mathscr{I}} := \{\tilde{S} \equiv \tilde{g}(X) \mid \tilde{g} \in C^{0,0}(D_X) \text{ and } (\tilde{g}(X),\tilde{g}^{-1}) \in \mathscr{I}\} \ \subseteq\ [S]_X. \tag{5.11}$$

In other words, each element in $\langle X\rangle_{\mathscr{I}}$ is a best-approximation of S under (5.10), and any two elements in $\langle X\rangle_{\mathscr{I}}$ are 'indistinguishable'[5] from S given $(X,\mathscr{I})$. Let us agree on terminology.

Definition 5.3.1 (BSS-triple on $\mathscr{I}$)**.** Let $\mathscr{I} \subseteq \mathfrak{C}$. A BSS-triple (X,S,f) with $(S,f) \in \mathscr{I}$ is called *a BSS-triple on* $\mathscr{I}$, in symbols: $(X,S,f)_{\mathscr{I}}$. The associated class $\langle X\rangle_{\mathscr{I}}$ in (5.11) is the *maximal solution* from X given $\mathscr{I}$, its elements are the *quasi sources* of X given $\mathscr{I}$.

A schematic illustration of these and the following definitions is given in Figure 5.1, while in Example 5.3.12 the notions of this section are fleshed out in the cocktail party context of p. 12.

Remark 5.3.2. (i) The maximal solution of a BSS-triple $(X,S,f)_{\mathscr{I}}$ is maximal wrt. set inclusion and (hence) unique by construction.

(ii) If the identifiability assumption $\mathscr{I}$ in (5.10) is given as a product set in $\mathcal{M}_1 \times C^{0,0}(\mathbb{R}^d)$, say $\mathscr{I} = \mathscr{S} \times \mathfrak{T}$, then the maximal solution $\langle X\rangle \equiv \langle X\rangle_{\mathscr{S}\times\mathfrak{T}}$ from X given $\mathscr{I}$ reads

$$\langle X\rangle = \mathrm{ev}_X(\mathfrak{T}) \cap \mathscr{S} \tag{5.12}$$

where $\mathrm{ev}_X(\mathfrak{T})$ is the image of the evaluation map $\mathrm{ev}_X : \mathfrak{T} \ni g \mapsto g(X)$, cf. (5.6).

[5] More precisely, this is captured by the equivalence relation $\big[\tilde{S} \sim_{(X,\mathscr{I})} \check{S} :\Leftrightarrow \tilde{S},\check{S} \in \langle X\rangle_{\mathscr{I}}\big]$ on $[S]_X$.

(iii) The [size of the] set $\langle X \rangle_{\mathscr{I}}$ depends on $\mathscr{I}$. In particular: the weaker its identifiability assumption (i.e. the larger the set $\mathscr{I}$), the larger is the class $\langle X \rangle_{\mathscr{I}}$ and hence the 'further away' the set of all quasi sources will be from the original source $S \cong \{S\}$. Accordingly, any S-containing subset $[S]_\circ \subseteq \langle X \rangle_{\mathscr{I}}$ [$\subseteq [S]^\circ$, resp.] defines an upper [resp. lower] bound to the 'accuracy' to which S can be recovered from X.

Central to the blind inversion task $X \mapsto S$ is the *accuracy* to which it shall be achieved — that is, which deviations [of the attained quasi sources] from the original source S are considered optimal or most tolerable? The answer is typically defined by the BSS-practitioner or their application context, and it translates to the constraint[6] $\langle X \rangle_{\mathscr{I}} \subseteq \lceil S \rceil$ for some a-priori given 'optimal' superset $\lceil S \rceil$ of S (the 'accuracy bound'). The primary difficulty, then, is to identify an IA $\mathscr{I}$ for which this inclusion holds, i.e. for which the set $\langle X \rangle_{\mathscr{I}}$ is identical to the optimum superset $[S]_\circ := \lceil S \rceil \cap \langle X \rangle_{\mathscr{I}}$ of S. In case of non-optimality, the deviation between $\langle X \rangle_{\mathscr{I}}$ and $[S]_\circ$ may be controlled by deriving a lower bound $[S]^\circ$ a posteriori to the choice of $\mathscr{I}$.

The above immediately suggests the following notions.

Definition 5.3.3 (Blind Inversion)**.** Given a map $\lceil \cdot \rceil : \mathcal{M}_1 \to 2^{\mathcal{M}_1}$, an IA $\mathscr{I} \subseteq \mathfrak{C}$ is called

$$\lceil \cdot \rceil\text{-}\textit{sufficient} \quad \text{if} \quad \langle X \rangle_{\mathscr{I}} \subseteq \lceil S \rceil \quad \text{for each BSS-triple } (X, S, f)_{\mathscr{I}}. \tag{5.13}$$

Further, for any $\mathscr{I} \subseteq \mathfrak{C}$ we call *inversion map on* $\mathscr{I}$ any function $\Phi : \mathcal{M}_1 \to 2^{\mathcal{M}_1}$ such that:

$$\emptyset \neq \Phi(X) \subseteq \langle X \rangle_{\mathscr{I}}, \quad \text{for each BSS-triple } (X, S, f)_{\mathscr{I}}. \tag{5.14}$$

Finally, a map $\varphi : \mathcal{M}_1 \to 2^{\mathcal{M}_1}$ shall be called *faithful* if $S \in \varphi(S)$ for each $S \in \mathcal{M}_1$.

The task of blind inversion now seeks to accomplish the following: For a given accuracy bound $\lceil \cdot \rceil$, can we find an identifiability assumption $\mathscr{I}$ (the weaker the better) which is strong enough to guarantee that for each cause $(S, f) \in \mathscr{I}$ the source S can be recovered from its mixture (5.7) up to an accuracy of $\lceil \cdot \rceil$, i.e. up to within an element of $\lceil S \rceil$? And if so, is there an explicitly computable algorithm that performs this inversion?

In formal terms, the (general) **Problem of Blind Source Separation** reads as follows:

	Given a faithful map $\lceil \cdot \rceil : \mathcal{M}_1 \to 2^{\mathcal{M}_1}$, determine:
$(\star)$	(a) a [weak] IA $\mathscr{I} \subseteq \mathfrak{C}$ that is $\lceil \cdot \rceil$-sufficient;
	(b) an explicit inversion map Φ on $\mathscr{I}$.

Any pair $(\mathscr{I}, \Phi)$ that satisfies (a) & (b) is called a *solution* to the BSS-problem for $\lceil \cdot \rceil$.

[6] So that each element $\tilde{S} \in \lceil S \rceil \cap \langle X \rangle_{\mathscr{I}}$ is an *optimal solution* to the such-constrained BSS problem.

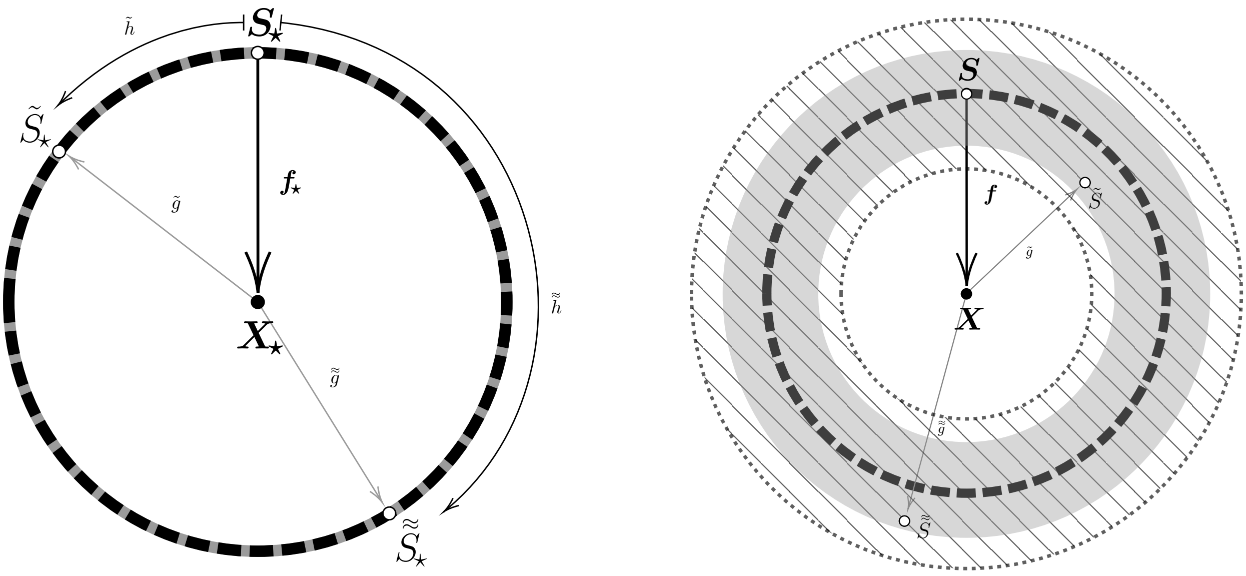

a. $[S_\star]_\circ = \langle X_\star \rangle_{\mathscr{I}_*} = \Phi(X_\star)$ b. $[S]_\circ \subsetneq \langle X \rangle_{\mathscr{I}}, \ \Phi(X) \neq [S]_\circ$

Figure 5.1: *Schematic illustration of two BSS-triples $(X_\star, S_\star, f_\star)_{\mathscr{I}_*}, (X, S, f)_{\mathscr{I}}$ and their respective maximal solution. Each point in the plane represents a signal, i.e. an element of $\mathcal{M}_1$, and the coloured areas visualise subsets of $\mathcal{M}_1$ that relate to the observables $X_\star, X$ or their sources $S_\star, S$. Depicted in this way are, for some accuracy-bound $\lceil \cdot \rceil$, the sets of accurate quasi sources $[S_\star]_\circ := \langle X_\star \rangle_{\mathscr{I}_*} \cap \lceil S_\star \rceil$ and $[S]_\circ := \langle X \rangle_{\mathscr{I}} \cap \lceil S \rceil$ (dashed black circles), the maximal solution $\langle X \rangle_{\mathscr{I}}$ (striped annulus with pointed boundary), as well as the outputs $\Phi(X_\star)$ and $\Phi(X)$ of some inversion map $\Phi \equiv \Phi_{\mathscr{I}_*}$ on $\mathscr{I}_*$ (grey annuluses). The IA $\mathscr{I}_*$ supporting $X_\star$ is $\lceil \cdot \rceil$-sufficient, which implies that the sets $\langle X_\star \rangle_{\mathscr{I}_*} = \Phi(X_\star)$ (i.e., the set of all $\mathscr{I}_*$-informed best guesses on $S_\star$ given $X_\star$) and $[S_\star]_\circ$ coincide, see panel a. In particular, each quasi source $\tilde{S}_\star \equiv \tilde{g}(X_\star)$ in $\langle X_\star \rangle_{\mathscr{I}_*}$ is an 'accurate' guess for $S_\star$ insofar as the residual $\tilde{h} = \tilde{g} \circ f_\star$ that relates $S_\star$ to $\tilde{S}_\star$ is in 'minimal distance to the identity' in a sense determined[7] by $\lceil \cdot \rceil$. In panel b., the IA $\mathscr{I}$ underlying X is too weak to be $\lceil \cdot \rceil$-sufficient and so the maximal solution $\langle X \rangle_{\mathscr{I}}$ is a proper superset of the optimum $[S]_\circ$. The* robustness *of $(\mathscr{I}_*, \Phi)$ asserts that the inversion property of Φ is resilient to deviations from $\mathscr{I}_*$ of the causes (in $\mathfrak{C}$) underlying its input signal. This amounts to a 'continuous transition from a. to b.', as captured by* (5.15)*: if, in some general sense, the causes $(S, f) \in \mathscr{I}$ and $(S_\star, f_\star) \in \mathscr{I}_*$ are 'not too different', can we provide quantitative guarantees that $\Phi(X)$ and $\lceil S \rceil$ are also not too far apart?*

Remark 5.3.4. As it stands, the above problem $(\star)$ does in fact comprise a whole family of inverse problems, even for a fixed $\lceil\cdot\rceil$, since the $\mathscr{I}$-qualifying adjective 'weak' in $(\star)_{(a)}$ is not explicitly specified.[8] This ambiguity can be avoided by demanding $\mathscr{I}$ to be the 'weakest' (i.e. globally maximal wrt. set inclusion) $\lceil\cdot\rceil$-sufficient IA in $\mathfrak{C}$, which would turn $(\star)_{(a)}$ into the stronger task of characterising $\lceil\cdot\rceil$-accurate invertibility. We refrain from this stronger formulation here to better capture the BSS-problem in the form in which it is currently discussed and approached in the literature, cf. e.g. [38, 81, 111] and the reference therein. ◆

Remark 5.3.5. The BSS-problem $(\star)$ consists of an abstract part, (a), which theoretically ensures the [sufficiently accurate] blind invertibility of an observable back for its cause, and a more procedural part, (b), which asks for an algorithm to perform this inversion in (both) theory and statistical practice. The documented solutions to $(\star)$ usually address the tasks (a) and (b) simultaneously, typically by first proving the $\lceil\cdot\rceil$-sufficiency (5.13) of an IA constructively and then deriving an explicit inversion map (5.14) in direct consequence to the employed strategy of proof, cf. [37, 38, 81, 111]. We also follow this general scheme in Section 5.6. ◆

In addition to the classical formulation $(\star)$, we propose to (provisionally) say that a solution

$$\begin{gathered}(\mathscr{I}, \Phi) \text{ is } \textbf{robust} \quad \text{if,} \quad \text{in some appropriate sense,} \quad \text{the maps} \\ (\tilde{S}, \tilde{f}) \mapsto \Phi(\tilde{f}(\tilde{S})) \;\text{ and }\; (\tilde{S}, \tilde{f}) \mapsto \operatorname{dist}_\varrho\Big(\Phi(\tilde{f}(\tilde{S})), \lceil\tilde{S}\rceil\Big) \quad \text{are continuous}^9 \text{ on } \mathscr{I}.\end{gathered} \tag{5.15}$$

(See [75] for the general notion of robustness that we follow.) Above uses the 'asymmetric distance'

$$\operatorname{dist}_\varrho(\mathcal{A}, \mathcal{B}) := \sup\nolimits_{a\in\mathcal{A}} \mathrm{d}_\varrho(a, \mathcal{B}) \quad \text{with} \quad \mathrm{d}_\varrho(a, \mathcal{B}) := \inf\nolimits_{b\in\mathcal{B}} \varrho(a, b) \tag{5.16}$$

where $\varrho : \mathcal{M}_1^{\times 2} \to \bar{\mathbb{R}}_+$, with $\varrho \circ (\mathrm{id} \times \mathrm{id}) = 0$, is some appropriate distance function between the outputs of Φ and the desired signals in $\lceil\cdot\rceil$; specific choices of ϱ will depend on $\lceil\cdot\rceil$.

Remark 5.3.6. From a purely conceptional angle, the (anticipated) notion of robustness (5.15) adds to the conventional analysis of problem $(\star)$ the classical aspects of (numerical and analytical) well-conditionedness and stability, which play a central role throughout many areas of pure and applied mathematics, see e.g. [63, 71, 75], [6, p. 139], or [57, Sect. IV.12.3].

---

[7] For instance, the popular choice (5.26) of $\lceil S\rceil$ as the monomial orbit $\mathrm{M}_d \cdot S$ mandates $\tilde{h} \in \mathrm{M}_d$.

[8] E.g., while $(\lceil S\rceil, \mathscr{I}) = \big(\{S\}, \{(S, f)\}\big)$ is trivial, the choice $(\lceil S\rceil, \mathscr{I}) = \big([S]_X, \mathfrak{C}\big)$ is meaningless; solutions to $(\star)_{(a)}$ that are of practical relevance or theoretical interest will lie strictly between these two extremes.

[9] As usual, a map Φ is called *continuous on* $\mathscr{I}$ (or $\mathscr{I}$*-continuous*) if Φ is pointwise continuous on $\mathscr{I}$.

5.3.2 Independent Component Analysis

As indicated above, the preliminary description of BSS-robustness (5.15) is only still an informal desideratum rather than an exact definition; for the latter, both ϱ and a suitable topology ('$\mathscr{I}$-continuity') need to be specified further. This gap is closed in the present subsection which provides an according completion of (5.15) for the most widely received special instance of ($\star$) known as Independent Component Analysis (ICA).

To prepare, denote by $\lceil\cdot\rceil_{\mathfrak{m}} : \mathcal{M}_1 \to 2^{\mathcal{M}_1}$ the function that sends a signal S to its orbit $\lceil S\rceil_{\mathfrak{m}} := \mathrm{M}_d \cdot S$ under the monomial group (action) $\mathrm{M}_d := \{\Lambda P \mid \Lambda \in \Delta_d, P \in \mathrm{P}_d\}$ on $\mathcal{M}_1$ (cf. (5.26)). Let similarly $\tilde{\mathrm{M}}_d := \{h : u \mapsto Mu + b \mid M \in \mathrm{M}_d,\, b \in \mathbb{R}^d\}$ be the subgroup [of $C^{1,1}(\mathbb{R}^d)$] of affinely monomial transformations, with $\lceil S\rceil_{\tilde{\mathfrak{m}}} := \tilde{\mathrm{M}}_d \cdot S$. Write also $\mathfrak{R} := \{g : \mathbb{R}^d \to \mathbb{R}^d \,|\, g \text{ is Borel}\}$ for the set of Borel measurable transformations on $\mathbb{R}^d$.

Below gives a refinement of (5.14), see (5.19). Note for this that if $\mathscr{I}\subseteq\mathfrak{C}$ is $\lceil\cdot\rceil_{\mathfrak{m}}$-sufficient, then for any inversion map Φ on $\mathscr{I}$ there is $\hat{\Phi} : \mathcal{M}_1 \to 2^{\mathfrak{R}}$, $X \mapsto \{B \in \mathfrak{R} \mid B \cdot X \in \Phi(X)\}$, with

$$\Phi(X) \;=\; \hat{\Phi}(X) \cdot X, \quad \text{for each BSS-triple } (X, S, f)_{\mathscr{I}}. \tag{5.17}$$

Among the inverse problems subsumed under ($\star$), the following variant is by far most established one to date, cf. e.g. [38, 80, 81, 111]; the reference set $\hat{\mathscr{S}}_\perp$ is defined in (5.27).

Definition 5.3.7 (Independent Component Analysis)**.** A set $\mathscr{I} \subseteq \mathfrak{C}$ is called an *ICA-condition* if

$$\mathscr{I} = \mathscr{S} \times \mathrm{GL}_d \quad \text{for some} \quad \mathscr{S} \supseteq \hat{\mathscr{S}}_\perp \quad \text{and} \quad \mathscr{I} \quad \text{is} \quad \lceil\cdot\rceil_{\mathfrak{m}}\text{-sufficient;} \tag{5.18}$$

a BSS-triple $(X, S, f)_{\mathscr{I}}$ is called an *ICA-triple* (*on* $\mathscr{I}$) if $\mathscr{I}$ is an ICA-condition, and in this case we call an *ICA-inversion on* $\mathscr{I}$ any function

$$\hat{\Phi} : \mathcal{M}_1 \to 2^{\mathfrak{R}} \quad \text{such that:} \quad \emptyset \neq \hat{\Phi}(f(S)) \circ f \;\subseteq\; \tilde{\mathrm{M}}_d, \quad \text{for each } (S, f) \in \mathscr{I}. \tag{5.19}$$

The Problem of *Independent Component Analysis (ICA)* is ($\star$) for $\lceil\cdot\rceil \equiv \lceil\cdot\rceil_{\tilde{\mathfrak{m}}}$ and with $\mathscr{I}$ and Φ of the form (5.18) and (5.19); call an according pair $(\mathscr{I}, \hat{\Phi})$ a *solution* to the ICA problem.

Remark 5.3.8. In other words, an ICA-triple $(X, S, f)_{\mathscr{I}}$ is a BSS-triple where the source S is to be recovered up to some permutation and scaling (and perhaps a constant offset) of its original components and where the prior information $\mathscr{I} \ni (S, f)$ for this consists of a condition on the source ($\mathscr{S}$: regularity of each source component and their statistical independence) and a linearity condition on the mixing (GL_d) which are imposed separately via (5.10). Note further that if $\hat{\Phi}$ is an ICA-inversion of the form (5.19) then $X \mapsto \hat{\Phi}(X) \cdot X$ is an inversion map on $\mathscr{I}$ in the sense of Definition 5.3.3, as is clear from (5.14) and (5.12) and provided that $\tilde{\mathrm{M}}_d \cdot \mathscr{S} \subseteq \mathscr{S}$. ◆

The specificity (5.18) and (5.19) of ICA enables a natural analytical interpretation of the generic robustness notion (5.15). For this, let $(\mathscr{I}, \hat{\Phi})$ be any fixed ICA-solution, set $\Phi : X \mapsto \hat{\Phi}(X) \cdot X$, and let us quantify the (deviations from) the accuracy condition (5.19) via the 'deviance function'

$$\partial_{\hat{\Phi}} : \mathfrak{C} \to \mathbb{R}_+ \quad \text{given by} \quad \partial_{\hat{\Phi}}(S, f) := \sup_{B \in \hat{\Phi}(f(S))} \inf_{(M,v) \in \mathrm{M}_d \times \mathbb{R}^d} \varrho_B((S,f),(M,v))$$
$$\text{for} \quad \varrho_B((S,f),(M,v)) := \sup_{u \in D_S^{(v)}} \tfrac{|B \circ f(u-v) - Mu|}{|Mu|} \mathbb{1}_{\times}(Mu), \tag{5.20}$$

with $\mathbb{1}_{\times} := \mathbb{1}_{\mathbb{R}^d \setminus \{0\}}$ and $D_S^{(v)} := v + D_S$. For inequality (5.21) below, recall notation (5.16).

Lemma 5.3.9. *The map $\partial_{\hat{\Phi}}$ vanishes on $\mathscr{I}$, i.e. $\partial_{\hat{\Phi}}\big|_{\mathscr{I}} = 0$, and for any function f and every stochastic process $S = (S_t)$ with $(S, f) \in \mathfrak{C}$ we have that, with probability one,*

$$\mathrm{dist}_{\tilde{\varrho}}(\Phi(f(S)), \lceil S \rceil_{\tilde{\mathfrak{m}}}) \leq \partial_{\hat{\Phi}}(S, f) \quad \textit{for} \quad \tilde{\varrho}(X, Y) := \sup_{t \in \mathbb{I}} \tfrac{|X_t - Y_t|}{|Y_t|} \mathbb{1}_{\times}(Y_t). \tag{5.21}$$

Proof. That $\partial_{\hat{\Phi}}(\mathscr{I}) \equiv 0$ is clear from (5.19). Write $\varphi(S, f)$ for the left-hand side of the inequality (5.21) and notice our slight abuse of notation, cf. Convention 5.1.1: here, $\varphi(S, f)$ is a function of the process $S = (S_t)$ in $\mathbb{R}^d$, while $\partial_{\hat{\Phi}}(S, f)$ is a function of the law, $\mathbb{P}_S \in \mathcal{M}_1$, of S. By (5.16),

$$\varphi(S, f) = \sup_{\tilde{X} \in \Phi(X)} \inf_{h \in \tilde{\mathrm{M}}_d} \tilde{\varrho}(\tilde{X}, h(S)) \quad \text{for} \quad X = f(S), \tag{5.22}$$

where $\Phi(X) = (\hat{\Phi}(X) \circ f)(S) \equiv \{(B \circ f)(S) \mid B \in \hat{\Phi}(X)\}$ by definition. Hence for any $\tilde{X} \in \Phi(X)$ and all $h \equiv M + v \in \tilde{\mathrm{M}}_d$ (some $M \in \mathrm{M}_d$, $v \in \mathbb{R}^d$) there is $B \equiv B(\tilde{X}) \in \hat{\Phi}(X)$ with

$$\tilde{\varrho}(\tilde{X}, h(S)) = \sup_{t \in \mathbb{I}} \tfrac{|(B \circ f)(S_t) - (MS_t + v)|}{|MS_t + v|} \mathbb{1}_{\times}(h(S_t)) = \sup_{t \in \mathbb{I}} \tfrac{|(B \circ f)(\tilde{S}_t - \tilde{v}) - M\tilde{S}_t|}{|M\tilde{S}_t|} \mathbb{1}_{\times}(M\tilde{S}_t)$$

where $\tilde{S} = (\tilde{S}_t)$ is given by $\tilde{S}_t := S_t + \tilde{v}$ with $\tilde{v} := M^{-1}v$. Hence and since the trace of a process is almost surely contained in its spatial support, we have that, with probability one,

$$\tilde{\varrho}(\tilde{X}, h(S)) \leq \sup_{u \in D_{\tilde{S}}} \tfrac{|(B \circ f)(u - \tilde{v}) - Mu|}{|Mu|} \mathbb{1}_{\times}(Mu) = \varrho_B((S,f),(M,\tilde{v})) \tag{5.23}$$

where $D_{\tilde{S}} = D_S^{(\tilde{v})}$. Statement (5.21) now follows from (5.23) and the definitions (5.22) and (5.20). □

With the relative error between the output of Φ and the desired accuracy $\lceil\cdot\rceil_{\tilde{\mathfrak{m}}}$ controlled by $\partial_{\hat{\Phi}}$, we can now deliver the anticipated formalisation of (5.15) for the special case of ICA. Notice that the following is a natural combination of the BSS-specific stability requirements (5.15) with the classical notion of statistical robustness as set out in [75, Sects. 1.2–1.4, 2.6].

Definition 5.3.10 (Robust ICA)**.** A solution $(\mathscr{I}, \hat{\Phi})$ to the ICA-problem is *robust* if

$$\partial_{\hat{\Phi}} : \mathfrak{C} \to \mathbb{R}_+ \quad \text{is continuous on } \mathscr{I} \tag{5.24}$$

with respect to a sufficiently coarse topology on the causal space $\mathfrak{C}$.

In addition to the qualitative assertion (5.24) it is often desirable to also arrive at an explicit *quantitative analysis* of this robustness (cf. e.g. [75, Sect. 1.3f.]). Provided the underlying topology on $\mathfrak{C}$ is sufficiently informative, this can be achieved by complementing (5.24) with a modulus of continuity of $\partial_{\hat{\Phi}}$ at $\mathscr{I}$, as done in Section 5.6 below.

Remark 5.3.11 ('Sufficiently Coarse')**.** What do we mean by the definiens 'sufficiently coarse'? Just as for other uses of the robustness concept [67, 75], the answer to 'how coarse is coarse enough' (cf. Rem. 5.3.3.2) typically depends on the specific (application) context in which the 'stability of Φ under violations of $\mathscr{I}$' is of interest; a one-size-fits-all topology – or universal level of coarseness – to optimally capture all thinkable violations is generally difficult to pinpoint. However, taking classical theory as a rough guide, see e.g. [75, Sects. 1.3, 2.6], we can hope to capture 'most' relevant violations by, for reference, calling a topology $\tau_{\mathfrak{C}}$ on $\mathfrak{C}$

$$\textit{sufficiently coarse} \quad \text{if the } (\tau_{\mathfrak{C}}, \tilde{\pi}_1)\text{-induced final topology on } \mathcal{M}_1 \text{ is coarser than } \tau^\star, \tag{5.25}$$

where $\tilde{\pi}_1 : \mathfrak{C} \to \mathcal{M}_1$, $(\tilde{S}, \tilde{f}) \mapsto \tilde{S}$, is the canonical projection from a cause to its signal and $\tau^\star$ is the topology of weak convergence on $\mathcal{M}_1(\mathcal{C}_d)$ (e.g. [18, Def. 8.1.1] and [16, Sect. 7]). Again, we emphasize that the specification (5.25) is merely a reference point to provide orientation and should not be read as a necessary or strictly binding constraint on (5.24); we do not restrict Definition 5.3.10 to topologies of only this kind. In fact, the following sections will establish (5.24) for a topology $\tau_{\mathfrak{C}}$ on $\mathfrak{C}$ which, while conforming to (5.25) in the setting of most applications, is generally slightly finer than (5.25) but nonetheless still sufficiently coarse to informatively capture (the 'convergence to exactness' of) many infringements of an IA that are of relevance in theory and practice, cf. e.g. Example 1.1.2 and Section 5.7.

5.3.3 Complementary Motivation

Let us illustrate the notions and definitions of the previous sections by an appeal to Example 1.1.2.

Example 5.3.12 (Cocktail Party Revisited)**.** The original speech signals pertaining to the speakers in Example 1.1.2 can be modelled as [drawn from] the components S^i of a stochastic process $S = (S^1, \ldots, S^d) \in \mathcal{M}_1$. The microphone recordings of these signals are then mixtures of the form $X^i = f_i(S^1, \ldots, S^d)$ that combine to an observable $X = (X^1, \ldots, X^d) = f(S)$ for some invertible map $f = (f_1, \ldots, f_d) : D_S \to \mathbb{R}^d$. The inverse problem from Example 1.1.2 thus translates into the BSS-problem ($\star$) for the triple (X, S, f), and is to be solved to an accuracy $\lceil \cdot \rceil$ not explicitly specified in the example. The example states, however, that the goal is to "recover from X what each of the d speakers have said". Thus, the goal is *not* necessarily an exact recovery[10] of S from X. Instead, the objective is to recover from X only the messages – that is, the informational content of their respective speeches – that the speakers communicated: which speaker it was that said something, or how loudly they originally said it, is irrelevant.

This abstracts from the original source S the order of its components [i.e., the attribution (via enumeration) of a speech signal to its original speaker] and the scale of its components [i.e., how loudly such a speech was originally communicated]. Consequently, any two signals $(b_1 S^{j_1}, \ldots, b_d S^{j_d})$ and $(\hat{b}_1 S^{k_1}, \ldots, \hat{b}_d S^{k_d})$ with $b_i, \hat{b}_i \neq 0$ and $\{j_i\} = \{k_i\} = [d]$ are both equally desirable outcomes of the blind inversion. I.o.w., the task in Example 1.1.2 is 'accurately' accomplished if from X we arrive at (not necessarily S but instead) any signal of the form

$$\tilde{S} \in \big\{ (\alpha_1 S^{\tau(1)}, \ldots, \alpha_d S^{\tau(d)}) \,\big|\, \alpha_1, \ldots, \alpha_d \in \mathbb{R}_\times, \, \tau \in S_d \big\} = \mathrm{M}_d \cdot S, \tag{5.26}$$

where $\mathrm{M}_d := \{\Lambda P \mid \Lambda \in \Delta_d, P \in \mathrm{P}_d\}$ is the subgroup [of GL_d] of monomial matrices. This establishes the monomial orbit $\lceil S \rceil := \mathrm{M}_d \cdot S$ as a suitable task-specific accuracy (5.13) for Ex. 1.1.2.

Since our prior knowledge about the speakers, S, and about 'the physics' of their recording situation, f, is very limited, we can generally pinpoint the original constellation (S, f) only as an element of some subset $\mathscr{I}$ in $\mathfrak{C}$ and not as a singleton in said space – this is our 'blindness'. Consequently, at this prior level, every signal that conforms to $\mathscr{I}$ (that is, every $\hat{S} := g(X)$ s.t. $(\hat{S}, g^{-1}) \in \mathscr{I}$) appears to us indistinguishable from the original speeches we

[10] ... i.e., to obtain from X not only the content of *what was said* but also *which speaker* (identified by their number i) it was that said it and *at which original sound volume* they said it; this is usually impossible.

are looking for. We thus dub signals of this kind 'quasi sources' and subsume them into the set (5.11).

The main question for us at this point is whether our limited knowledge $\mathscr{I}$ of the speakers and their mixing is sufficient to, at least in principle, recover from X the signal S up to the desired accuracy (5.26). This translates into the requirement $\langle X\rangle_{\mathscr{I}} \subseteq \lceil S\rceil$ on $\mathscr{I}$, and if this requirement is met we want to perform the inversion $X \mapsto S$ accordingly. The latter asks for an explicitly computable inversion procedure $\Phi : \mathcal{M}_1 \to 2^{\mathcal{M}_1}$ which, since it should make optimal use of the data $(X, \mathscr{I})$, we naturally expect to satisfy $\Phi(X) \neq \emptyset$ and $\Phi(X) \subseteq \langle X\rangle_{\mathscr{I}}$.

Now due to complex interdependencies between the speakers $S^1, \ldots, S^d$ or since the recording situation f is somewhat complicated or both, we might fail to formulate an assumption $\mathscr{I}$ on (S, f) for which the required accuracy condition (5.13) is met. We may, however, often find that (S, f) is still more or less 'well approximated' by an idealised constellation $(S_\star, f_\star)$ that conforms to an assumption $\mathscr{I}_*$ in $\mathfrak{C}$ for which the accuracy condition is met. In our example, a viable such approximation of (S, f) might be to choose $S_\star$ to have the same component signals as S but with no statistical dependence between them, and to take $f_\star$ as the linearisation of f at some point $x_0 \in D_S$, in symbols: $(S_\star, f_\star) \equiv (\mathbb{P}_{S^1} \otimes \cdots \otimes \mathbb{P}_{S^d}, D_f(x_0))$ (cf. Section 5.3.2).

This idealisation then pays off as follows. The $\lceil \cdot \rceil$-sufficiency of $\mathscr{I}_*$ allows for a map $\Phi \equiv \Phi_{\mathscr{I}_*}$ of the form (5.14) that achieves an exact blind inversion on $\mathscr{I}_*$-caused observables. Provided that the inversion property $\Phi(\tilde{f}(\tilde{S})) \subseteq \lceil \tilde{S} \rceil$ of this map is *stable* around the idealisation $(\tilde{S}, \tilde{f}) \equiv (S_\star, f_\star) \in \mathscr{I}_*$ as per (5.15), the proximity $(S, f) \approx (S_\star, f_\star)$ will then guarantee that $\Phi(X)$ is approximately a subset of $\lceil S \rceil$, as desired. More specifically, given an estimate on the 'deviation' Δ between (S, f) (the original cause) and $(S_\star, f_\star)$ (the $\mathscr{I}_*$-idealisation of that cause), we can provide results asserting, in terms of precise and informative error bounds with entirely explicit constants, that each element in $\Phi(X)$ equals an element of $\lceil S \rceil$ up to a relative error of order $\mathcal{O}(\Delta)$. The merits this has for our example application are clear: While the actual cause (S, f) of X may be too complex to blind-recover S exactly, we may find a 'structurally simpler approximation' to (S, f), say $\mathscr{I}_* \ni (S_\star, f_\star)$ with inversion Φ on $\mathscr{I}_*$ as above, and obtain that the estimated speeches $\tilde{S}^\Delta \equiv \Phi(X)$ will still be $\lceil \cdot \rceil$-accurate up to a controlled relative distance to (5.26). Provided that the idealisation error Δ is small enough,[11] the obtained signal $\tilde{S}^\Delta$ may show acoustic distortions but still be close enough to the optimum (5.26) (the clearly audible speeches) so as to recognize from it the original information that the speakers communicated. ◆

[11] Which one can try to achieve, e.g., by adjusting the microphone positions or attenuating external noise.

While the above perspectives and ideas on the robustness of BSS are relevant and intuitive, the current literature does not seem to offer any flexible and informative theoretical framework to formalise and exploit them. The approach that we develop in this chapter is intended as a first proposal in this direction.

5.3.3.1 Supplement to Definition 5.3.7

For its brief use in (5.18), we define the set ('smoothness class')

$$\hat{\mathscr{S}}_{\perp} := \{\mu \in \mathscr{S}_{\perp} \mid \sharp\{i \in [d] \mid \mathbb{E}\big[(\mu^{i}_{0,1})^{3}\big] = 0\} \leq 1 \text{ and } \sharp\,\sigma(\mathrm{Cov}(\mu_0, \mu_1)) = d\} \tag{5.27}$$

of all IC signals whose incremental covariance has pairwise distinct eigenvalues and for which at most one of their third incremental moments vanishes. These signals are of 'conventional regularity' in the sense that they are in particular non-Gaussian, hence (linearly) identifiable by the classical paradigm [37], and also of a sufficiently rich time-structure to be recoverable via the methods in [116]. I.o.w., (5.27) is a class of sources which are so regular ('smooth') that a blind inversion procedure should be able to recover them in order to be called an ICA-method.

5.3.3.2 On Definition 5.3.10

Intuitively speaking, any topology on $\mathfrak{C}$ for which condition (5.24) holds is a model for how 'sensitively' the $\lceil\,\cdot\,\rceil_{\tilde{\mathfrak{m}}}$-accuracy of Φ reacts to violations of the assumption $\mathscr{I}$: the coarser such a topology [i.e., the fewer open sets it contains], the less sensitively ("the more robustly") does Φ vary under general perturbations of $\mathscr{I}$. Since the continuity requirement (5.24) is redundant if understood wrt. the discrete, i.e. the finest, topology on $\mathfrak{C}$, a certain level of coarseness of the (5.24)-supporting topology on $\mathfrak{C}$ is clearly necessary. There is, however, no immediate or canonical choice for such a topology on $\mathfrak{C}$ so that setting a threshold on 'how coarse is coarse enough' is notoriously a matter of judgement, as further discussed in Remark 5.3.11. ◆

With an operable formulation of both the BSS-problem (5.1) – approached via ICA (5.18) – and its associated notion of robustness (5.15) at hand (see Definitions 5.3.7 and 5.3.10), the overarching goal for the rest of this chapter is to bring the above framework to life and fruition: Section 5.5 provides a natural, easy-to-handle and informative (premetric) coarse topology on $\mathfrak{C}$ that supports the stability (5.15) for a generically derived ICA-solution $(\mathscr{I}, \hat{\Phi})$ (Section 5.6) and allows the resulting BSS-robustness (5.24) of this solution to be thoroughly quantified via explicit moduli of continuity (Theorems 5.6.10 and 5.6.14). A number of application-relevant corollaries to this robustness illustrate the practical significance of the theory (Section 5.7).

5.4 Coordinatisable Signals and Their Convergences

Foundational to the announced topology on $\mathfrak{C}$, and thus to our above-mentioned robustness program altogether, is a distinguished set of 'coordinates' to describe the signals in $\mathcal{M}_1$. This leads to a structured and explicitly computable quantification of $\mathcal{M}_1$ which, alongside the topologies that this coordinatisation both captures and imposes, is studied in the present section.

The announced coordinates take the form of moment-like integral statistics and are defined by iteratively integrating the sample paths of a signal against each other. The latter to be well-defined requires a certain regularity of the realisations of a signal, which is accounted for by certain 'regularity classes' within $\mathcal{C}_d$ that the samples of a coordinatisable signal are assumed to be contained in. An expressive and convenient such class[12] is the regular subspace

$$\mathcal{C}^1 \equiv \mathcal{C}^1_d := \left\{x : [0,1] \to \mathbb{R}^d \;\middle|\; x \text{ is absolutely continuous}^{13}\right\} \tag{5.28}$$

of $\mathcal{C}_d$, and the signals whose samples are contained in this class are the elements of the subspace

$$\mathfrak{D} := \{\mu \in \mathcal{M}_1(\mathcal{C}_d) \mid \mu(\mathcal{C}^1) = 1\} \tag{5.29}$$

of $\mathcal{M}_1$, i.e. the set of all Borel probability measures on $(\mathcal{C}_d, \|\cdot\|_\infty)$ supported on $\mathcal{C}^1$. Writing $C^{1,1}(\mathbb{R}^d) =: C^{1,1}$ for the group of all C^1-diffeomorphisms on $\mathbb{R}^d$, note that (by the chain rule, cf. (5.30)) the space $\mathfrak{D}$ is left invariant under the $C^{1,1}$-action (5.6) on $\mathcal{M}_1(\mathcal{C}_d)$.

After establishing a few essential structural properties of the space $\mathcal{C}^1$ and exposing how it enables the inclusion of discrete-time signals (Sect. 5.4.1), we are ready to introduce the announced signal coordinates on $\mathfrak{D}$ (Sect. 5.4.2) and study how sensitively some natural topologies on $\mathcal{M}_1$ (Sect. 5.4.3) are captured by these coordinates (Sect. 5.4.4).

5.4.1 Properties of the Sample Space

The path space $\mathcal{C}^1$ enjoys a few topological regularity properties that make it a convenient choice to act as the sample space of our signals. We note, however, that the main ideas of this chapter can also be readily extended to signals with much less regular sample paths, see Remark 5.4.9 (and the references therein) at the end of this subsection. The contents of the following lemma are well-known.

[12] Though certainly not the largest possible regularity class, see Remark 5.4.9.
[13] To recall the definition of absolute continuity, see e.g. [54, Definition 1.17].

Lemma 5.4.1. *The above space $\mathcal{C}^1$ of absolutely continuous paths in $\mathbb{R}^d$ can be written as*

$$\mathcal{C}^1 = \left\{ x \in \mathcal{C}_d \;\middle|\; \exists\,!\, \dot{x} \in L^1([0,1];\mathbb{R}^d) \,:\, x = x_0 + \int_0^{\cdot} \dot{x}_s \,\mathrm{d}s \right\}; \tag{5.30}$$

this space is a Borel subset of $(\mathcal{C}_d, \|\cdot\|_\infty)$ and a separable Banach space with respect to the 1-variation norm

$$\|x\|_{1\text{-var}} \;=\; |x_0| \;+\; \|\dot{x}\|_{L^1}. \tag{5.31}$$

Proof. The identity (5.30) holds by [54, Propositions 1.31 & 1.32]. In fact, [54, Prop. 1.31] asserts that for $\mathfrak{X} := \mathbb{R}^d \times L^1([0,1];\mathbb{R}^d)$ and $\mathfrak{Y} := \mathcal{C}_d$, the map $f : \mathfrak{X} \to \mathfrak{Y}$ given by $f(c,v) := c + \int_0^{\cdot} v_s \,\mathrm{d}s$ is a Banach space isomorphism (which also proves (5.31)). From this, [90, Theorem 15.1] implies that the image $f(\mathfrak{X}) = \mathcal{C}^1$ is a Borel subset of $(\mathcal{C}^1, \|\cdot\|_\infty)$, i.e. that $\mathcal{C}^1 \in \mathcal{B}(\mathcal{C}_d)$. That $(\mathcal{C}^1, \|\cdot\|_{\text{-var}})$ is separable and Banach is stated as [54, Corollary 1.35]. □

Remark 5.4.2. Recall that $\|\cdot\|_{1\text{-var}}$ in fact dominates a whole family $(\|\cdot\|_{p\text{-var}})_{p\in[1,\infty]}$ of p-variation norms on $\mathcal{C}^1$ which, in generalisation of (5.31) (cf. [54, Sect. 5.1]), are given by

$$\|x\|_{p\text{-var}} := |x_0| + \left[\sup_{(t_\nu)} \sum_\nu |x_{t_{\nu+1}} - x_{t_\nu}|^p \right]^{1/p} \qquad \text{for} \quad 1 \le p < \infty \tag{5.32}$$

and $\|\cdot\|_{\infty\text{-var}} := \|\cdot\|_\infty$, where the above supremum runs over all (finite) dissections (t_ν) of $[0,1]$. Since these norms become weaker as p increases, that is $\|\cdot\|_{p\text{-var}} \ge \|\cdot\|_{q\text{-var}}$ for each $1 \le p \le q \le \infty$ (e.g. [54, Prop. 5.3] and [104, Lem. 1.6.4.]), we have the inclusions of subspaces

$$C_b(\mathcal{C}^1; \|\cdot\|_{1\text{-var}}) \supseteq C_b(\mathcal{C}^1; \|\cdot\|_{p\text{-var}}) \supseteq C_b(\mathcal{C}^1; \|\cdot\|_{q\text{-var}}) \qquad \text{for} \quad 1 \le p \le q \le \infty. \tag{5.33}$$

It is useful to observe that while the uniform topology on $\mathcal{C}^1$ is weaker than its 1-variation topology, both induce the same measurable structure on $\mathcal{C}^1$. This follows essentially by definition (5.31) of the 1-variation norm and the above fact that $\mathcal{C}^1$ is separable wrt. $\|\cdot\|_{1\text{-var}}$.

(Further topological properties of the spaces (5.33) are discussed in [54, Sect. 5.3.3].)

Due to its separability, the space $(\mathcal{C}^1, \|\cdot\|_{1\text{-var}})$ will allow us to compare the classical (i.e., $\|\cdot\|_\infty$-induced) weak topology of signals to a spectrum of finer ($\|\cdot\|_{p\text{-var}}$-induced) signal topologies that are fine enough to support the continuity of certain convenient coordinates on $\mathcal{M}_1$ (Section 5.4.3). A cornerstone for this comparison is the following lemma.

Lemma 5.4.3. *The spaces $(\mathcal{C}^1, \|\cdot\|_\infty)$ and $(\mathcal{C}^1, \|\cdot\|_{1\text{-var}})$ have the same Borel σ-algebra.*

Proof. Note first that since $\|\cdot\|_{1\text{-var}} \geq \|\cdot\|_\infty$ (which is easy to see) we find that the 1-variation topology on $\mathcal{C}^1$ is finer than the uniform topology on $\mathcal{C}^1$, which of course implies that $\mathcal{B}_{1\text{-var}} := \sigma(\mathcal{C}^1, \|\cdot\|_{1\text{-var}}) \supseteq \sigma(\mathcal{C}^1, \|\cdot\|_\infty) =: \mathcal{B}_\infty$. Since the separability of $(\mathcal{C}^1, \|\cdot\|_{1\text{-var}})$ guarantees that the σ-algebra $\mathcal{B}_{1\text{-var}}$ is generated by the closed $\|\cdot\|_{1\text{-var}}$-balls, the converse inclusion $\mathcal{B}_{1\text{-var}} \subseteq \mathcal{B}_\infty$ follows if we can show that

$$B^1_r(x) := \{y \in \mathcal{C}^1 \mid \|y - x\|_{1\text{-var}} \leq r\} \in \mathcal{B}_\infty \quad \text{for every} \quad x \in \mathcal{C}^1 \text{ and any } r \geq 0. \tag{5.34}$$

To see that this holds, fix any $x \in \mathcal{C}^1$ and $r \geq 0$ and recall that, by definition,

$$\|z\|_{1\text{-var}} = \sup_{\mathcal{I} \in \mathfrak{I}} V_{\mathcal{I}}(z) \quad \text{with} \quad V_{(t_\nu)}(z) := |z_0| + \sum_\nu |z_{t_{\nu+1}} - z_{t_\nu}|$$

and where $\mathfrak{I} := \{\mathcal{I} = (t_\nu) \mid \mathcal{I} \text{ is a (finite) dissection of } [0,1]\}$. Given any $\mathcal{I} \in \mathfrak{I}$ it is clear that the function $Q_{\mathcal{I}} : \mathcal{C}^1 \ni y \mapsto V_{\mathcal{I}}(y - x)$ is continuous wrt. $\|\cdot\|_\infty$, whence the level set $C_{\mathcal{I}} := \{y \in \mathcal{C}^1 \mid Q_{\mathcal{I}}(y) \leq r\}$ is $\|\cdot\|_\infty$-closed. Combined with this, the immediate identity

$$B^1_r(x) = \bigcap\nolimits_{\mathcal{I} \in \mathfrak{I}} C_{\mathcal{I}} \quad \text{implies that} \quad B^1_r(x) \text{ is closed wrt. } \|\cdot\|_\infty\,,$$

which shows that (5.34) holds as desired. □

Remark 5.4.4 (Discrete Signals). Sometimes observables are modelled a priori as random vectors or discrete time series in $\mathbb{R}^d$ rather than as continuous-time stochastic processes. Such discrete-time models are still special cases of our formalism, i.e. can be seen as signals in $\mathfrak{D}$, as per the identifications (for $Z := \mathbb{R}^d$ and $\hat{\mathcal{C}}_0 := \{x \in \mathcal{C}_{(0,1)} \mid x_0 = 0\}$)

$$\hat{\iota}_1 \,:\, \mathbb{R}^d \ni u \overset{\sim}{\longmapsto} (t \cdot u)_{t \in [0,1]} \in \hat{\mathcal{C}}_0 \quad \text{and} \quad \hat{\iota}_n \,:\, Z^{\times n} \overset{\sim}{\longrightarrow} \hat{\mathcal{C}}_{\mathcal{E}_n} \quad (n \geq 2), \tag{5.35}$$

where for a dissection $\mathcal{I} = (t_\nu \mid 0 = t_0 < \ldots < t_{n-1} = 1)$ of $[0,1]$ we set the subspace (of $\mathcal{C}^1$)

$$\hat{\mathcal{C}}_{\mathcal{I}} := \Big\{x = (x_t)_{t \in [0,1]} \;\Big|\; x_t = x_{t_\nu} + \tfrac{t_\nu - t}{t_{\nu+1} - t_\nu}(x_{t_{\nu+1}} - x_{t_\nu}) \text{ for } t \in [t_\nu, t_{\nu+1}),\ \nu \in [n-1]_0\Big\} \tag{5.36}$$

(the space of all $\mathcal{I}$-piecewise linear paths from $[0,1]$ to $\mathbb{R}^d$) and denote by $\hat{\iota}_n$ the map that sends a tuple $(z_\nu) \in Z^{\times n}$ to its piecewise-linear interpolation (cf. Rem. 5.2.1) along $\mathcal{E}_n := \big((\nu-1)/(n-1) \mid \nu \in [n]\big)$, the equidistant dissection of $[0,1]$. (The maps in (5.35) are indeed identifications as they are (isometric for $n = 1$) Banach isomorphisms with inverses $\hat{\pi}_n := \hat{\iota}_n^{-1}$ given by $\hat{\pi}_1 := \pi_1$ and $\hat{\pi}_n = \pi_{\mathcal{E}_n}$ for $n \geq 2$.) For any 'discrete signal' $\upsilon \in \mathcal{M}_1(Z^{\times n})$, the pushforward $\hat{\upsilon} := (\hat{\iota}_1)_*\upsilon \in \widehat{\mathfrak{D}}_1 :\cong \{\mu \in \mathcal{M}_1(\mathcal{C}^1) \mid \mu(\hat{\mathcal{C}}_0) = 1\}$ (if $n = 1$) or

$$\hat{\upsilon} := (\hat{\iota}_n)_*\upsilon \,\in\, \widehat{\mathfrak{D}}_{\mathcal{E}_n} := \{\mu \in \mathcal{M}_1(\mathcal{C}^1) \mid \mu(\hat{\mathcal{C}}_{\mathcal{E}_n}) = 1\} \quad (\text{for } n \geq 2)$$

is its continuous-time identification; for the ('timeless') case $n = 1$, the signature moments (5.38) of $\hat{\upsilon}$ coincide with the classical [multivariate] moments of υ up to the factor $1/m!$. ◆

5.4.2 Signature Moments

Let us now define the announced coordinates of a signal μ in $\mathfrak{D}$.

The key idea behind these coordinates is that each sample path $x \in \mathcal{C}^1$ of μ admits itself a concise non-local (in-time) description by way of the iterated-integral[14] statistics

$$x = (x_t^1, \cdots, x_t^d)_{t\in[0,1]} \quad\longmapsto\quad \int_{0\leq t_1\leq\cdots\leq t_m\leq 1} \mathrm{d}x_{t_1}^{i_1}\cdots \mathrm{d}x_{t_m}^{i_m} \;=:\; \mathfrak{sig}_{i_1\cdots i_m}(x) \tag{5.37}$$

for $i_1, \ldots, i_m \in [d]$, $m \in \mathbb{N}$, and $\mathfrak{sig}_\emptyset(x) := 1$. These are the so-called *signature coefficients* of x, and they combine to a multi-indexed family $\mathfrak{sig}(x) := (\mathfrak{sig}_w(x) \mid w \in [d]^\star)$ which is known as the *signature* of x (see Definition 2.2.6). Notice the formal similarity of $\mathfrak{sig}_{i_1\cdots i_m}$ with the moment map $x \mapsto x_{0,1}^{i_1}\cdots x_{0,1}^{i_m}$, motivating us to refer to (5.37) as the 'noncommutative moments' of x. See Section 2.2 for further details, and recall Section 2.2.3 for the definition below.

By duality, the signature-based description (5.37) of a path can be naturally transferred from $\mathcal{C}^1$ to the signals in $\mathfrak{D}$. The resulting 'dual' coordinates on $\mathfrak{D}$ are introduced below.

Lemma 5.4.5. *The map $(\mathcal{C}^1, \|\cdot\|_\infty) \ni x \mapsto \mathfrak{sig}_w(x) \in \mathbb{R}$ is Borel-measurable for each $w \in [d]^\star$.*

Proof. By Lemma 5.4.3 and the fact that $\mathfrak{sig}_w$ is continuous wrt. $\|\cdot\|_{\text{1-var}}$ (e.g. [105, Theorem 3.1.3]). □

Definition 5.4.6 (Signature Moments)**.** Given a signal μ in $\mathfrak{D}$, the numbers (in $\bar{\mathbb{R}}$)

$$\langle\mu\rangle_{i_1\ldots i_m} \;:=\; \int_{\mathcal{C}_d} \mathfrak{sig}_{i_1\ldots i_m}(x)\,\mu(\mathrm{d}x) \qquad (i_1, \ldots, i_m \in [d],\; m \in \mathbb{N}) \tag{5.38}$$

are called the *signature moments of* μ. We refer to the *coredinates of* μ as the matrices

$$[\mu]_0 := \big(\langle\mu\rangle_{ij}\big)_{i,j\in[d]} \quad\text{ and }\quad [\mu]_\nu := \big(\langle\mu\rangle_{ij\nu}\big)_{i,j\in[d]} \quad (\nu = 1, \ldots, d), \tag{5.39}$$

and denote by $[\mu]_\mathfrak{c} := ([\mu]_0, [\mu]_1, \cdots, [\mu]_d)$ their ordered collection. We say that $[\mu]_\mathfrak{c}$ *exists*, in symbols: $[\mu]_\mathfrak{c} < \infty$, if all the signature moments in (5.39) exist in $\mathbb{R}$.

(The word 'coredinates' is a portmanteau of 'core coordinates'.)

Owing to the above-mentioned characteristicness of the signature-based path coordinatisation (5.37), the signature moments (5.38), if contained in $\mathbb{R}$, characterise a signal uniquely, see [34]. As a result of the algebraic interrelations[15] within $\{\mathfrak{sig}_w \,|\, w \in [d]^\star\}$, the moments

[14] Since every $x \in \mathcal{C}^1$ is of bounded variation, the integrals (5.37) exist in the sense of Lebesgue-Stieltjes.

[15] The non-commutative moments $\{\mathfrak{sig}_w \mid w \in [d]^\star\}$ form a point-separating subalgebra of $\mathbb{R}^{\mathcal{C}^1}$.

(5.38) also efficiently encode any transformative action $f : \mathcal{C}^1 \to \mathcal{C}^1$ inflicted (via push-forward, cf. (5.6)) upon $\mathfrak{D}$. In fact, the coredinates (5.39) already sufficiently capture the *linear* action

$$\mathbb{R}^{d\times d} \ni A \longmapsto A\cdot\mu \equiv A_*\mu := \mu\circ A^{-1}\,; \tag{5.40}$$

flattening the third-order tensors $(\langle\mu\rangle_{ijk})$ to the matrices $[\mu]_\nu$ in (5.39) yields a family of affine equivariant statistics (see below). This equivariance is a simple yet extremely useful algebraic property of $[\mu]_{\mathfrak{c}}$, and when combined with mild statistical assumptions on μ can be used to study the action (5.40) for a wide class of (non-IC) signals and under nonlinear perturbations.

Lemma 5.4.7 (Affine Equivariance)**.** *Let $\mu \in \mathfrak{D}$ and define $f \equiv f(u) := A\cdot u + b$ on $\mathbb{R}^d$ for some $A = (a_{ij}) \in \mathbb{R}^{d\times d}$ and $b = (b_i) \in \mathbb{R}^d$. Then for any $i_1, \ldots, i_m \in [d]$ and $m \in \mathbb{N}$,*

$$\langle f_*\mu\rangle_{i_1\ldots i_m} \;=\; \sum_{j_1,\ldots,j_m=1}^{d} a_{i_1j_1}\cdots\, a_{i_mj_m}\cdot\langle\mu\rangle_{j_1\ldots j_m}. \tag{5.41}$$

Proof. This is clear by the multilinearity of iterated Stieltjes integration. Indeed: Fix any multiindex $(i_1, \ldots, i_m) \in [d]^{\times m}$. Then by said multilinearity, we find for each $x \in \mathcal{BV}$ that

$$\begin{aligned}
\mathfrak{sig}_{i_1\cdots i_m}(f(x)) \;&=\; \int_{\Delta_m} \mathrm{d}\big[f(x)\big]^{i_1}\cdots\mathrm{d}\big[f(x)\big]^{i_m}\\
&=\; \int_{\Delta_m} \mathrm{d}\big[\textstyle\sum_{j_1=1}^{d} a_{i_1j_1}x^{j_1}\big]\cdots\mathrm{d}\big[\textstyle\sum_{j_m=1}^{d} a_{i_mj_m}x^{j_m}\big]\\
&=\; \sum_{j_1,\ldots,j_m=1}^{d} a_{i_1j_1}\cdots a_{i_mj_m}\mathfrak{sig}_{j_1\cdots j_m}(x)
\end{aligned}$$

for the m-simplex $\Delta_m := \{t \in [0,1]^m \mid t_1 \leq \cdots \leq t_m\}$. So by definitions (5.38) and (5.6),

$$\begin{aligned}
\langle f_*\mu\rangle_{i_1\ldots i_m} \;&=\; \int_{\mathcal{C}_d} \mathfrak{sig}_{i_1\cdots i_m}(x)\, f_\star\mu(\mathrm{d}x) \;=\; \int_{\mathcal{BV}} \mathfrak{sig}_{i_1\cdots i_m}(f(x))\,\mu(\mathrm{d}x)\\
&=\; \sum_{j_1,\ldots,j_m=1}^{d} a_{i_1j_1}\cdots\, a_{i_mj_m}\cdot\langle\mu\rangle_{j_1\ldots j_m},
\end{aligned}$$

where the second identity is the change-of-variables formula for pushforward measures. □

Of course, the signature statistics (5.39) share their equivariance property (5.41) with many classical signal statistics of comparable order such as, e.g., covariance matrices or classical multivariate moments or cumulants, which is what underlies the latter's utility in many classical ICA approaches [38, 81, 111]. Hence and if preferred one may also just use these classical statistics in lieu of (5.38), in which case the topological robustness approach developed in this chapter remains equally valid, mutatis mutandis, upon replacement of (5.39) by any of their (affine equivariant) classical counterparts; cf. also Rem. 5.4.4. This leads us to the following remark.

Remark 5.4.8 (Why Signatures?)**.** In stark contrast to the classical ('commutative') moment statistics of a stochastic process, the coredinates (5.39) of a signal μ are 'global' [in time] statistics of this signal, i.e. functions that exhaust (the law of) the *whole* signal and not just, as with classical statistics, the signal's finite-dimensional distributions probed at a small number of preselected time-points. This globality of the signal statistics (5.38) enables them to capture non-local statistical effects (such as, e.g., time-dependent noise or estimation errors due to asynchronously sampled channel data, cf. Section 5.7) much more naturally and efficiently than with classical ('fixed-t') moment statistics (cf. e.g. Sect. 4.2). What is more, the derived statistics (5.39) do in general[16] also carry more information per matrix than classical statistics of the same dimension, which reflects in the fact the matrices (5.39) are generally asymmetric and, hence, contain about twice as much information as same-dimensional matrix statistics obtained from classical ('commutative') moments, which are all symmetric. These aspects in particular, which will be developed and leveraged throughout the following, make the signal coordinates (5.38) and their derived statistics (5.39) a very natural and informative quantisation of statistical dependencies within and between multidimensional signals and their (non)linear transformations. Owing to these qualities, we will see that these coordinates are an excellent basis for constructing the desired robustness topology in (5.24). ◆

Remark 5.4.9 (Extensions to Rougher Paths)**.** While the sample space (5.28) is fairly general already and, in particular, covers all signals that one encounters in statistical practice (cf. Rem. 5.4.4, Prop. 5.4.12), note that the considerations of this chapter can be readily extended to spaces of even more irregular paths by using the theory of rough paths [54, 104, 105]. As this is mainly a technical exercise, we refrain from implementing these extensions here and instead refer the interested reader to the above references to not exceed the scope of this chapter. ◆

5.4.3 Convergence of Signals

In this section we formalise how different signals can be topologically related to each other by introducing a natural notion of p-variation-graded weak convergence on $\mathfrak{D}$. The resulting signal topology is at most minimally stronger than the topology of classical weak-convergence on $\mathfrak{D}$ and in fact equal to the latter in the context of many applications.

[16] An exception to this are the coredinate matrices of mean-stationary product streams and their linear transformations (due to (5.85)), which do coincide with classical moment statistics; see Lemma 5.5.9.

To begin, recall[17] that a net of signals (μ_α) in $\mathfrak{D}$ is said to converge *weakly* to $\mu \in \mathfrak{D}$ if

$$\lim_\alpha \mu_\alpha(\tilde{f}) = \mu(\tilde{f}) \quad \text{for each } \tilde{f} \in C_\mathrm{b}(\mathcal{C}^1; \|\cdot\|_\infty), \quad \text{where} \quad \mu_\alpha(\tilde{f}) := \int_{\mathcal{C}^1} \tilde{f}\, \mathrm{d}\mu_\alpha \tag{5.42}$$

and $C_b(\mathcal{C}^1; \|\cdot\|_\infty)$ denotes the space of bounded $\|\cdot\|_\infty$-continuous functions $\tilde{f} : \mathcal{C}^1 \to \mathbb{R}$.

As is apparent from its above definition, the concept of weak convergence is an inherently topological notion: While different topologies on $\mathcal{C}^1$ may induce the same Borel σ-algebra, the set of (bounded) continuous functions they provide – and hence the 'modes of weak convergence' that they induce on $\mathfrak{D}$ – may well differ.

Lemma 5.4.3 pertains to this, as it implies that the signal spaces $\mathcal{M}_1(\mathcal{C}^1) \equiv \mathcal{M}_1((\mathcal{C}^1, \|\cdot\|_\infty)) \cong \mathfrak{D}$ and $\mathcal{M}_1((\mathcal{C}^1, \|\cdot\|_{1\text{-var}}))$ coincide. This allows us to integrate each measure $\mu \in \mathcal{M}_1(\mathcal{C}^1)$ against any element of the space $C_b(\mathcal{C}^1; \|\cdot\|_{1\text{-var}})$ of bounded $\|\cdot\|_{1\text{-var}}$-continuous functions $\tilde{f} : \mathcal{C}^1 \to \mathbb{R}$.

As before, cf. [18, Def. 8.1.1], this gives rise to a notion of pointwise ("weak") convergence

$$\mu_\alpha \xrightarrow{1\text{-var}} \mu \quad :\Longleftrightarrow \quad \Big[\lim_\alpha \mu_\alpha(\tilde{f}) = \mu(\tilde{f}), \quad \forall\, \tilde{f} \in C_b(\mathcal{C}^1; \|\cdot\|_{1\text{-var}})\Big] \tag{5.43}$$

for a net $(\mu_\alpha) \subset \mathcal{M}_1(\mathcal{C}^1)$ and $\mu \in \mathcal{M}_1(\mathcal{C}^1)$. Thanks to the inclusions (5.33), this convergence can be systematically generalised to the following weaker notions of convergence.

Definition 5.4.10 (p-Weak & Weak' Topology)**.** Let $p \in [1, \infty)$. A net $(\mu_\alpha) \subset \mathfrak{D}$ will be called *p-weakly convergent* to a signal $\mu \in \mathfrak{D}$, in symbols: $\mu_\alpha \xrightarrow{p\text{-var}} \mu$, if

$$\lim_\alpha \mu_\alpha(\tilde{f}) = \mu(\tilde{f}) \quad \text{for each } \tilde{f} \in C_b(\mathcal{C}^1; \|\cdot\|_{p\text{-var}}). \tag{5.44}$$

Given $1 \leq \alpha \leq \infty$, we say that (μ_α) is $[\alpha]$-*weakly convergent* to μ, in symbols: $\mu_\alpha \xrightarrow{[\alpha]} \mu$, if

$$\mu_\alpha \xrightarrow{q\text{-var}} \mu \quad \text{for } \textit{some } \ 1 \leq q < \alpha. \tag{5.45}$$

The convergences of nets (5.44) and (5.45) each canonically[18] define topologies on $\mathfrak{D}$ which we call the *p-weak topology* and the $[\alpha]$-*weak topology* on $\mathfrak{D}$, respectively. We call *weak'-topology* the [2]-weak topology, and refer to the $[\infty]$-weak topology as the *variationally-weak topology.*

In other words, (5.44) is precisely the weak convergence of (μ_α) to μ wrt. the p-variation topology on $\mathcal{C}^1$ (cf. [18, Section 8.1]). Note further that the convergences (5.44), (5.45) and (5.42) stand in natural hierarchy to each other: For any $1 \leq p \leq q < r < \infty$, the p-weak

[17] See for instance [18, Definition 8.1.1] and [16, Section 7].
[18] See for instance [91, Chapter 2 (pp. 73)].

topology is stronger (due to (5.33)) than the q-weak topology which, in turn, is stronger than the $[r]$-weak topology which is stronger than the variationally-weak topology. The latter, in turn, is stronger than the classical (i.e. (5.42)-induced) weak topology on $\mathcal{M}_1(\mathcal{C}^1)$ but only minimally so: these last two topologies are fact 'almost equivalent' as per the following remark.

Remark 5.4.11. The norms $\|\cdot\|_{\text{p-var}}$ $(1 < p < \infty)$ and $\|\cdot\|_\infty$ can be related via a classical interpolation inequality (e.g. [54, Proposition 5.5]) which, for every $x \in \mathcal{C}_d$, implies that

$$\|x\|_\infty \;\leq\; \|x\|_{p\text{-var}} \;\leq\; 2\|x\|_{1\text{-var}}^{\frac{1}{p}} \cdot \|x\|_\infty^{1-\frac{1}{p}} \qquad \text{for each } p \in (1, \infty) \tag{5.46}$$

and hence that the norms $\|\cdot\|_{p\text{-var}}$ and $\|\cdot\|_\infty$ are asymptotically equivalent in the limit $p \to \infty$. In particular, the convergence-defining sets of test functions $C_b(\mathcal{C}^1; \|\cdot\|_{p\text{-var}})$ and $C_b(\mathcal{C}^1; \|\cdot\|_\infty)$ are 'asymptotically coincidental' as $p \to \infty$, and in this sense we may regard the convergence topologies defined by $(5.45)|_{alpha=infty}$ resp. (5.42) as 'essentially equivalent'. ◆

The a priori gap (in strength) between the p-weak topologies and the classical [i.e. (5.42)-induced] weak topology on $\mathfrak{D}$ is mainly a theoretical one with no relevance for many practical applications. This is because for many practically relevant signal spaces the p-weak topologies and the classical weak topology coincide, as the following example shows.

Proposition 5.4.12. *For $p \in [1, \infty)$ and $K > 0$, consider the 'compactified' signal spaces*

$$\mathfrak{D}_K^p := \{\mu \in \mathcal{M}_1(\mathcal{C}^1) \mid \mu(\mathcal{C}_{p,K}^1) = 1\} \quad \textit{for} \quad \mathcal{C}_{p,K}^1 := \{x \in \mathcal{C}^1 \mid \|x\|_{p\text{-var}} \leq K\}, \tag{5.47}$$

and for any dissection $\mathcal{I}$ of $[0,1]$ consider the signals

$$\widehat{\mathfrak{D}}_{\mathcal{I}} := \{\mu \in \mathcal{M}_1(\mathcal{C}^1) \mid \mu(\hat{\mathcal{C}}_{\mathcal{I}}) = 1\} \;\; \textit{supported on} \;\; \hat{\mathcal{C}}_{\mathcal{I}} \;\; (\textit{see } (5.36)). \tag{5.48}$$

On these spaces the variationally-weak topology and the classical weak topology are equivalent, and the classical weak topology is further equivalent to: the q-weak topology on $\mathfrak{D}_K^p$ for any $q > p$, and to the 1-weak topology on $\widehat{\mathfrak{D}}_{\mathcal{I}}$.

Proof. This essentially follows from a generalisation of (5.46). Indeed: Note first that the above subspaces are well-defined since $\mathcal{C}_{p,K}^1$ and $\hat{\mathcal{C}}_{\mathcal{I}}$ are both Borel-measurable wrt. $\|\cdot\|_\infty$.

Indeed: $\mathcal{C}_{p,K}^1$ is $\|\cdot\|_{1\text{-var}}$-closed [as a sublevel set] by (5.33) (these inclusions of course remain valid also if the 'b's are dropped) and $\mathcal{B}(\mathcal{C}^1; \|\cdot\|_{1\text{-var}}) = \mathcal{B}(\mathcal{C}^1; \|\cdot\|_\infty)$ by Lemma 5.4.3, and the $\mathcal{C}^1$-subspace $\hat{\mathcal{C}}_{\mathcal{I}}$ is $\|\cdot\|_\infty$-closed because it is finite-dimensional.

Using next that (5.46) can be generalised to $\|\cdot\|_\infty \leq \|\cdot\|_{q\text{-var}} \leq 2\|\cdot\|_{p\text{-var}}^{\frac{p}{q}} \|\cdot\|_\infty^{1-\frac{p}{q}}$ for any finite $q > p$, see [54, Prop. 5.5 (i)], we obtain that the norms $\|\cdot\|_\infty$ and $\|\cdot\|_{q\text{-var}}$ [have the same convergent sequences on $\mathcal{C}^1_{p,K}$ and thus] are equivalent on $\mathcal{C}^1_{p,K}$. In particular $C_b(\mathcal{C}^1_{p,K};\|\cdot\|_\infty) = C_b(\mathcal{C}^1_{p,K};\|\cdot\|_{q\text{-var}})$, whence for any net $(\mu_\alpha)\subset\mathfrak{D}^p_K\cong\mathcal{M}_1(\mathcal{C}^1_{p,K})$ we have: $\mu_\alpha\to\mu\in\mathfrak{D}^p_K$ weakly iff $\mu_\alpha\xrightarrow{q\text{-var}}\mu$ (cf. (5.43) and (5.44), with $\mathcal{C}^1$ replaced by $\mathcal{C}^1_{p,K}$). The latter convergence is equivalent to $\mu_\alpha\xrightarrow{[\infty]\text{-var}}\mu$ in $\mathfrak{D}^p_K$ because of $1<q<\infty$ and $C_b(\mathcal{C}^1_{p,K};\|\cdot\|_{q\text{-var}}) = C_b(\mathcal{C}^1_{p,K};\|\cdot\|_\infty)$ and (5.33) [with $\mathcal{C}^1$ replaced by $\mathcal{C}^1_{p,K}$]. The lemma's first two assertions now follow by the fact that two topologies coincide iff they have the same convergent nets.

The final assertion follows likewise upon noting that $C_b(\hat{\mathcal{C}}_\mathcal{I};\|\cdot\|_\infty) = C_b(\hat{\mathcal{C}}_\mathcal{I};\|\cdot\|_{1\text{-var}})$, which holds by the equivalence of $\|\cdot\|_\infty$ and $\|\cdot\|_{1\text{-var}}$ on $\hat{\mathcal{C}}_\mathcal{I}$. With $\|\cdot\|_{1\text{-var}}\geq\|\cdot\|_\infty$ known, the latter follows from $\|x\|_{1\text{-var}} = |x_0| + \sum_{\nu=1}^{n-1}|x_{t_\nu}-x_{t_{\nu-1}}|\leq(2|\mathcal{I}|-1)\|x\|_\infty$ (any $x\in\hat{\mathcal{C}}_\mathcal{I}$). □

The above signal spaces (5.47) and (5.48) owe their practical relevance to the fact that real-world signal processing systems are always subject to capacity limitations such as bounded data storage and finite resolution of the registered observables (e.g. [42, Sects. 5, 12]). The former type of constraints gives rise to (5.47), while the latter type mandates (finite horizon) time-discretizations of the signals as represented by (5.48), cf. e.g. [121].

5.4.4 Convergence of Signature Moments

This section shows that under a natural growth condition, the signature coordinates of a signal are continuous wrt. the above signal topologies.

To combine the consideration of Sections 5.4.1 and 5.4.3 into an informative topology for (5.24), we will need to link the convergence (5.45) of signals to the convergence of their moments (5.38). As usual for integral statistics, this can be done by a growth condition.

Definition 5.4.13. A subset $\mathfrak{U}$ of $\mathfrak{D}$ is called *uniformly signature integrable of order m* if

$$\inf_{a>0}\sup_{\mu\in\mathfrak{U}}\int_{\{|\mathfrak{sig}_w|>a\}}|\mathfrak{sig}_w|\,\mathrm{d}\mu \;=\; 0 \qquad \text{for each } w\in[d]^\star_{\leq m}. \tag{5.49}$$

We call $\mathfrak{U}$ *uniformly signature integrable* if $\mathfrak{U}$ is uniformly signature integrable of order 3.

Put plainly, a set of signals is uniformly signature integrable (of order m) if its elements attain with uniformly lower probability those sample paths whose signature coefficients (up to order m) are very large. The growth assumption (5.49) ensures that on such spaces, the moment coordinates (5.38) are continuous with respect to the weak' topology $(5.45)|_{\alpha=2}$.

Lemma 5.4.14. *Let $\mathfrak{U}$ be uniformly signature integrable of order m, and $(\mu_\alpha)\subset\mathfrak{U}$ be any net such that $\mu_\alpha\to\mu\in\mathfrak{U}$ in the weak'-topology. Then $\langle\mu_\alpha\rangle_w\to\langle\mu\rangle_w$ for each $w\in[d]^\star_{\leq m}$.*

Proof. Since $\langle\tilde{\mu}\rangle_w = \int_{\mathcal{C}^1} \mathfrak{sig}_w \, \mathrm{d}\tilde{\mu}$ and because $\mathfrak{sig}_w$ is continuous wrt. the p-variation topology on $\mathcal{C}^1$ for any $p \in [1,2)$, e.g. [105, Thm. 3.1.3], the assertion holds by [18, Lemma 8.4.3]. $\square$

As an obvious[19] example, uniformly signature integrable spaces (of any order) include all subsets of $\mathfrak{D}$ whose elements have uniformly $\|\cdot\|_{p\text{-var}}$-bounded support, cf. Proposition 5.4.12. Drawing on the classical concept of uniform integrability, these bounded-support examples of uniformly signature integrable spaces can be easily generalised as follows.

Proposition 5.4.15. *Let $m \in \mathbb{N}$, and let $\mathfrak{U}$ be a subset of $\mathfrak{D}$ for which there is $q > 1$ s.t.*

$$\sup_{\mu\in\mathfrak{U}} \int_{\mathcal{C}_d} \|x\|_{1\text{-var}}^{mq} \, \mu(\mathrm{d}x) \; < \; \infty. \tag{5.50}$$

Then $\mathfrak{U}$ is uniformly signature integrable of order m.

Proof. Let $\mathfrak{U} \subseteq \mathcal{M}_1(\mathcal{C}^1)$ satisfy (5.50) for some $q > 1$. Then for any $\mu \in \mathfrak{U}$ and with r the conjugate index of q,

$$\int_{\{|\mathfrak{sig}_w|>a\}} |\mathfrak{sig}_w| \, \mathrm{d}\mu \;\leq\; \frac{c}{m!} \int_{\mathcal{C}_d} \|x\|_{1\text{-var}}^{m} \cdot \mathbb{1}_{\{|\mathfrak{sig}_w|>a\}} \, \mu(\mathrm{d}x) \tag{5.51}$$

$$\leq\; \frac{c}{m!} \left(\int_{\mathcal{C}_d} \|x\|_{1\text{-var}}^{mq} \, \mu(\mathrm{d}x) \right)^{\frac{1}{q}} \cdot \Big(\mu(\{|\mathfrak{sig}_w| > a\}) \Big)^{\frac{1}{r}} \;\overset{a\to\infty}{\longrightarrow}\; 0 \tag{5.52}$$

for all $w \in [d]^{\leq m}$, where the first inequality is due to the classical signature bound [104, Theorem 3.7] (the constant c is specified there), the second inequality is due to Hölder, and the final convergence to zero follows by Chebyshev's inequality. The latter in fact implies

$$\sup\nolimits_{\mu\in\mathfrak{U}} \mu(\{|\mathfrak{sig}_w| > a\}) \;\leq\; \frac{1}{a} \sup\nolimits_{\mu\in\mathfrak{U}} \int_{\mathcal{C}_d} |\mathfrak{sig}_w(x)| \, \mu(\mathrm{d}x) \;\lesssim\; \frac{1}{a} \sup\nolimits_{\mu\in\mathfrak{U}} \int_{\mathcal{C}_d} \|x\|_{1\text{-var}}^{m} \, \mu(\mathrm{d}x) \;\to\; 0$$

as $a \to \infty$, where the last inequality is again due to the cited signature bound and the last supremum is finite by assumption (5.50) (recalling $L^q(\mu) \subseteq L^1(\mu)$, as the measures μ are all finite). Both (5.52) and the above $\mathfrak{U}$-uniformity of this convergence combine to (5.49). $\square$

5.4.4.1 Integrability

An analysis of the robustness (5.24) requires to capture how nonlinear transformations of $\mathbb{R}^d$ (via (5.6)) affect close-by source signals. For this – and only for this, see Rems. 5.5.15, 5.8.10 (ii) – it will be opportune to sharpen the notion of locality in $\mathfrak{D}$ by enforcing the comparability of (coredinates of) signals via uniform integrability (5.50) of their strong moments:

[19] Upon recalling that $|\mathfrak{sig}_w(x)| \lesssim \|x\|_{p\text{-var}}^{p|w|}$ for any $x \in \mathcal{C}^1$ and $p \in [1,2)$, e.g. [104, Thm. 3.7 & Sect. 1.2.2]; notice that said inequality also (and for all $\mu \in \mathcal{M}_1$) implies that: $\int_{\mathcal{C}_d} \|x\|_{p\text{-var}}^{3} \, \mu(\mathrm{d}x) < \infty \implies [\mu]_{\mathfrak{c}} < \infty$.

Notation. A set $\mathfrak{U}$ in $\mathfrak{D}$ is called *core integrable* if it is uniformly signature integrable and

$$K_{\mathfrak{U}|\beta} := \sup_{\mu\in\mathfrak{U}} \int_{\mathcal{C}_d} \|x\|_{1\text{-var}}^{\beta}\, \mu(\mathrm{d}x) \,<\, \infty \quad \text{for some } \; \beta \in (5/2, 3). \tag{5.53}$$

For any fixed β as in (5.53), denote $p_\beta := (\beta-2)^{-1}$ and $\alpha = \alpha_\beta := 1 - 1/p_\beta$ $[\in (0, \frac{1}{2})]$.

We say that a subset $\mathfrak{U} \subseteq \mathfrak{D}$ is *core-integrable as per* $(K_{\mathfrak{U}|\beta}, \beta, p, \alpha)$ if (5.53) holds for some fixed $\beta \in (5/2, 3)$ and associated constants $p = p_\beta$ and $\alpha = \alpha_\beta$ as above, and for any such tuple we further define $c_p := (1 - 2^{1-2/p})^{-1}$ and $C_p := 2^{\alpha+1} c_p (1 + K_{\mathfrak{U}|\beta})$. ◆

As a trivial example, a subspace $\mathfrak{U} \subseteq \mathfrak{D}$ is core integrable if it satisfies (5.50) for $m = 3$. We conclude with a bit of context on the condition (5.50) and how it relates to the usual integrability conditions for classical moments of random vectors in $\mathbb{R}^d$.

Remark 5.4.16. The sets $\mathfrak{U}$ in $\mathfrak{D}$ that satisfy (5.50) are precisely the bounded subsets of the mq-Wasserstein space on $(\mathcal{C}^1, \|\cdot\|_{1\text{-var}})$; for a discussion of those, see e.g. [122, Sects. 2.1, 2.2]. Other than that, signals satisfying (5.50) are also known as Radon measures of order mq, and they are central to the theory of Radonifying mappings developed by Laurent Schwartz [142, 143] and others. Such signals and their uniformly bounded families are described in many sources, see e.g. [17, 56, 103, 122, 146] and associated references. For the special case $\mu \equiv \hat{\nu} \in \widehat{\mathfrak{D}}_1$ of random vectors in $\mathbb{R}^d$ (cf. Remark 5.4.4), the integrals in (5.49) and (5.53) reduce to the classical statistics

$$\frac{1}{|w|!} \int_{\{|u^w| > |w|! a\}} |u^w|\, \nu(\mathrm{d}u) \quad \text{ and } \quad \int_{\hat{\mathcal{C}}_0} \|x\|_{1\text{-var}}^{\tilde{q}}\, \hat{\nu}(\mathrm{d}x) = \int_{\mathbb{R}^d} |u|^{\tilde{q}}\, \nu(\mathrm{d}u),$$

respectively, with $u^{i_1\cdots i_m} := u_1^{\sharp\{j|i_j=1\}} \cdots u_d^{\sharp\{j|i_j=d\}}$ for $u = (u_i) \in \mathbb{R}^d$. Both (5.49) and (5.53) thus appear as natural generalisations from $\mathbb{R}^d$ to $\mathcal{C}_d$ of classical integrability conditions for multivariate moments, cf. e.g. [16, pp. 30]. ◆

Our approach towards ICA-robustness ((5.15), (5.24)) will make use of the coredinates (5.39) of a signal, which is why for the following we focus on the 'coredinatisable' space

$$\dot{\mathfrak{D}} \,:=\, \{\mu \in \mathfrak{D} \mid [\mu]_{\mathfrak{c}} < \infty,\, \langle\mu\rangle_{ii} \neq 0 \;\text{ for } i = 1, \ldots, d\}\,. \tag{5.54}$$

Note that this is a very large subspace of $\mathfrak{D}$: The required existence of a signal's coredinates includes (e.g.) the full 3-Wasserstein space on $(\mathcal{C}^1, \|\cdot\|_{1\text{-var}})$ (Footnote 19; cf. [122, Sect. 2.1]), while the condition $\langle\mu\rangle_{ii} \neq 0$ of non-vanishing 2^{nd}-order diagonal moments is merely a mild non-degeneracy assumption that can be imposed with essentially[20] no loss of generality.

[20] Indeed: Since for each $\mu = (\mu^1, \ldots, \mu^d) \in \mathfrak{D}$ we have $\langle\mu\rangle_{ii} = \frac{1}{2}\mathbb{E}[(\mu^i_{0,1})^2]$, the condition $\langle\mu\rangle_{ii} \neq 0$ asks for the increments $\mu^i_{0,1} := \mu^i_1 - \mu^i_0$ to have non-zero variance (see Sect. 5.2.2 for notation). Practically this is no restriction since one can simply cut-off any ('non-degenerate', cf. Sect. 5.2.1) signal in $\mathfrak{D}$ to a (regular) interval over whose end-points its increment has positive variance and then rescale this cut-off signal back to $[0, 1]$.

We are now in the position to endow the signal space $\dot{\mathfrak{D}}$ with an explicit and well-structured topology that will help us to arrive at an insightful quantification of the robustness (5.24).

5.5 An ICA-Tailored Topology on Causal Space

We describe an informative and explicitly computable premetric topology on the causal space $\mathfrak{C}$ that relates to the coarseness requirement (5.25) while still being strong enough to, in a conveniently quantifiable manner, support the robustness (5.24) for a generic ICA-solution $(\mathscr{I}, \hat{\Phi})$ such as the one derived in Section 5.6. As the first of three main steps to this end, Section 5.5.1 shows how the signature moments from Definition 5.4.6 enable a systematic, divergence-like quantisation of the (statistical) 'distance' between two signals in $\dot{\mathfrak{D}}$, and that the information thus encoded amounts to a topological profile of statistical dependence within and between signals that is algebraically and analytically flexible and sensitive to both spatial (non)linear transformations and more intrinsic statistical variations, including those of Sect. 5.4.3. This topological landscape is then connected to the ICA-specific identifiability structure (5.18) in Section 5.5.2, where the subset of accurately identifiable source signals is charted as (belonging to) the minimizing set of this landscape. Finally, Section 5.5.3 makes good on the final announcement of Sect. 5.3.2 by extending the premetric signal topology to a coarse topology on the whole causal space $\mathfrak{C}$, for which robustness (5.24) will then be established in Section 5.6.

5.5.1 A Premetric Topology on Signal Space

Out of the countable infinitude (5.38) of coordinates that the signature moments provide, the lower-order arrangement (5.39) captures those aspects of a signal which are most expressive of any linear action (5.40) inflicted upon it. In addition, and similar still to covariance matrices or cumulants for random vectors, the coredinate matrices (5.39) describe the (second- and third-order) statistical dependence between the components of a signal by how 'pronounced' their off-diagonal structure is relative to their diagonal, cf. Lemma 5.5.9. This, combined with the inversion theory of Section 5.6.1, will allow for an explicit topological coordinatisation of ICA-identifiable causes (Fig. 5.1 a.) and the causes close to them (Fig. 5.1 b.). Let $\dot{\mathfrak{D}}$ be the signal space from (5.54).

We propose to implement this charting by way of the following distance function. (Recall that $\mu \equiv (\mu_t) \in \mathcal{M}_1$ is called mean-stationary if $\mathbb{E}\mu_t = \mathbb{E}\mu_0$ for each $t \in [0, 1]$ (Sect. 5.2.2).)

Definition 5.5.1 (ICA-Premetric)**.** For any signals $\mu, \tilde{\mu} \in \dot{\mathfrak{D}}$ with coredinates (5.39), define

$$\delta(\mu, \tilde{\mu}) \;:=\; \sqrt{\alpha_{\mu,\tilde{\mu}} + \beta_{\mu,\tilde{\mu}}} \qquad \text{for} \tag{5.55}$$

$$\alpha_{\mu,\tilde{\mu}} := \sum_{i,j=1}^{d}\left(\frac{\langle\tilde{\mu}\rangle_{ij} - \langle\mu\rangle_{ij}}{\sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}}}\right)^2 \quad \text{and} \quad \beta_{\mu,\tilde{\mu}} := \sum_{\nu=1}^{d}\sum_{i,j=1}^{d}\left(\frac{\langle\tilde{\mu}\rangle_{ij\nu} - \langle\mu\rangle_{ij\nu}}{\sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}\langle\mu\rangle_{\nu\nu}}}\right)^2.$$

For a mean-stationary $\mu = (\mu^1, \ldots, \mu^d) \in \dot{\mathfrak{D}}$ we refer to the *IC-defect* of μ as the number

$$\delta_{\perp\!\!\!\perp}(\mu) \;:=\; \delta(\mu^1 \otimes \cdots \otimes \mu^d, \mu)\,; \tag{5.56}$$

for a non mean-stationary signal, this defect is defined[21] by equation (5.62) below.

For signals $\mu, \tilde{\mu} \in \dot{\mathfrak{D}}$ with coredinates $[\mu]_{\mathfrak{c}}, [\tilde{\mu}]_{\mathfrak{c}}$ as in (5.39) and the standardisation matrix $N_\mu := \operatorname{ddiag}(\langle\mu\rangle_{11}, \cdots, \langle\mu\rangle_{dd})^{1/2}$, notice that (5.55) can be written more compactly as

$$\delta(\mu, \tilde{\mu}) \;=\; \left\| N_\mu^{-1} \cdot \left([\tilde{\mu}]_0 - [\mu]_0, \frac{[\tilde{\mu}]_1 - [\mu]_1}{\sqrt{\langle\mu\rangle_{11}}}, \ldots, \frac{[\tilde{\mu}]_d - [\mu]_d}{\sqrt{\langle\mu\rangle_{dd}}}\right) \cdot N_\mu^{-1} \right\|_{\mathcal{V}} \tag{5.57}$$

where $\|(C_0, \cdots, C_d)\|_{\mathcal{V}} := \sqrt{\sum_{\nu=0}^{d} \|C_\nu\|^2}$ is the ℓ_2-norm on the space $\mathcal{V} := (\mathbb{R}^{d\times d})^{\oplus (d+1)}$ derived from the Frobenius norm $\|\cdot\|$ on $\mathbb{R}^{d\times d}$.

The map $(\mu, \tilde{\mu}) \mapsto \delta(\mu, \tilde{\mu})$ defines a *premetric* on $\dot{\mathfrak{D}}$ (cf. [7, Def. 2.4.4]), which is to say that

$$\delta(\mu, \tilde{\mu}) \geq 0 \quad \text{and} \quad \delta(\mu, \mu) = 0 \tag{5.58}$$

for any two $\mu, \tilde{\mu} \in \dot{\mathfrak{D}}$. This premetric canonically (e.g. [23, Section 2]) induces a topology τ_δ on the signal space $\dot{\mathfrak{D}}$, of which we note the following basic facts.

Lemma 5.5.2. *For the premetric topology τ_δ on $\dot{\mathfrak{D}}$ induced by* (5.55)*, the following holds.*

(i) *For each subset $\mathcal{O}$ in $\dot{\mathfrak{D}}$, we have $\mathcal{O} \in \tau_\delta$ iff for each $\mu \in \mathcal{O}$ there is an $r > 0$ such that*

$$B_\delta(\mu, r) := \{\tilde{\mu} \in \dot{\mathfrak{D}} \mid \delta(\mu, \tilde{\mu}) < r\}, \tag{5.59}$$

the μ-centred δ-ball of radius r, is contained in $\mathcal{O}$.

(ii) *The topological space $(\dot{\mathfrak{D}}, \tau_\delta)$ is sequential.*

(iii) *For any metric space (W, d_W), a map $\phi : (\dot{\mathfrak{D}}, \tau_\delta) \to W$ is continuous at $\mu \in \dot{\mathfrak{D}}$ iff:*

$$\forall\, \tilde{\varepsilon} > 0 \,:\, \exists\, \tilde{\delta} > 0 \,:\, \forall\, \tilde{\mu} \in \dot{\mathfrak{D}} \,:\; \delta(\mu, \tilde{\mu}) < \tilde{\delta} \implies d_W(\phi(\mu), \phi(\tilde{\mu})) < \tilde{\varepsilon}. \tag{5.60}$$

[21] Lemma 5.5.9 guarantees that the two definitions (5.56) and (5.62) are consistent (Remark 5.5.6).

Proof. The first two points follow directly from (5.58) – for reference, item (i) is reported in [7, p. 23] while statement (ii) is [7, Prop. 2.4.9]. Item (iii) follows from the statements (i) and (ii) and [7, Prop. 3.1.3]; or alternatively from [23, Lemma 2.2 and Cor. 2.3]. □

Remark 5.5.3. The ε-δ characterisation (5.60) of (topological) continuity holds more generally if ϕ is a metric-space-valued map on any premetric space, see e.g. [23, Lem. 2.2 & Cor. 2.3].

Proposition 5.5.4. *For any uniformly signature integrable subset $\tilde{\mathfrak{D}}$ of $\dot{\mathfrak{D}}$, the (subspace) topology τ_δ on $\tilde{\mathfrak{D}}$ is coarser than the weak' topology on $\tilde{\mathfrak{D}}$.*

Proof. Let $\tilde{\mathfrak{D}} \subseteq (\dot{\mathfrak{D}}, \tau_\delta)$ be uniformly signature integrable, and recall that the τ_δ-subspace topology on $\tilde{\mathfrak{D}}$ is induced by the restricted premetric $\delta|_{\tilde{\mathfrak{D}}^{\times 2}}$. The assertion then holds if for any net (μ_α) in $\tilde{\mathfrak{D}}$ we have that (recalling Definition 5.4.10 for notation):

$$\mu_\alpha \xrightarrow{[2]} \mu \in \tilde{\mathfrak{D}} \quad \text{implies} \quad \delta(\mu, \mu_\alpha) \to 0, \tag{5.61}$$

see also [23, Lemma 2.1]. By the definition (5.55) of δ, the desired implication (5.61) clearly holds if, for any (μ_α) and μ as above and any $i_1, \ldots, i_m \in [d]$ for $m \leq 3$, we have that:

$$\mu_\alpha \to \mu \;\text{ weakly'} \quad \text{implies} \quad \langle \mu_\alpha \rangle_{i_1 \cdots i_m} \to \langle \mu \rangle_{i_1 \cdots i_m};$$

this last implication, however, is guaranteed by Lemma 5.4.14, which concludes the proof. □

The idea behind the (5.57)-induced topologization of $\mathcal{M}_1$ is closely related to the natural idea of maximum mean discrepancy (MMD) over the finite set of test functions $\mathcal{H} := \{\mathfrak{sig}_w(\cdot) \mid |w| \leq 3\}$ and combined with an appropriate standardization that absorbs the $\tilde{\mathrm{M}}_d$-inflicted ('redundant') degrees of freedom which are due to the $\lceil \cdot \rceil_{\tilde{\mathfrak{m}}}$-controlled indistinguishability of the true source (Definition 5.3.7).

While not required here, note that the premetric (5.57) can be extended to a full metric via the systematic addition of higher-order signature moments.

Remark 5.5.5. The premetric (5.57) does not satisfy the identity of indiscernibles ("$\delta(\mu, \tilde{\mu}) = 0$ iff $\mu = \tilde{\mu}$") as would be required of (a quasisemimetric or) a statistical divergence on $\mathfrak{D}$. This 'missing' identity can be restored, however (see for instance [34, Section 5.4] for a proof), if instead of only its 2^{nd} and 3^{rd}-order moments one takes into account all the signature moments (5.38) of a signal and computes their ℓ_2-distance in the associated (Hilbert) coordinate space $\mathcal{V}^\infty := \bigoplus_{m \geq 1} (\mathbb{R}^d)^{\otimes m}$ which is an isometric extension of the above coordinate space $\mathcal{V} \cong (\mathbb{R}^d)^{\otimes 2} \oplus (\mathbb{R}^d)^{\otimes 3}$ that houses the coredinates (5.39). While higher-order extensions of this and other kinds are possible (cf. Section 5.9.4), the following sections show that for the present case of classical [i.e. linear] BSS/ICA, the premetric topology τ_δ associated [via (5.55)] to the third-order cap (5.39) is already sufficient. ◆

Let us include the following addition to Definition 5.5.1.

Remark 5.5.6. By Lemma 5.5.9, the IC-defect of a mean-stationary signal $\mu \in \dot{\mathfrak{D}}$ reads

$$\delta_{\perp\!\!\!\perp}(\mu) \;=\; \left\| N_\mu^{-1}\left([\mu]_0 - (\delta_{ij}\langle\mu\rangle_{ij}), \frac{[\mu]_1 - (\delta_{ij1}\langle\mu\rangle_{ij1})}{\sqrt{\langle\mu\rangle_{11}}}, \ldots, \frac{[\mu]_d - (\delta_{ijd}\langle\mu\rangle_{ijd})}{\sqrt{\langle\mu\rangle_{dd}}} \right) N_\mu^{-1} \right\|_{\mathcal{V}} ; \tag{5.62}$$

for a signal $\mu \in \dot{\mathfrak{D}}$ that is not mean-stationary, we declare its IC-defect $\delta_{\perp\!\!\!\perp}(\mu)$ to be defined by the right-hand side of (5.62). Clearly $\mu \in \dot{\mathfrak{D}}$ is orthogonal iff $\delta_{\perp\!\!\!\perp}(\mu) = 0$, and $\delta_{\perp\!\!\!\perp}$ is invariant under the M_d-action on $\dot{\mathfrak{D}}$ as follows directly from (5.62) and Lemma 5.4.7. ◆

The IC-defect $\delta_{\perp\!\!\!\perp}$ defined by (5.55) and (5.62), is an auxiliary function on $\dot{\mathfrak{D}}$ with which the blind identifiability (5.13) of a (non)linearly transformed signal can be quantitatively controlled: Source signals at which this defect vanishes (Sect. 5.5.2) will be exactly identifiable, while the blind inversion for a source becomes increasingly inaccurate – or impossible altogether – as its IC-defect increases, see Section 5.6.

In anticipation of this, the next lemma paves the way for an informative robustness analysis (5.24) by establishing moduli of continuity for $\delta_{\perp\!\!\!\perp}$ that relate to, respectively, the above premetric topology on $\dot{\mathfrak{D}}$ and the uniform topology on the space of spatial mixing transformations.

Here and in the sequel, the Jacobian D_R of a map $R \in C^1(\mathbb{R}^d;\mathbb{R}^d)$ is normed via $\|D_R\|_\infty := \sup_{x\in\mathbb{R}^d}\|D_R(x)\|_2$ where $\|\cdot\|_2$ is the Euclidean operator norm ('spectral norm') on $\mathbb{R}^{d\times d}$.
Further, the notation from Subsection 5.4.4.1 is in use.

Lemma 5.5.7. *Let $\hat\mu \in \dot{\mathfrak{D}}$ and $B_r \equiv B_\delta(\hat\mu, r)$ be the $\hat\mu$-centered δ-ball of radius $r > 0$. (See (5.59)). Suppose that $r < \hat r := \rho_0 \wedge (\rho_0/\rho_1)$, with $\rho_0 := \min_{i\in[d]}\langle\hat\mu\rangle_{ii}$ and $\rho_1 := \max_{i\in[d]}\langle\hat\mu\rangle_{ii}$.*

(i) *There are constants $\mathfrak{m}_{\hat\mu|r} = \mathfrak{m}_{\hat\mu|r}(\hat\mu, r), K_{\hat\mu|r} = K_{\hat\mu|r}(\hat\mu, r) \geq 0$, growing monotonously in r, such that for each $\mu \in B_r$ we have $\max_{|w|=2,3}\left|\langle\mu\rangle_w - \langle\hat\mu\rangle_w\right| \leq \mathfrak{m}_{\hat\mu|r}\cdot r$ and*

$$\left|\delta_{\perp\!\!\!\perp}(\mu) - \delta_{\perp\!\!\!\perp}(\hat\mu)\right| \;\leq\; K_{\hat\mu|r}\sqrt{r + r^2}\,; \tag{5.63}$$

(ii) *if $B_\delta(\hat\mu, \rho_0)$ is core-integrable as per (K_0, β, p, α), then for each $R \in C^1(\mathbb{R}^d;\mathbb{R}^d)$ there exists a constant $\tilde K_{\hat\mu|r} = \tilde K_{\hat\mu|r}(\hat\mu, r, K_0, R) \geq 0$, growing monotonously in r, such that: for each $\mu \in B_r$ we have, for the auxiliary functions $\phi_k(\cdot)$ from Lemma 5.8.1 (iv),*

$$\left|\delta_{\perp\!\!\!\perp}((I+R)_\star\mu) - \delta_{\perp\!\!\!\perp}(\mu)\right| \;\leq\; \tilde K_{\hat\mu|r}\cdot\phi_3(\|D_R\|_\infty)^{\frac{1}{2}}\left(\|D_R\|_\infty\|R\|_\infty\right)^{\frac{\alpha}{2}} \tag{5.64}$$

provided that $\rho_R := C_p\phi_2(\|D_R\|_\infty)(\|D_R\|_\infty\|R\|_\infty)^\alpha < \rho_0 - \rho_1 r$.

For any given $c, C > 0$ *with* $c < \rho_0$*, the constant* $\tilde{K}_{\hat{\mu}|r}$ *in* (5.64) *can be chosen to apply uniformly – that is, dependent only on* $\hat{\mu}, K_0$ *and* c, C *but independent of* R*,* μ *and* r *– on the set*

$$\mathcal{R}_c^C := \{(\mu, R) \in \dot{\mathfrak{D}} \times C^1(\mathbb{R}^d; \mathbb{R}^d) \mid \rho_1\delta(\hat{\mu}, \mu) + \rho_R \leq c \text{ and } \rho_R \vee \|D_R\|_\infty \leq C\}. \tag{5.65}$$

All of the above constants are explicit.

Proof. Note that any functional φ which acts on signals in $\dot{\mathfrak{D}}$ via their coredinates (5.39) (such as φ_κ, $\mathfrak{m}_\mu$ etc. in Lemma 5.8.1), factorizes as $\varphi = \phi_\varphi \circ [\,\cdot\,]^{\times k}$ with $[\,\cdot\,]^{\times k} : \dot{\mathfrak{D}}^{\times k} \to \mathcal{V}^{\times k}$ and $\phi_\varphi : \mathcal{V}^{\times k} \to \mathbb{R}$ (for $(\mathcal{V}, \|\cdot\|_\mathcal{V})$ as in (5.57) and k the number of arguments of φ), where $[\,\cdot\,] : \mu \mapsto ([\mu]_\nu) \in \mathcal{V}$ is the coredinate map (see (5.38)) and ϕ_φ is the coredinate action of φ.

Note further that by definition (5.59) of the δ-balls and the restriction $r < \hat{r}$, the image

$$[B_r] \equiv [\,\cdot\,](B_r) \quad \text{is contained in} \quad \hat{\mathfrak{B}} := \bar{B}_{q\rho_0}([\hat{\mu}]_0) \times \bar{B}_{q\rho_0\sqrt{\rho_1}}([\hat{\mu}]_{\nu\geq 1}) \equiv \hat{\mathfrak{B}}_0 \times \hat{\mathfrak{B}}_1 \tag{5.66}$$

for some $q \equiv q_r \in (0,1)$, where we defined $\hat{\mathfrak{B}}_0 \equiv \bar{B}_{q\rho_0}([\hat{\mu}]_0) := \{C \in \mathbb{R}^{d\times d} \mid \|C - [\hat{\mu}]_0\| \leq q\rho_0\}$ and $\hat{\mathfrak{B}}_1 \equiv \bar{B}_{q\rho_0\sqrt{\rho_1}}([\hat{\mu}]_{\nu\geq 1}) := \{C' \in (\mathbb{R}^{d\times d})^{\oplus d} \mid \|(0, C') - (0, [\hat{\mu}]_1, \ldots, [\hat{\mu}]_d)\|_\mathcal{V} \leq q\rho_0\sqrt{\rho_1}\}$.

Clearly $\hat{\mathfrak{B}} \subset \mathcal{V}$ compact, so that (by the extreme value theorem) we have

$$K_{\varphi|r} := \sup_{\mu, \tilde{\mu} \in B_r} |\varphi(\mu, \tilde{\mu})| \leq \sup_{(\mathfrak{q}, \tilde{\mathfrak{q}}) \in \hat{\mathfrak{B}}^{\times 2}} |\phi_\varphi(\mathfrak{q}, \tilde{\mathfrak{q}})| < \infty \quad \text{if} \quad \phi_\varphi\big|_{\hat{\mathfrak{B}}^{\times 2}} \text{ is continuous.} \tag{5.67}$$

This allows us to derive the asserted inequalities from Lemma 5.8.1, as done below.

(i) : Take any $\tilde{\mu} \in B_r$, for which then (5.143) and (5.144) (for $\mu := \hat{\mu}$ and $\eta := r$). Hence

$$\sup_{\mu \in B_r}\Big[\max_{|w|=2,3} |\langle\mu\rangle_w - \langle\hat{\mu}\rangle_w|\Big] \leq \mathfrak{m}_{\hat{\mu}|r} r \quad \text{and} \quad \sup_{\mu \in B_r}\Big[|\delta_\perp(\mu) - \delta_\perp(\hat{\mu})|\Big] \leq K_{\hat{\mu}|r}\sqrt{r + r^2}$$

for $\mathfrak{m}_{\hat{\mu}|r} := \sup_{\mu \in B_r} \mathfrak{m}_\mu$ and $K_{\hat{\mu}|r} \equiv K_{\hat{\varphi}|r} := \sup_{\mu \in B_r} |\hat{\varphi}(\mu, \hat{\mu})|$, with $\mathfrak{m}_\mu$ and $\hat{\varphi}$ as in Lemma 5.8.1 (iii). Now since – as noted at the beginning of this proof – both $\mathfrak{m}_\mu$ and $\hat{\varphi}$ factor through the map $[\,\cdot\,]$ resp. $[\,\cdot\,]^{\times 2}$ with coredinate actions $\phi_\mathfrak{m}$ resp. $\phi_{\hat{\varphi}}$, the finiteness of $\mathfrak{m}_{\hat{\mu}|r}$ and $K_{\hat{\mu}|r}$ holds by (5.67), which is applicable as the restricted actions $\phi_{\hat{\varphi}}|_{\hat{\mathfrak{B}}^{\times 2}}$ and $\phi_\mathfrak{m}|_{\hat{\mathfrak{B}}^{\times 2}}$ are both continuous; the latter follows from[22] (5.66) resp. the definition of $\mathfrak{m}_\mu$.

[22] Indeed: Clearly $\phi_{\hat{\varphi}} = \sqrt{\max\{3\phi_\vartheta, \phi_\mathfrak{m}, \phi_\mathfrak{m}^2\}\sum_{k,i,j}\phi_{\varphi_{ijk}}}$ with $\phi_{\varphi_{ijk}} \circ [\,\cdot\,]^{\times 2} = \varphi_{ijk}$, (cf. (5.131)), and the inclusion (5.66), specifically $\{[\mu]_0 \mid \mu \in B_r\} \subseteq \hat{\mathfrak{B}}_0$, implies that the functions $\phi_{\varphi_{ijk}}\big|_{\hat{\mathfrak{B}}^{\times 2}}$ are each continuous since their denominators are bounded away from zero. (To be explicit, the latter is seen from

$$\langle\mu\rangle_{ii} \geq \langle\hat{\mu}\rangle_{ii} - |\langle\hat{\mu}\rangle_{ii} - \langle\mu\rangle_{ii}| \geq \langle\hat{\mu}\rangle_{ii} - \|[\hat{\mu}]_0 - [\mu]_0\| \geq \langle\hat{\mu}\rangle_{ii} - q\rho_0 \geq (1-q)\rho_0 > 0$$

which holds for each $\mu \in B_r$ and all $i \in [d]$.) The continuity of the remaining constituents of $\phi_{\hat{\varphi}}$, i.e. ϕ_ϑ and $\phi_\mathfrak{m}$, is clear by definition of ϑ and $\mathfrak{m}$, respectively (see Lemma 5.8.1, where said definitions are given).

(ii): This assertion follows similarly from Lemma 5.8.1 (iv). Indeed: Let $\mu \in B_r$ be arbitrary. Then $\|[\mu]_0 - [\hat{\mu}]_0\| \leq \rho_1 r$, thus $\langle\mu\rangle_{ii} \geq \langle\hat{\mu}\rangle_{ii} - |\langle\hat{\mu}\rangle_{ii} - \langle\mu\rangle_{ii}| \geq \rho_0 - \rho_1 r$ and, hence,

$$\langle\tilde{\mu}\rangle_{ii} \;\geq\; \langle\mu\rangle_{ii} - |\langle\tilde{\mu}\rangle_{ii} - \langle\mu\rangle_{ii}| \;\geq\; \rho_0 - \rho_1 r - \rho_R \;>\; 0 \quad \text{for} \quad \tilde{\mu} := (I+R)_\star\mu \tag{5.68}$$

and any $i \in [d]$ (recall that the penultimate inequality holds by (5.133)). The lower bounds on the $\langle\mu\rangle_{ii}$ and $\langle\tilde{\mu}\rangle_{ii}$ prevents the functional $\varphi_R : \mu \mapsto \varphi_{\tilde{\kappa}}(\mu, (I+R)_\star\mu)\hat{\phi}_p(R)$ (with $\hat{\phi}_p$ and $\varphi_{\tilde{\kappa}}$ as in (5.134) but for $\tilde{\kappa} := 2^{\alpha+1}c_p(1+K_0) \geq \tilde{K}_{p,3}$), resp. its coredinate action ϕ_{φ_R}, from blowing up on the domain B_r, resp. on $\tilde{\mathfrak{B}} \equiv \bar{B}_{\rho_1 r}([\hat{\mu}]_0) \times \hat{\mathfrak{B}}_1 \supseteq [B_r]$. As for point (i) this ensures that the action $\phi_{\varphi_R}|_{\tilde{\mathfrak{B}}}$ is continuous, whence from (5.134) we obtain as desired that

$$\sup\nolimits_{\mu\in B_r} \left|\delta_{\perp\!\!\!\perp}((I+R)_\star\mu) - \delta_{\perp\!\!\!\perp}(\mu)\right| \;\leq\; \tilde{K}_{\hat{\mu}|r} \cdot \phi_3(\|D_R\|_\infty)^{\frac{1}{2}}(\|D_R\|_\infty\|R\|_\infty)^{\frac{\alpha}{2}}$$

for the finite – by (5.67) – constant $\tilde{K}_{\hat{\mu}|r} \equiv \tilde{K}_{\hat{\mu}|r,R} := \sup_{\mu\in B_r}|\varphi_R(\mu)| < \infty$.

For any $C, c > 0$ with $c < \rho_0$ and $\mathcal{R}_c^C$ as in (5.65), the lemma's last point follows by replacing the above $\tilde{K}_{\hat{\mu}|r}$ with $\tilde{K}_{\hat{\mu}|c,C} := \sup_{(\mu,R)\in\mathcal{R}_c^C}|\varphi_R(\mu)|$, which is clearly finite as is seen from the definition of φ_R, the (lower and upper) bounds (5.68) [and its preceding line] and (5.133). □

The following section describes the signals whose IC-defect $\delta_{\perp\!\!\!\perp}$ is minimal. These will then constitute a (large) class $\mathscr{I}_*$ of ICA-identifiable (cf. (5.10) and (5.19)) causes in $\mathfrak{C}$, see Section 5.6.

5.5.2 Orthogonal Signals

It is well-known that general signals are not identifiable from their blind linear mixtures but that this recovery becomes possible if the components of the source are mutually independent or, in the time-dependent case, at least pairwise uncorrelated. Here we show that the coredinates of such signals take a specially simple algebraic form.

Definition 5.5.8 (Product Signal)**.** An element $\mu \equiv (\mu^1, \cdots, \mu^d) \in \mathfrak{D}$ is called is called a *product signal* if it is the product of its marginals, that is if $\mu = \mu^1 \otimes \cdots \otimes \mu^d$.

Clearly, a process Y in $\mathbb{R}^d$ is a product signal iff its components $Y^1, \ldots, Y^d$ are mutually independent. The coredinates (5.39) of a product signal, if mean-stationary, are all diagonal.

Lemma 5.5.9. *For $\mu \in \mathfrak{D}$ a mean-stationary product signal, we have*

$$\langle\mu\rangle_{ij} = \delta_{ij} \cdot \langle\mu\rangle_{ij} \quad \textit{and} \quad \langle\mu\rangle_{ijk} = \delta_{ijk} \cdot \langle\mu\rangle_{ijk} \tag{5.69}$$

where δ_{ij} and δ_{ijk} are the Kronecker deltas in dimension two and three, respectively.

Definition 5.5.10. A signal $\mu \in \mathfrak{D}$ for which (5.69) holds will be called (3rd-order) *orthogonal.*

Proof of Lemma 5.5.9. Let $\mu = \mu^1 \otimes \cdots \otimes \mu^d \in \mathfrak{D}$ be a mean stationary (signed) measure on the path space $\mathcal{BV}$. The assertion (5.69) then follows from the more general statement that:

$$\langle \mu \rangle_{i_1 \dots i_m} = 0 \qquad (m \in \mathbb{N}) \tag{5.70}$$

for every multiindex $(i_1, \dots, i_m) \in [d]^{\times m}$ of such kind that at least one of its entries i_ν appears exactly once, i.e. for every multiindex contained in the set

$$J := \{(i_1, \dots, i_m) \in [d]^{\times m} \mid \exists\, \nu_0 \in [m] : i_{\nu_0} \neq i_\nu \text{ for each } \nu \neq \nu_0\}.$$

To prove (5.70), recall from Lemma 5.4.1 that each path $x \equiv (x_t^1, \dots, x_t^d)_{t \in \mathbb{I}} \in \mathcal{C}^1$ admits an integrable derivative $\dot{x} \equiv (\dot{x}^i)$ almost everywhere (set $\dot{x}$ to zero everywhere else) with

$$x_t^i = x_s^i + \int_s^t \dot{x}_r^i \,\mathrm{d}r \qquad (0 \le s \le t \le 1) \tag{5.71}$$

for all $i \in [d]$. Now by the mean-stationarity of μ we for all $i \in [d]$ have

$$0 = \mathbb{E}[\mu_{s,t}^i] = \int_{\mathcal{C}_d} x_{s,t}\, \mu^i(\mathrm{d}x) = \int_{\mathcal{C}_d} (x_t^i - x_s^i)\, \mu(\mathrm{d}x)$$

for each $0 \le s \le t \le 1$ (see Sect. 5.2.2 for notation), whence from (5.71) and Fubini we find that

$$0 = \int_s^t \phi_\mu^i(r)\,\mathrm{d}r \quad \text{ for } \quad \phi_\mu^i(r) := \int_{\mathcal{C}_d} \dot{x}_r^i\, \mu(\mathrm{d}x).$$

Since $s, t \in [0,1]$, $s \le t$, have been arbitrary, it follows

$$\phi_\mu^i = 0 \quad \text{a.e. on } \mathbb{I}, \qquad \text{for all } \; i \in [d]. \tag{5.72}$$

Now as $\phi_\mu^i(r) = \int_{\mathcal{C}_1} \dot{z}_r\, \mu^i(\mathrm{d}z)$, observe that for every $\boldsymbol{i} \equiv (i_1, \dots, i_m) \in J$ we have

$$\begin{aligned}
\langle \mu \rangle_{i_1 \dots i_m} &= \int_{\mathcal{C}_d} \xi_{i_1 \dots i_m}(x)\, \mu(\mathrm{d}x) = \int_{\Delta_m} \int_{\mathcal{C}_d} \dot{x}_{t_1}^{i_1} \cdots \dot{x}_{t_m}^{i_m}\, \mu(\mathrm{d}x)\, \mathrm{d}^m \boldsymbol{t} \\
&= \int_{\Delta_m} \int_{\mathcal{C}_d} \dot{x}_{t_1}^{i_1} \cdots \dot{x}_{t_m}^{i_m}\, \mu^1 \otimes \cdots \otimes \mu^d(\mathrm{d}(x^i))\, \mathrm{d}^m \boldsymbol{t} \\
&= \int_{\Delta_m} \left[\int_{\mathcal{C}_1} \dot{x}_{t_{\nu_0}}^{i_{\nu_0}}\, \mu^{i_{\nu_0}}(\mathrm{d}x^{i_{\nu_0}}) \cdot \prod_{i \in [d] \setminus \{i_{\nu_0}\}} \int_{\mathcal{C}_1} z^{\boldsymbol{q}_i(\boldsymbol{t})}\, \mu^i(\mathrm{d}z) \right] \mathrm{d}^m \boldsymbol{t} \\
&= \int_{\Delta_m} \phi_\mu^{i_{\nu_0}}(t_{i_{\nu_0}}) \cdot Q_\mu^{\boldsymbol{i}}(\boldsymbol{t})\, \mathrm{d}^m \boldsymbol{t},
\end{aligned}$$

where the integral in the last line can be written as, for $\boldsymbol{t}' := (t_2, \dots, t_m)$,

$$\begin{aligned}
\int_{\Delta_m} \phi_\mu^{i_{\nu_0}}(t_{i_{\nu_0}}) \cdot Q_\mu^{\boldsymbol{i}}(\boldsymbol{t})\, \mathrm{d}^m \boldsymbol{t} &= \int_{\mathbb{I}^{d-1}} \int_{\mathbb{I}} \phi_\mu^1(t_1) Q_\mu^{\boldsymbol{i}}(\boldsymbol{t}') \cdot \mathbb{1}_{\Delta_m}(\boldsymbol{t})\, \mathrm{d}t_1 \mathrm{d}^{m-1} \boldsymbol{t}' \\
&= \int_{\mathbb{I}^{d-1}} Q_\mu^{\boldsymbol{i}}(\boldsymbol{t}') \left[\int_{\mathbb{I}} \phi_\mu^1(t_1) \cdot \mathbb{1}_{\Delta_m}(\boldsymbol{t})\, \mathrm{d}t_1 \right] \mathrm{d}^{m-1} \boldsymbol{t}' = 0 \qquad \text{(by (5.72))},
\end{aligned}$$

where for convenience we assumed $i_{\nu_0} = 1$ without loss of generality, otherwise permuting $\Delta_m := \{(t_1, \cdots, t_m) \mid 0 \leq t_1 \leq t_2 \leq \ldots \leq t_m \leq 1\}$ accordingly. For each $i \in [d]$, the product $Q^{\boldsymbol{i}}_\mu(\boldsymbol{t}) := \prod_{i\in[d]\setminus\{i_{\nu_0}\}} \int_{\mathcal{C}_1} z^{\boldsymbol{q}_i(\boldsymbol{t})}\, \mu^i(\mathrm{d}z)$ involves the factors

$$z^{\boldsymbol{q}_i(\boldsymbol{t})} := \prod\nolimits_{\nu\in[m]\,:\,i_\nu=i} \dot{z}^{i_\nu}_{t_\nu} \qquad (z \in \mathcal{C}_1)$$

where $z^{\boldsymbol{q}_i(\boldsymbol{t})} := 1$ if the product is empty. This proves (5.70) and hence Lemma 5.5.9. □

Remark 5.5.11. Given the fact (verifiable by a direct computation) that the signature moments (5.38) of a mean-stationary product signal μ coincide with its signature cumulants [19] up to order $m = 3$, the assertion of Lemma 5.5.9 can also be derived from [19, Theorem 1.2].

In the following subsection we introduce an identifiability space E, as outlined in Section 5.3.2, that is endowed with a (coarse) extension of the above signature topology on $\mathfrak{D}$.

5.5.3 A Coarse Topology on the Space of Causes

This section extends the signal topology of Section 5.5.1 to an 'identifiability' topology on the causal space $\mathfrak{C}$, based on which the robustness property (5.24) can then be appropriately analysed.

Exploiting the premetric structure τ_δ from Sect. 5.5.1 and for technical convenience, said topology will be supported on 'regular' subspaces E of $\mathfrak{C}$ of the product form

$$E \;=\; \mathfrak{S} \times C^{1,1}(\mathbb{R}^d) \qquad \text{with} \qquad \mathfrak{S} \subseteq \dot{\mathfrak{D}} \text{ core-integrable} \tag{5.73}$$

as per $(K_\mathfrak{S}, \beta, p, \alpha) =: c_\mathfrak{S}$, see Subsection 5.4.4.1 for notation.

The factor $\mathfrak{S}$ is endowed with the τ_δ-induced subspace topology τ'_δ on $\mathfrak{S}$, i.e. $\tau'_\delta = \{\mathcal{O} \cap \mathfrak{S} \mid \mathcal{O} \in \tau_\delta\}$,[23] while $C^{1,1}(\mathbb{R}^d) =: C^{1,1}$ is assigned the premetric $\bar{d} : C^{1,1} \times C^{1,1} \to \mathbb{R}_+$ given by

$$\bar{d}(\theta_1, \theta_2) := \min\Big((\|\theta_1^{-1} \circ \theta_2 - \mathrm{id}\|_\infty + 1) \cdot \|D(\theta_1^{-1} \circ \theta_2 - \mathrm{id})\|_\infty,\, 1 \Big) \tag{5.74}$$

which measures the residual deviation between two transformations, in adaptation of the classical (pseudo)norm $\min(\|\cdot\|_\infty + \|D(\,\cdot\,)\|_\infty, 1)$ on $C^1(\mathbb{R}^d; \mathbb{R}^d)$. Both δ and $\bar{d}$ combine to a premetric $\mathbb{d} : E^{\times 2} \to \mathbb{R}_+$ on the subspace (5.73), which we define by

$$\mathbb{d}((\mu_1, \theta_1), (\mu_2, \theta_2)) := \delta(\mu_1, \mu_2) \,+\, \bar{d}(\theta_1, \theta_2). \tag{5.75}$$

[23] Clearly then, τ'_δ is the topology induced by the restriction of δ to the subset $\mathfrak{S} \times \mathfrak{S}$, which is again a premetric on $\mathfrak{S}$. (We use the same symbol for this restriction and the original premetric (5.55), for convenience.)

Being the sum of the premetrics (5.55) and (5.74), the premetric[24] $\mathbb{d}$ combines the marginal toplogies $(\mathfrak{S}, \delta)$ and $(C^{1,1}(\mathbb{R}^d), \bar{d})$ to a coarse joint topology $\tau_{\mathbb{d}}$ on E. The lemma below shows how this topology relates to the 'identifiability controlling' defect $\delta_{\perp\!\!\!\perp}$ of a signal.

The premetric space $(E, \mathbb{d})$ is the domain on which we will show and analyse the robustness (5.24). (Accordingly, our corresponding IA $\mathscr{I}$ will then be a subset of E.) To stay consistent with the setup of Sect. 5.3.2, we may (trivially) complement $\tau_{\mathbb{d}}$ into a topology on $\mathfrak{C}$ as follows.

Definition 5.5.12 (Causal Topology)**.** Given E as in (5.73) and $E^{\mathrm{c}} := \mathfrak{C} \setminus E$, we call

$$\text{the disjoint union topology} \quad \tau_{\mathfrak{C}} \quad \text{on} \quad (E, \tau_{\mathbb{d}}) \sqcup (E^{\mathrm{c}}, \tau_{\emptyset}) \tag{5.76}$$

the ICA-supporting *causal topology* on $\mathfrak{C}$. (Here, $\tau_{\emptyset} = \{\emptyset, E^{\mathrm{c}}\}$ is the trivial topology on E^{c}.)

Remark 5.5.13. By definition, the topology $\tau_{\mathfrak{C}}$ on $\mathfrak{C}$ is simply the disjoint union $\tau_{\mathfrak{C}} = \tau_{\mathbb{d}} \sqcup \tau_{\emptyset}$. Hence and by construction of (5.75), the topology $\tau_{\mathfrak{C}}$ relates to the referential coarseness (5.25) via Proposition 5.5.4. Consequently still and by Remark 5.5.3, a map $\phi : (\mathfrak{C}, \tau_{\mathfrak{C}}) \to (W, d_W)$ [the latter a metric space] is seen to be continuous at a point $\mathfrak{c}_\star \in E \subseteq \mathfrak{C}$ if for every $\varepsilon > 0$ there is a $\tilde{\delta} > 0$ such that $d_W(\phi(\mathfrak{c}), \phi(\mathfrak{c}_\star)) < \varepsilon$ for each $\mathfrak{c} \in \{\tilde{\mathfrak{c}} \in E \mid \mathbb{d}(\mathfrak{c}_\star, \tilde{\mathfrak{c}}) < \tilde{\delta}\}$. ◆

Lemma 5.5.14. *For a fixed orthogonal signal $\mu_\star \in E$ and $A \in \mathrm{GL}_d$, consider the $\mathbb{d}$-ball*

$$\mathbb{B}_r^\star \equiv \mathbb{B}_r(\mu_\star, A) := \{\mathfrak{p} \in E \mid \mathbb{d}((\mu_\star, A), \mathfrak{p}) < r\} \tag{5.77}$$

for $r > 0$. Let further $\rho_0 := \min_{i\in[d]} \langle \mu_\star \rangle_{ii}$ and $\rho_1 := \max_{i\in[d]} \langle \mu_\star \rangle_{ii}$. Provided $r < r_0 := \big(\rho_0/(\rho_1 + 2C_p)\big)^{1/\alpha}$, there are then explicit constants $\hat{K} = \hat{K}(\mu_\star, C_p) \geq 1$ and $K_r = K_r(\mu_\star, r, C_p)$, increasing in r, such that for any given $(\mu, f) \in \mathbb{B}_r^\star$ there exists a unique $\tilde{\mu} \in \mathfrak{D}$ such that:

$$f_*\mu = A_*\tilde{\mu} \quad \textit{with} \quad \delta_{\perp\!\!\!\perp}(\tilde{\mu}) \leq K_r \cdot r^{\alpha/2} \quad \textit{and} \quad \big\|[\tilde{\mu}]_{\mathfrak{c}} - [\mu_\star]_{\mathfrak{c}}\big\|_{\mathcal{V}} \leq \hat{K} r^\alpha. \tag{5.78}$$

Proof. Fix any $(\mu, f) \in \mathbb{B}_r^\star$. Then $\delta(\mu_\star, \mu) + \bar{d}(A, f) < r$ by defs. (5.77) and (5.75). Further, $\bar{d}(A, f) \geq \|D_R\|_\infty \|R\|_\infty$ for $R := A^{-1} \circ f - \mathrm{id}$, by (5.74) (note that $r_0 < (\rho_0/\rho_1)^{1/\alpha} \leq 1$). Abbreviating $r_\delta \equiv \delta(\mu_\star, \mu)$ and $\rho_R := C_p \phi_2(\|D_R\|_\infty)\big(\|D_R\|_\infty \|R\|_\infty\big)^\alpha$, we thus find that

$$r_\delta + \|D_R\|_\infty \|R\|_\infty < r < r_0 \qquad \text{and hence} \qquad \rho_1 r_\delta + \rho_R < c_0 r^\alpha < \rho_0$$

[24] Since $(E, \mathbb{d})$ is a premetric space, Lemma 5.5.2 (i) - (iii) hold as stated upon replacing $(\dot{\mathfrak{D}}, \delta)$ with $(E, \mathbb{d})$.

for $c_0 := \rho_1 + 2C_p$ (note that $\|D_R\|_\infty \leq 1$ implies $\phi_2(\|D_R\|_\infty) \leq 2$ and $\phi_3(\|D_R\|_\infty) \leq 4$). Thus $(\mu, R) \in \mathcal{R}_c^C$, cf. (5.65), for $c := c_0 r^\alpha$ and $C := \max\{2C_p, 1\}$. Hence by Lemma 5.5.7 (ii) there is a constant $\tilde{K}_r = \tilde{K}_r(\mu_\star, r, K_{\mathfrak{D}}) \geq 0$, independent of μ and r_δ and R, such that

$$|\delta_{\perp\!\!\!\perp}((I+R)_*\mu) - \delta_{\perp\!\!\!\perp}(\mu)| \;\leq\; \tilde{K}_r \cdot r_R^{\alpha/2} \tag{5.79}$$

for $r_R := \|D_R\|_\infty \|R\|_\infty$. Denoting $\tilde{\mu} := (\mathrm{I} + R)_*\mu$, we also find that

$$f_*\mu = [A(I+R)]_*\mu = A_*\tilde{\mu} \qquad \text{(by definition of } R\text{)}$$

and, by combination of (5.79) with Lemma 5.5.7 (i), using $\delta_{\perp\!\!\!\perp}(\mu_\star) = 0$ (cf. Lemma 5.5.9),

$$|\delta_{\perp\!\!\!\perp}(\tilde{\mu})| \;\leq\; |\delta_{\perp\!\!\!\perp}(\tilde{\mu}) - \delta_{\perp\!\!\!\perp}(\mu)| + |\delta_{\perp\!\!\!\perp}(\mu) - \delta_{\perp\!\!\!\perp}(\mu_\star)| \;\leq\; K_r' \left(\sqrt{r_R^\alpha} + \sqrt{r_\delta}\right) \;\leq\; K_r \cdot r^{\alpha/2} \tag{5.80}$$

for the constants $K_r' := \max\{\tilde{K}_r, \sqrt{2}K_{\mu_\star|r}\}$, for $K_{\mu_\star|r}$ as in (5.63), and for $K_r := 2K_r'$.

As to the distance between $[\tilde{\mu}]_{\mathfrak{c}}$ and $[\mu_\star]_{\mathfrak{c}}$, note that by (5.133) and (5.132),

$$\begin{aligned} &\big\|[\tilde{\mu}]_{\mathfrak{c}} - [\mu_\star]_{\mathfrak{c}}\big\|_{\mathcal{V}} \;\leq\; \big\|[\tilde{\mu}]_{\mathfrak{c}} - [\mu]_{\mathfrak{c}}\big\|_{\mathcal{V}} + \big\|[\mu]_{\mathfrak{c}} - [\mu_\star]_{\mathfrak{c}}\big\|_{\mathcal{V}} \\ = &\sqrt{\textstyle\sum_{|w|=2,3} \Delta_w(\tilde{\mu}, \mu)^2} + \sqrt{\textstyle\sum_{|w|=2,3} \Delta_w(\mu, \mu_\star)^2} \;\leq\; c_d[5C_p r_R^\alpha + \mathfrak{m}_{\mu_\star} r_\delta] \;\leq\; \hat{K} r^\alpha \end{aligned} \tag{5.81}$$

where we denoted $\Delta_w(\mu^{(1)}, \mu^{(2)}) := |\langle \mu^{(1)} \rangle_w - \langle \mu^{(2)} \rangle_w|$ (so that the identity holds by definition of $\|\cdot\|_{\mathcal{V}}$) as well as $c_d := (\sum_{|w|=2,3} 1)^{1/2} = d\sqrt{d+1}$ and $\hat{K} := c_d(5C_p + \mathfrak{m}_{\mu_\star})$. □

Inherent to the definition (5.76) of the causal topology $\tau_{\mathfrak{C}}$ is an element of choice pertaining to the selection of $\mathfrak{S}$ in (5.73): That we resort to an *integrable* subset $\mathfrak{S}$ of $\dot{\mathfrak{D}}$ is to ensure uniform comparability of different nonlinearly transformed signals under the residual metric (5.74), cf. Section 5.4.4.1 and (5.79),(5.80); the larger the enclosed threshold (5.53), the more 'fine grained' will be the $\tau_{\mathfrak{C}}$-supported variation of (5.24) on $\mathfrak{C}$. The 'necessity of choice' behind (5.73) can be avoided by replacing (5.74) with a 'non-incremental' (e.g., quotient) [pre]metric on $C^{1,1}$ or by restricting the latter space to a 'tamed' subspace of transformations (so as to arrive at a μ-uniformly integrable bound on the RHS of (5.148)). No integrability is needed (i.e., $\mathfrak{S} = \dot{\mathfrak{D}}$ works out) if only variations in the source (and not in its transformation) are considered.

Remark 5.5.15. Taking $\mathfrak{S} = \dot{\mathfrak{D}}$, the proof of Lemma 5.5.14 shows that on the A-sections

$$E_{A,r} := (\dot{\mathfrak{D}} \times \{A\}) \cap \mathbb{B}_r(\mu_\star, A) = B_\delta(\mu_\star, r) \times \{A\},$$

i.e. when the hidden relation A between source and observable is linear and the same for all signals so that any inexactness in the observable-to-source inversion (5.24) is due to source deviations (from $\mu_\star$) only, then no integrability assumptions on $\mathfrak{S}$ of the form (5.53) are required and the estimates (5.78) improve to:

Lemma. *Let* $\mu_\star \in \dot{\mathfrak{D}}$ *be orthogonal and* $A \in \mathrm{GL}_d$*, and let* $0 < r < \rho_0/\rho_1$*. Then there are constants* $\hat{L} = \hat{L}(\mu_\star)$ *and* $L_r = L_r(\mu_\star, r)$*, the latter increasing in* r*, with*

$$\delta_{\perp\!\!\!\perp}(\mu) \leq L_r\sqrt{r} \quad \text{and} \quad \big\|[\mu]_{\mathfrak{c}} - [\mu_\star]_{\mathfrak{c}}\big\|_{\mathcal{V}} \leq \hat{L}r \quad \text{for each } \mu \in B_\delta(\mu_\star, r). \tag{5.82}$$

Proof. Revisiting the proof of Lemma 5.5.14 on $E_{A,r}$, we have $R \equiv 0$ and hence note (5.82) for the constants $L_r := \sqrt{2}K_{\mu_\star|r}$ via (5.80), and $\hat{L} := \mathfrak{m}_{\mu_\star}$ via (5.81). □

The above simplification is useful in many applications, as the relation between source and observable can often be modelled as exactly linear but with the underlying source signal deviating from orthogonality; see for instance Section 5.7 for a few examples. ◆

5.6 Robust Independent Component Analysis

We introduce a quantifiably robust statistical procedure to recover (the inverse of) a matrix from its action on an unobserved [non]orthogonal source signal. The general strategy for this is based on the classical idea of jointly diagonalising a set of derived equivariant matrix statistics, given in our case by the coredinates (5.39) of the source signal (Section 5.6.1). By linking our inversion procedure to the causal topology from Definition 5.5.12, we obtain both an identifiability theorem which provides explicit stability bounds for the recovery of signals that may deviate from orthogonality in any way (Theorem 5.6.10), and, as our main result, a readily implementable new ICA-solution which is provably robust in the sense of Definition 5.3.10 with strong and general robustness guarantees that can be expressed as explicitly computable and naturally interpretable moduli of continuity (Theorem 5.6.14).

5.6.1 ICA-Inversion from Coredinates

Throughout this section, we follow the classical ICA-paradigm (5.18) and consider two linearly related signals ζ and χ in $\dot{\mathfrak{D}}$, that is

$$\chi = A \cdot \zeta \quad \text{with} \quad A \equiv (a_{ij}) \equiv (a_1|\cdots|a_d) \in \mathbb{R}^{d\times d} \tag{5.83}$$

some fixed invertible $d \times d$ matrix with columns a_i. For instance, we can always think of

$$\zeta = \mathbb{P}_S \;\text{ and }\; \chi = \mathbb{P}_X \quad \text{for a BSS-triple } (X, S, A) \text{ as in (5.8)} \tag{5.84}$$

with $S = (S_t)$ and $X = (X_t)$ two stochastic processes in $\mathbb{R}^d$.

The Problem of ICA (Def. 5.3.7) is then to recover the inverse A^{-1} from the input χ, that is: to find a map $\hat{\Phi}$ on $\dot{\mathfrak{D}}$ such that $\hat{\Phi}(\chi) \doteq A^{-1}$ up to some minimal ambiguity, see (5.19).

In principle, one may expect such a map $\hat{\Phi}$ to operate on certain 'pieces of information' on χ that efficiently relate the action of A to some relevant statistical properties of ζ, cf. (5.10).

The following ensures that the coredinates $[\chi]_0, \ldots, [\chi]_d$ of χ can be used to this effect.

Proposition 5.6.1. *Let χ and ζ be as in (5.83) with coredinates $[\chi]_{\mathfrak{c}}$ and $[\zeta]_{\mathfrak{c}}$, respectively. Provided that $[\zeta]_0$ is invertible, we have for the inverse $\theta := A^{-1}$ that*

$$\theta \cdot [\chi]_0 \cdot \theta^\intercal = [\zeta]_0 \quad \text{and} \quad \theta \cdot [\chi]_\nu \cdot \theta^\intercal = \sum_{\ell=1}^{d} a_{\nu\ell}[\zeta]_\ell \tag{5.85}$$

for each $\nu \in [d]$, and for each $c \equiv (c_0, \underline{c}) \in \mathbb{R}^{1+d}$ with $\underline{c} \equiv (\tilde{c}_1, \ldots, \tilde{c}_d)$ find that

$$\theta^{-1} \cdot [\chi]^c_\odot \cdot \theta \;=\; [\zeta]^{A|c}_\odot \qquad \text{for the matrices} \tag{5.86}$$

$$[\chi]^c_\odot \;:=\; c_0[\chi]_0^{-1} \cdot \sum_{\nu=1}^{d} \tilde{c}_\nu [\chi]_\nu \quad \text{and} \quad [\zeta]^{A|c}_\odot \;:=\; c_0[\zeta]_0^{-1} \cdot \sum_{\ell=1}^{d} \langle a_\ell, \underline{c}\rangle_2 [\zeta]_\ell.$$

If ζ is 3rd-order orthogonal, then $[\zeta]^{A|c}_\odot$ is diagonal.

Proof. This is a consequence of Lemmas 5.4.7 and Lemma 5.5.9. Indeed: Applied to the identity $\chi = A \cdot \zeta$ and written out at order $m = 2, 3$, the above equivariance of the signature moments reads

$$\begin{aligned} \langle\chi\rangle_{ij} &= \textstyle\sum_{\alpha,\beta=1}^{d} a_{i\alpha}a_{j\beta} \cdot \langle\zeta\rangle_{\alpha\beta} && ((i,j) \in [d]^2) \qquad \text{and} \\ \langle\chi\rangle_{ijk} &= \textstyle\sum_{\alpha,\beta,\gamma=1}^{d} a_{i\alpha}a_{j\beta}a_{k\gamma} \cdot \langle\zeta\rangle_{\alpha\beta\gamma} && ((i,j,k) \in [d]^3) \end{aligned} \tag{5.87}$$

which, recalling (5.39), is a system of equations equivalent to the matrix congruences

$$[\chi]_0 = A[\zeta]_0 A^\intercal \quad \text{and} \quad [\chi]_k = A\Big[\textstyle\sum_{\gamma=1}^{d} a_{k\gamma}[\zeta]_\gamma\Big]A^\intercal \tag{5.88}$$

for each $k \in [d]$. The identities (5.85) follow. As to the conjugacy (5.86), this holds by (5.88) upon noting that $[\chi]^c_\odot = c_0[\chi]_0^{-1} \cdot [\chi]_{\underline{c}}$ for the linear combination

$$[\chi]_{\underline{c}} \;\equiv\; \textstyle\sum_{\nu=1}^{d} \tilde{c}_\nu [\chi]_\nu \;=\; A\Big[\sum_{\gamma=1}^{d}\big(\sum_{\nu=1}^{d} a_{\nu\gamma}c_\nu\big)[\zeta]_\gamma\Big]A^\intercal \;=\; A\Big[\sum_{\gamma=1}^{d}(a_\gamma \cdot \underline{c})[\zeta]_\gamma\Big]A^\intercal$$

as this implies that

$$[\chi]^c_\odot \;=\; A^{-\intercal}c_0[\zeta]_0^{-1}A^{-1} \cdot A\Big[\textstyle\sum_{\gamma=1}^{d}(a_\gamma \cdot \underline{c})[\zeta]_\gamma\Big]A^\intercal \;=\; A^{-\intercal} \cdot [\zeta]^{A|c}_\odot \cdot A^\intercal.$$

The asserted diagonality of $[\zeta]^{A|c}_\odot$ is clear by Lemma 5.5.9. □

Remark 5.6.2. The main takeaway from Proposition 5.6.1 is how the congruences (5.85) between the 'input' and the 'target' statistics $[\chi]_\nu$ and $[\zeta]_\nu$ combine to the A-carried conjugacy (5.86) between the (c-parametrized families of) contracted statistics $[\chi]^c_\odot$ and $[\zeta]^{A|c}_\odot$. This contraction of the coredinates $[\chi]_{\mathfrak{c}}$, $[\zeta]_{\mathfrak{c}}$ opens up additional degrees of freedom, parametrised by c, that can be conveniently exploited to, via (5.86), characterise A^{-1} up to monomial ambiguity if the eigenspectrum of $[\zeta]_{\odot}^{A|c}$ is non-degenerate. Said spectrum, in turn, can be controlled by the amount of statistical dependence between the components of ζ, cf. Lem. 5.5.9 and (5.56), provided that these components satisfy some mild statistical non-degeneracy condition, cf. e.g. (5.103). ◆

The strategy of Rem. 5.6.2 underlies the proof of Theorem 5.6.10 below. Its implementation, i.e. the construction of an ICA-inversion (5.19) from Proposition 5.6.1, is next.

5.6.1.1 Auxiliary and Core Statistics

For a given $\mu \in \mathfrak{D}$ we define $C_\mu := \frac{1}{2}([\mu]_0 + [\mu]_0^\intercal) \in \mathbb{R}^{d\times d}$, which is symmetric and positive semidefinite. In the context of (5.83), suppose that

$$C_\chi \in \mathrm{GL}_d \quad \text{and} \quad R \equiv R_{(\chi)} := C_\chi^{-1/2} \text{ is the square root of the inverse } C_\chi^{-1}. \tag{5.89}$$

The matrix R exists and is unique, e.g. [87, Korollar 1.10], and, by its definition, $R^2 = C_\chi^{-1}$ and hence $C_{R\cdot\chi} = RC_\chi R^\intercal = \mathrm{I}$. This latter identity is the central purpose of R, and its choice as an inverse square root (5.89) is also known in the literature as Mahalanobis (or ZCA) whitening.

Remark 5.6.3. From the spectral theorem we know that

$$R_{(\chi)} \equiv \mathcal{R}(C_\chi) \;=\; Q\Lambda_\chi^{-1/2}Q^\intercal \tag{5.90}$$

for some $Q = Q(C_\chi) \in \mathrm{O}_d(\mathbb{R})$ and $\Lambda_\chi := \mathrm{ddiag}(\lambda_1^\chi, \ldots, \lambda_d^\chi)$, where $\lambda_1^\chi \geq \ldots \geq \lambda_d^\chi > 0$ are the eigenvalues of C_χ (enumerated in descending order, counting multiplicities). Further

$$C_\chi \;=\; A\big[\tfrac{1}{2}([\zeta]_0 + [\zeta]_0^\intercal)\big]A^\intercal \;=\; AC_\zeta A^\intercal, \tag{5.91}$$

so that C_χ is invertible iff C_ζ is invertible; the latter is violated only for the degenerate case that the increment $\zeta_{0,1}$ is supported on a hyperplane of $\mathbb{R}^d$ (Lemma 5.8.2). ◆

For usage below, consider the pre-transformed matrices $A_R := RA$ and

$$B_R \equiv (b_1|\cdots|b_d)^\intercal := (A_R)^{-1}, \;\text{ and }\; \bar{B}_R := \mathrm{ddiag}(|b_1|, \cdots, |b_d|)^{-1}B_R \tag{5.92}$$

for the matrix B_R rescaled to unit rows.

The central components for our inversion procedure are the normalised statistics

$$\mathfrak{x}_0(\theta) := N_{\theta\cdot\chi}^{-1} \cdot [\theta\cdot\chi]_0 \cdot N_{\theta\cdot\chi}^{-1} \quad \text{and} \quad \mathfrak{x}_\nu(\theta) := N_{\theta\cdot\chi}^{-1} \cdot \frac{[\theta\cdot\chi]_\nu}{\sqrt{\langle\theta\cdot\chi\rangle_{\nu\nu}}} \cdot N_{\theta\cdot\chi}^{-1} \quad (\nu \in [d]) \tag{5.93}$$

for $\theta \in \mathrm{GL}_d =: \Theta$ and with $N_\mu = \mathrm{ddiag}(\langle\mu\rangle_{11}, \cdots, \langle\mu\rangle_{dd})^{1/2}$ for each $\mu \in \mathfrak{D}$, as before.

Remark 5.6.4. The statistics in (5.93) are all well-defined. Indeed, since C_χ is invertible by assumption and hence $C_\chi \in \mathrm{Sym}_d^+$, we have $C_{\theta\cdot\chi} = \theta C_\chi \theta^\intercal \in \mathrm{Sym}_d^+$ and thus $\langle \theta \cdot \chi \rangle_{ii} = e_i^\intercal C_{\theta\cdot\chi} e_i > 0$ for each $i \in [d]$ and all $\theta \in \Theta$.

Lemma 5.6.5. *The statistics* $\mathfrak{r}_\nu(\theta)$ *in* (5.93) *are each* scale-invariant *for any* $\chi \in \mathfrak{D}$*, i.e.*

$$\mathfrak{r}_\nu(\Lambda \cdot \theta) \;=\; \mathfrak{r}_\nu(\theta) \quad \textit{for all} \;\; \theta, \Lambda \in \mathrm{GL}_d \;\; \textit{with } \Lambda \textit{ positive diagonal.} \tag{5.94}$$

Proof. This is a direct consequence of Lemma 5.4.7. Indeed: As χ is arbitrary and the GL_d-action (5.6) leaves $\mathfrak{D}$ invariant, it suffices to show (5.94) for $\theta := \mathrm{id}$ and any $\Lambda \equiv \mathrm{ddiag}[\lambda_1, \ldots, \lambda_d]$ with $\lambda_i > 0$. Then by Lemma 5.4.7, cf. (5.87),

$$\begin{aligned} \langle \Lambda \cdot \chi \rangle_{ij} &= \lambda_i \lambda_j \textstyle\sum_{\alpha,\beta=1}^d \delta_{i\alpha}\delta_{j\beta} \langle \chi \rangle_{\alpha\beta} = \lambda_i\lambda_j \langle \chi \rangle_{ij} \\ \langle \Lambda \cdot \chi \rangle_{ijk} &= \lambda_i\lambda_j\lambda_k \textstyle\sum_{\alpha,\beta,\gamma=1}^d \delta^{ijk}_{\alpha\beta\gamma} \langle \chi \rangle_{\alpha\beta\gamma} = \lambda_i\lambda_j\lambda_k \langle \chi \rangle_{ijk} \end{aligned}$$

for $i, j, k \in [d]$. In particular, $N_{\Lambda\cdot\chi} = \Lambda \cdot N_\chi$ as well as

$$\left[\Lambda \cdot \chi\right]_0 = \Lambda [\chi]_0 \Lambda \quad \text{and} \quad \left[\Lambda \cdot \chi\right]_\nu = \lambda_\nu \Lambda [\chi]_\nu \Lambda \quad \text{for each } \nu \in [d],$$

which implies (5.94) by way of the definitions (5.93). □

5.6.1.2 Blind Inversion via Contrast Optimization

Similar in spirit to classical ICA-approaches, cf. [38, 81] and Section 5.6.2.1, we aim to recover the inverse of the hidden mixing transformation as minimizer of a specially constructed cost function: Quantifying the 'off-diagonality' of a matrix θ via $\|\theta\|_\times := \|\theta - \mathrm{ddiag}(\theta)\|$, for $\|\cdot\|$ the Frobenius norm on $\mathbb{R}^{d\times d}$, we combine the statistics (5.93) to the objective (or 'contrast') function

$$\Theta \ni \theta \;\longmapsto\; \phi_\chi(\theta) \;:=\; \sum_{\nu=0}^{d} \left\|\mathfrak{r}_\nu(\theta \cdot R)\right\|_\times^2 \tag{5.95}$$

which is minimal over the '(approximate) joint diagonalisers' of the matrices (5.93), see below. The objective function ϕ_χ is monomially invariant, which means that

$$\phi_\chi(M\theta) = \phi_\chi(\theta) \quad \text{for each } M \in \mathrm{M}_d, \;\theta \in \Theta. \tag{5.96}$$

Indeed, since each $M \in \mathrm{M}_d$ is of the form $M = P\Lambda$ for $P \in \mathrm{P}_d^\pm$ and some diagonal $\Lambda > 0$, the asserted invariance follows directly from Lemmas 5.4.7, 5.6.5 (cf. (5.88)) and the definition of the Frobenius norm. Since the desired demixing matrices will be identified as minimisers of (5.95), a suitable restriction ('compactification') of the original domain Θ will be opportune.[25]

[25] In classical ICA, where the components of the source ζ are assumed (at least pairwise) independent, such a compactification of the domain is usually reached via 'whitening' of χ; indeed, this reduces the original domain Θ to the compact subdomain O_d of orthogonal matrices. This whitening step, however, requires the intercomponental covariance structure of ζ to diagonalise, which we do not assume of our (potentially non-orthogonal) source.

We propose to consider for this the set

$$\Xi_1^0 := \{\theta \equiv (\theta_1|\cdots|\theta_d)^\intercal \in \mathbb{R}^{d\times d} \,|\, \|\theta_i\|_2 = 1, \forall\, i \in [d]\} \tag{5.97}$$

of all matrices in $\mathbb{R}^{d\times d}$ whose rows constitute a unit basis of $\mathbb{R}^d$, and define as our parameter domain for the optimisation of ϕ_χ the superset of O_d which is given by

$$\Xi_1 \;:=\; \Xi_1^0 \cap \{\theta \in \mathrm{GL}_d \;|\; \kappa_2(\theta) \le \kappa_0\} \tag{5.98}$$

for any fixed condition bound $\kappa_0 \ge \kappa_2(\bar{B}_R)$ (imposed as an a priori assumption[26]). We note that condition-restricted sets of the form (5.98) are compact.

Lemma 5.6.6. *For Ξ_1^0 as in* (5.97) *and $\gamma \ge 1$, the set*

$$\Xi_1^\gamma := \Xi_1^0 \cap \{\theta \in \Theta \mid \kappa_2(\theta) \le \gamma\} \;\; \textit{is compact in} \;\; \mathbb{R}^{d\times d}.$$

Proof. The set $\Xi_1^\gamma \subset \mathbb{R}^{d\times d}$ is bounded (by definition of Ξ_1^0), so it remains to show that it is also closed. For this, let $(\theta_n) \subset \Xi_1^\gamma$ and $\theta \in \mathbb{R}^{d\times d}$ be such that $\theta_n \to \theta$ in $\mathbb{R}^{d\times d}$. Then necessarily also $\theta \in \Xi_1^0$ and $\lim_{n\to\infty} \|\theta_n\|_2 = \|\theta\|_2 \neq 0$. The latter implies that $\|\theta_n\|_2 \ge c$ for some $c > 0$ and almost all $n \in \mathbb{N}$ (say for all $n \ge n_0$, with some $n_0 \in \mathbb{N}$). Suppose now that θ is singular, i.e. $\theta \notin \Theta$. Then and by the fact that $\|\cdot\|_2$ is submultiplicative (from which we obtain the last inequality below, see e.g. [74, Problem 5.6.P47 (p. 369)]), we must have

$$\gamma \;\ge\; \kappa_2(\theta_n) = \|\theta_n\|_2 \|\theta_n^{-1}\|_2 \;\ge\; c\|\theta_n^{-1}\|_2 \;\ge\; c/\|\theta_n - \theta\|_2 \quad \text{for all } n \ge n_0,$$

which is clearly a contradiction. This implies $\theta \in \Theta$ and hence, since $\kappa_2(\cdot)$ is continuous on Θ, also $\kappa_2(\theta) = \lim_{n\to\infty} \kappa_2(\theta_n) \le \gamma$ and thus $\theta \in \Xi_1^\gamma$, as desired. □

Since Ξ_1 is compact and ϕ_χ is continuous in θ (even continuously differentiable, cf. Lemma 5.4.7), we have

$$\mathfrak{I}_\chi := \left[\operatorname*{arg\,min}_{\theta\,\in\,\Xi_1} \phi_\chi(\theta)\right] \cdot R_{(\chi)} \;\in\; \mathfrak{F} \tag{5.99}$$

for the subset $\mathfrak{F} := \{\mathcal{M} \subseteq \mathrm{GL}_d \;|\; \mathcal{M} \text{ is non-empty and compact}\}$ of 2^{GL_d}. Recall here that

$$(\mathfrak{F}, \mathfrak{d}) \;\text{ is a metric space } \quad \text{wrt.} \quad \mathfrak{d}(\mathcal{A}, \mathcal{B}) := \max\{\sup\nolimits_{a\in\mathcal{A}} d(a, \mathcal{B}),\, \sup\nolimits_{b\in\mathcal{B}} d(\mathcal{A}, b)\}, \tag{5.100}$$

the classical Hausdorff distance on $\mathfrak{F}$, where $d(a, \mathcal{B}) := \inf_{b\in\mathcal{B}} \|b - a\|$. We set

$$\hat{\Phi} \;:\; \dot{\mathfrak{D}} \to \mathfrak{F}, \quad \mu \to \mathfrak{I}_\mu, \tag{5.101}$$

and claim that $\hat{\Phi}$ is not only an ICA-inversion (5.19) on a yet to be specified IA $\mathscr{I}_\star$, but in fact a robust such inversion in the sense of (5.15) and Def. 5.3.10. The next two sections show this.

[26] That is, the assumption that the observer does not underestimate the condition number $\kappa_2(\bar{A}_R)$.

Remark 5.6.7 (On the Optimization (5.99))**.** As further detailed in Section 5.6.2.1, the optimization of contrasts of the form (5.95) over constrained subsets of Θ such as (5.98), is an established research topic of its own with a strong and active research activity. In addition to row constraints (5.97) (e.g. [144]), the additional imposition of condition constraints (5.98) in the optimization of (5.95) (as formulated in (5.99)) is studied and free of any subtleties of its own, see e.g. [158]. ◆

Remark 5.6.8. Since the identifiability map (5.101) returns invertible transformations of its input observables, it might also appear natural to use as an alternative topology on $\mathfrak{F}$ the 'residual' semimetric which is given by (5.100) but for the alternative point-set distances

$$\tilde{d}(a,\mathcal{B}) := \inf_{b\in\mathcal{B}} \tilde{d}(a,b) \quad \text{with} \quad \tilde{d}(a,b) := \|a\circ b^{-1} - \mathrm{id}\| \vee \|b\circ a^{-1}-\mathrm{id}\| \tag{5.102}$$

for id the identity matrix in GL_d. Continuity wrt. the multiplicative setting (5.102) is implied by [continuity wrt.] (5.100), however, as the metric topology (5.100) is finer than the one induced by (5.102) as can be seen upon recalling that matrix inversion is Lipschitz at each point.

5.6.2 Blind Inversion for (Non-)Orthogonal Sources

The main goal of this section is to prove that the blind inversion of signals can be robust – or stable – under general perturbations of the source, including violations of its orthogonality ('IC') assumption. Recalling that infringements of the latter type are quantified by the defect $\delta_{\perp\!\!\!\perp}(\zeta)$ from (5.62), the following stability analysis of the relation between $\hat{\Phi}(\chi)$ and the true inverse A^{-1} (up to prefactors in M_d) is an important step towards a quantitative understanding of the desired general robustness (5.24). It involves a few constants and auxiliary functions which we introduce next.

The following considerations are about source signals in $\mathcal{M}_1$ that lie within the set

$$\mathscr{S}_0 := \{\mu\in\dot{\mathfrak{D}} \mid C_\mu\in\mathrm{GL}_d \text{ and } \sharp\{i\in[d]\,|\,\langle\mu\rangle_{iii}=0\}\leq 1\} \tag{5.103}$$

of measures μ such that C_μ is invertible (Lem. 5.8.2) and $\langle\mu\rangle_{iii}=0$ for at most one $i\in[d]$.[27]

Notation 5.6.9. For a BSS-triple $(\chi,\zeta,A)_{\mathscr{S}_0\times\mathrm{GL}_d}$, let $\gamma:=1+\sqrt{5}$ and $k_d:=\sqrt{\tfrac{1}{6}(d-1)d(2d-1)}$ and $R\equiv R_{(\chi)}$ as in (5.90), and consider the (ζ,A)-dependent $\varepsilon_0:=q_0/(1+q_0)>0$ with

$$q_0 := (\gamma k_d r_0)^{-1} \quad\text{and}\quad r_0 := \kappa_0\varsigma_1\left[\frac{\|B_R\|}{\sqrt{d}}(1+\kappa_0+(1+\xi d)\kappa_0\kappa_2(B_R))+\kappa_2(B_R)\varsigma\right] \tag{5.104}$$

[27] Note that $\langle\mu\rangle_{iii}=\mathbb{E}[(\mu_{0,1}^i)^3]$. Further, we may (wlog – upon re-enumeration of components) assume for convenience that it is $\langle\mu\rangle_{111}$ which may vanish; note that this assumption underlies the definition of ς_1.

let further $c_1 := 2dk_d r_0$ and $c_2 := \sqrt{d}\kappa_2(B_R)c_1$, and note that (5.104) involves the terms

$$\varsigma_1 := k_d^{-1}\sqrt{\sum_{i=1}^{d}\frac{(i-1)^2\langle\zeta\rangle_{ii}^3}{\langle\zeta\rangle_{iii}^2}},\quad \xi := \sqrt{\sum_{\nu=1}^{d}\|[A_R\cdot\zeta]_\nu\|^2},\quad\text{and}\quad \varsigma := \sqrt{\sum_{i=1}^{d}\langle\zeta\rangle_{iii}^2/\langle\zeta\rangle_{ii}^3},$$

cf. (5.38), (5.89). In the following, these constants should not be taken all too seriously with regards to optimality as we did not attempt to optimize the sharpness of our inequalities.

The following theorem is a cornerstone for our main robustness result (Theorem 5.6.14).

Theorem 5.6.10 (Blind Inversion)**.** *Let (χ, ζ, A) be a BSS-triple on $\mathscr{S}_0 \times \mathrm{GL}_d$ and $\hat{\Phi}$ as in* (5.101)*. If $\delta := \delta_{\perp\!\!\!\perp}(\zeta) \leq \varepsilon_0$, then for each $\theta_\star \in \hat{\Phi}(\chi)$ there is $M \in \mathrm{M}_d$ and $E \in \mathbb{R}^{d\times d}$ such that*

$$\theta_\star = M(\mathrm{I}+E)A^{-1} \quad \textit{with} \quad \|E\| \leq c_1\tfrac{\delta}{1-\delta} \quad \textit{and} \quad \tfrac{\|\theta_\star - MA^{-1}\|}{\|MA^{-1}\|} \leq c_2\tfrac{\delta}{1-\delta}\,. \tag{5.105}$$

Moreover, the deviance (5.20) *satisfies $\partial_{\hat{\Phi}}(\zeta, A) \leq c_2\frac{\delta}{1-\delta}$, and $(\mathscr{I}_\star, \hat{\Phi})$ is an ICA-solution for*

$$\mathscr{I}_\star := \mathscr{S}_\star \times \mathrm{GL}_d \quad \textit{with} \quad \mathscr{S}_\star := \{\mu \in \mathscr{S}_0 \mid \mu \textit{ is orthogonal }\}. \tag{5.106}$$

Proof. The proof of (5.105) combines standard arguments from matrix analysis with recent perturbation bounds for eigenspaces of (nonsymmetric) matrices; it is rather lengthy and hence deferred to Section 5.8.3. The remaining claims are corollaries of this proof and (5.105):

Statement (5.105) implies that for each $B \in \hat{\Phi}(\chi)$ there is $M \in \mathrm{M}_d$ and $E \in \mathbb{R}^{d\times d}$ with

$$\tfrac{|BAu - Mu|}{|Mu|} = \tfrac{|M\cdot E\cdot M^{-1}(Mu)|}{|Mu|} \leq \|M\cdot E\cdot M^{-1}\|_2 \leq c_2\tfrac{\delta}{1-\delta}, \quad \text{for each} \quad u \in \mathbb{R}^d\setminus\{0\}, \tag{5.107}$$

where the last inequality holds by the same argument as for (5.204). This together with the definition (5.20) of the deviance function $\partial_{\hat{\Phi}}$ then implies the asserted $\partial_{\hat{\Phi}}(\zeta, A)$-inequality.

The $\lceil\cdot\rceil_{\mathfrak{m}}$-sufficiency of (5.106) follows from (5.105) and the fact that $\delta_{\perp\!\!\!\perp}(\mathscr{S}_\star) = 0$, cf. (5.62). Indeed: For any BSS-triple $(\chi, \zeta, A)_{\mathscr{I}_\star}$ we have, in the notation of (5.161), that $(\chi, \tilde{\zeta}, \tilde{A})_{\mathscr{I}_\star}$ and hence $\mathrm{I} = RC_\chi R^\intercal = (R\tilde{A})C_{\tilde{\zeta}}(R\tilde{A})^\intercal = \tilde{A}_R\tilde{A}_R^\intercal$ and thus $\tilde{B}_R := \tilde{A}_R^{-1} \in \mathrm{O}_d \subseteq \Xi_1$ (cf. (5.98)) for each (Ξ_1-defining choice of) $\kappa_0 \geq 1$. Statement (5.105) thus applies to (the triple $(\chi, \tilde{\zeta}, \tilde{A})_{\mathscr{I}_\star}$ and) the given map $\hat{\Phi}$ in particular, implying $\hat{\Phi}(\chi) \subseteq \mathrm{M}_d\cdot\tilde{A}^{-1} = \mathrm{M}_d\cdot A^{-1}$ as desired. □

As is to be expected, the deviation (5.105) between $\mathrm{M}_d\cdot A^{-1}$ and the recovered inverses $\theta_\star \in \hat{\Phi}(\chi)$ computed from χ — and likewise the deviation $\mathrm{dist}_{\tilde{\varrho}}(\hat{\Phi}(X)\cdot X, \lceil S\rceil_{\mathfrak{m}})$, via (5.21), between the accurate quasi sources and their estimates — are expressed on a multiplicative scale [resp. in terms of relative error] rather than on an additive scale [resp. in absolute error], owing to the multiplicative nature of the group GL_d and its action on $\mathcal{M}_1$.

We note further that the proof of (5.105) (Section 5.8) allows to also quantify the breakdown point of the blind inversion procedure (5.101) as being reached at precisely those causes (ζ, A) in $\mathfrak{C}$ for which one of the eigenvalues $(\tilde{\lambda}_i)$ in (5.196) attains an algebraic multiplicity of strictly above one. (We may interpret these causes as the 'singularities' of $\hat{\Phi}$ in the causal space $\mathfrak{C}$.)

Remark 5.6.11. Let us reconsider the usual case (5.84) where ζ and χ are (the laws of) some $\mathbb{R}^d$-valued stochastic processes $S = (S_t)_{t\in\mathbb{I}}$ and $X = (X_t)_{t\in\mathbb{I}}$ (resp.) defined on some probability space $(\Omega, \mathscr{F}, \mathbb{P})$. Denoting $\hat{S} := \theta_\star \cdot X$ for any $\theta_\star \in \hat{\Phi}(X)$, Theorem 5.6.10 and (5.21) then yield[28]

$$\tilde{\varrho}(\hat{S}, MS) \equiv \sup_{t\in\mathbb{I}} \tfrac{|\hat{S}_t - MS_t|}{|MS_t|} \leq c_2\,\hat{\delta} \quad \text{for some} \quad M = M(\theta_\star) \in \mathrm{M}_d \tag{5.108}$$

and with $\hat{\delta} := \frac{\delta_{\perp\!\!\!\perp}(S)}{(1-\delta_{\perp\!\!\!\perp}(S))}$ and $c_2 \equiv c_2(S, A)$ as above, provided that $\delta_{\perp\!\!\!\perp}(S) \leq \varepsilon_0 \equiv \varepsilon_0(S, A)$. If in addition to the 'global' deviation (5.108) one is also interested in the reconstruction error for an individual component $S^i = (S^i_t)_{t\in\mathbb{I}}$ of the source ($i \in [d]$), then one may consider the proportionality factor $p^i_t := |\hat{S}_t|/|\hat{S}^i_t| \cdot \mathbb{1}_\times(\hat{S}^i_t)$ and note that by (5.108), with $M =: (\beta_i \cdot \delta_{\sigma(i)j})_{ij}$,

$$\frac{\left|\hat{S}^i_t - \beta_i S^{\sigma(i)}_t\right|}{\left|\beta_i S^{\sigma(i)}_t\right|} \cdot \imath \leq \frac{p^i_t \cdot \delta'}{1 - p^i_t \cdot \delta'} \qquad \text{for each} \quad t \in \mathbb{I} \quad \text{with} \quad (p^i_t \cdot \delta') \vee \tilde{\delta} < 1,$$

where $\imath := \mathbb{1}_\times(\hat{S}^i_t)\mathbb{1}_\times(S^{\sigma(i)}_t)$ and $\delta' := \tilde{\delta}/(1-\tilde{\delta})$ and $\tilde{\delta} := c_2\hat{\delta}$; see e.g. [59, Prop. 1.12]. ◆

5.6.2.1 Related Work

As mentioned earlier, the general strategy to perform a blind inversion of linearly mixed stochastic sources via a joint diagonalisation of suitably chosen equivariant tensor statistics of the observed signal – as done in (5.99) – is a well-established part of the current theory of blind source separation, see e.g. [25, 28, 151] or [38, 111, 117, 125, 150, 151] for an overview. Its relevance in the BSS-context has brought this joint diagonalisation problem (JDP) – which underlies the minimisation (5.99) of the contrast (5.95) over the manifold Θ and over certain subsets such as (5.98) – significant popularity and much research attention in numerical analysis and optimization communities, cf. the referenced literature and the references therein.

[28] Notice that if $|M \cdot S_t| = 0$ then $\hat{S}_t = \theta_\star A S_t = 0$, and for such $t \in \mathbb{I}$ the quotient in (5.108) is set to $0/0 := 0$.

Yet despite the manifest relevance of this actively researched problem, the perturbation analysis of JDP, just as related sensitivity analyses for tensor decompositions in general, has been notoriously underexplored, cf. [24, Sect. 1.4]. Among the very few available perturbation results, approaches that anticipate our own bound (5.105) most closely are [1], who obtained first-order perturbation bounds for joint diagonalizers of symmetrically perturbed symmetric matrices, and [144], who under the same row-normalisation constraints as ours, but once again in an entirely symmetric setting only, obtained bounds between symmetrically perturbed and exact joint diagonalizers, and finally [24], who obtained upper bounds related to ours in the more general context of perturbed joint block diagonalization. For completeness, we also mention the recent but not directly related work [100], where again under the symmetry constraint some numerical stability properties of joint approximate diagonalizers are studied.

We emphasize that while of a similar nature, none of the above works renders our bound (5.105) or its underlying perturbation analysis (Section 5.8.3) redundant: The perturbation results of Theorem 5.6.10 are a cornerstone for our proof of Theorem 5.6.14 where the robustness of (5.101) is established via the continuity of maps between premetric spaces, and the latter requires an understanding of how general and thus also asymmetric perturbations of the statistics (5.93) affect the minimizers (5.99) (which leaves only [24] applicable). Based on recently obtained second-order bounds [88] for the eigenspaces of generically perturbed matrices, the tailor-made perturbation analysis in Section 5.8.3 is methodically independent from all prior approaches and provides this understanding in an informative, concise and essentially self-contained way. As remarked above, however, the resulting bound (5.105) has not been optimized so that a combination of Sect. 5.8.3 with some of the referenced approaches is likely to give improvements.

5.6.3 A Robust ICA-Inversion

In this section, we 'globalise' the identifiability result of Section 5.6.1 by embedding it into the robustness framework of Sections 5.3.2 and 5.5.1. More specifically, we show that the inversion map $\hat{\Phi}$ from (5.101) can be readily extended to an ICA-solution $(\mathscr{I}, \hat{\Phi})$ that is robust in the general sense of both (5.15) and Definition 5.3.10.

Definition 5.6.12 ($(\mathscr{I}_\star, \hat{\Phi})$)**.** Let $\mathscr{I}_\star$ be as in (5.106) and consider the set-valued inversion map

$$\hat{\Phi} \,:\, \dot{\mathfrak{D}} \longrightarrow \mathfrak{F}, \quad \mu \,\mapsto\, \hat{\Phi}(\mu) := \left[\operatorname*{arg\,min}_{\theta\,\in\,\tilde{\Xi}_1} \phi_\mu(\theta)\right] \cdot R_{(\mu)} \tag{5.109}$$

defined as in (5.101) and (5.99) but with Ξ_1 generalised to compact domains $\tilde{\Xi}_1$ of the form (5.111). (Since the difference between (5.101) and (5.109) is minor, we give both maps the same symbol.)

Recall from Theorem 5.6.10 that the above pair $(\mathscr{I}_\star, \hat{\Phi})$ is an ICA-solution, that is

$$\hat{\Phi}(A \cdot \tilde{\mu}) \cdot A \subseteq \mathrm{M}_d \qquad \text{for each} \quad (\tilde{\mu}, A) \in \mathscr{I}_\star,$$

and recall from Lemma 5.5.9, cf. also Remark 5.4.4, that the family $\mathscr{I}_\star$ of exactly identifiable causes is far from empty and in fact very large.

In accordance with its abstract blueprint (5.14), the inversion map (5.17) associated to the M-estimator (5.109) essentially parametrises (a subset of) the maximal solution (5.12). To ensure an informative robustness analysis à la (5.15) of this map,[29] we follow (5.98) and exhaust the set $\mathscr{I}_\star$ of identifiable causes by regular sublevel sets (5.110) related to condition numbers of the cause.

Remark 5.6.13 (Condition Bounds)**.** Recall that the optimisation domains $\Xi_1 = \Xi_1(\kappa_0)$ in (5.98) are defined in dependence of a prior condition bound $\kappa_0 < \infty$. To 'globalise' this definition for use in (5.109), note that the 'identifiability superset' $G := \mathfrak{D} \times \mathrm{GL}_d$ can be exhausted via

$$G = \bigcup\nolimits_{b \geq 1} G_b \qquad \text{for} \qquad G_b := \{(\mu, A) \in G \mid \kappa_2(\bar{B}_{(A,\mu)}) \leq b\} \tag{5.110}$$

where for any given $(\mu, A) \in G$ the matrix $\bar{B}_{(A,\mu)}$ is defined analogous to (5.92), that is: $\bar{B}_{(A,\mu)} := \mathrm{ddiag}(|b_1|, \ldots, |b_d|)^{-1} B_{(A,\mu)} \in \Xi_1^0$ with $B_{(A,\mu)} \equiv (b_1| \cdots |b_d)^\intercal := (R_\mu A)^{-1}$ for $R_\mu := R_{(A\mu)}$ the square root of the inverted (half) covariance $C_{A\mu}^{-1}$, cf. (5.89). For any $b \geq 1$ and $\Delta_\kappa > 0$, we set

$$\tilde{\Xi}_1 := \Xi_1^0 \cap \{\theta \in \mathrm{GL}_d \mid \kappa_2(\theta) \leq b + \Delta_\kappa =: \tilde{\kappa}_0\}. \tag{5.111}$$

For any priorly chosen regularity parameter b and tolerance $\Delta_\kappa > 0$, the compact (see Lem. 5.6.6) sets (5.111) are the domains we use in (5.109). Note that since for each $(\mu, A) \in G_b$ the condition number $\kappa_2(\bar{B}_{(A,\mu)})$ is continuous[30] in $[\mu]_0$, there is an (explicitly computable) $r_* = r_*((\mu, A), \Delta_\kappa) > 0$ such that $\sup\{\kappa_2(\bar{B}_{(A,\tilde{\mu})}) \mid \tilde{\mu} \in \mathfrak{D} \,:\, \|[\tilde{\mu}]_0 - [\mu]_0\| \leq r_*\} \leq b + \Delta_\kappa$. ◆

[29] For formal consistency with Section 5.3.2, Theorem 5.6.14 below considers a 'full-domain' extension $\hat{\Phi} : \mathcal{M}_1 \to \mathfrak{F}$ of (5.109), where $\hat{\Phi} : \dot{\mathfrak{D}} \to \mathfrak{F}$ is as in (5.109) and $\hat{\Phi} : \mathcal{M}_1 \setminus \dot{\mathfrak{D}} \to \mathfrak{F}$ is defined arbitrarily. More meaningful extensions of (5.109) can, for instance, be defined as per Remark 5.4.9.

[30] More precisely, we have $\kappa_2(\bar{B}_{(A,\mu)}) = \tilde{\varphi}_A([\mu]_0)$ for the continuous (at the point $[\mu]_0 \in \mathrm{GL}_d$; see Lemma 5.8.8) function $\tilde{\varphi}_A(\mathfrak{a}) := \kappa_2(\mathcal{R}(\mathcal{C}_\mathfrak{a})A \cdot \mathrm{ddiag}_i(|(A^{-1}\mathcal{R}(\mathcal{C}_\mathfrak{a})^{-1})^\intercal \cdot e_i|)) \in \bar{\mathbb{R}}$; here we set $\kappa_2(\mathfrak{b}) = +\infty$ if $\mathfrak{b} \notin \mathrm{GL}_d$.

Let us fix any E as in (5.73) and denote by $\tau_{\mathfrak{C}}$ its associated causal topology (5.76). Let further $\tilde{\mathscr{I}} := \mathscr{I}_\star \cap G_b$ for any fixed $b \geq 1$ which then defines $\tilde{\Xi}_1$ via (5.111), and set $\mathscr{I} := \tilde{\mathscr{I}} \cap E$.

The next result asserts that the ICA-solution $(\mathscr{I}, \hat{\Phi})$ is robust in the sense of (5.15) and (5.24); it also quantifies this robustness by providing explicit moduli for the underlying continuities.

Theorem 5.6.14 (Robustness)**.** *The ICA-solution $(\mathscr{I}, \hat{\Phi})$ is robust in the sense of* (5.15) & (5.24). *Specifically: Both the map $\mathfrak{I} : (\mathfrak{C}, \tau_{\mathfrak{C}}) \to (\mathfrak{F}, \mathfrak{d})$ given by $(\mu, f) \mapsto \hat{\Phi}(f_*\mu)$, as well as the deviance function $\partial_{\hat{\Phi}} : (\mathfrak{C}, \tau_{\mathfrak{C}}) \to \mathbb{R}_+$ of $\hat{\Phi}$, are continuous on $\mathscr{I}$. In fact, for any point $(\mu_\star, A) \in \mathscr{I}$ there are explicit constants $\tilde{c}_i = \tilde{c}_i(\mu_\star, A, c_{\mathfrak{S}}, \tilde{\kappa}_0) > 0$, $i = 0, 1, 2, 3, 4$, with the following property: Given $\varepsilon > 0$ define $\tilde{\delta} := \tilde{c}_0 \wedge \big[\tilde{c}_1^{2/\alpha}\big(\varepsilon/(\tilde{c}_2 + \varepsilon)\big)^{2/\alpha}\big]$, then for each point $\mathfrak{c} \in \mathfrak{C}$ the inclusion $\mathfrak{c} \in \mathbb{B}_{\tilde{\delta}}(\mu_\star, A)$ implies the following:*

$$\begin{gathered} \textit{for every } \ \theta \in \mathfrak{I}(\mathfrak{c}) \ \textit{ there is } \ M = M(\theta) \in \mathrm{M}_d \ \textit{ with } \ \|M\| \leq \tilde{c}_3 \ \textit{ and} \\ E = E(\theta) \in \mathbb{R}^{d\times d} \ \textit{ with } \ \|E\| \leq \varepsilon \ \textit{ such that } \ \theta = M(\mathrm{I} + E)A^{-1}; \end{gathered} \tag{5.112}$$

$$\textit{further,} \qquad \partial_{\hat{\Phi}}(\mathfrak{c}) \,\leq\, \tilde{c}_4\big(\varepsilon + \bar{d}(A, f)(1 + \varepsilon)\big). \tag{5.113}$$

If in fact $\mu \in \mathfrak{D}$ and $f \equiv A$, then the above holds as stated but for $\alpha = 1$ and simpler constants $\tilde{c}_0, \ldots, \tilde{c}_4$ that can each be chosen independent of $c_{\mathfrak{S}}$.

Proof. The proof of the $\mathscr{I}$-continuity of $\mathfrak{I}$ is rather technical and hence delegated to Section 5.8.4, whose notation we adopt here. As we will see next, the asserted constants $\tilde{c}_0, \ldots, \tilde{c}_4$ can be distilled from said proof. Indeed: Going through Section 5.8.4 for any fixed $\mathfrak{p}_\star \equiv (\mu_\star, A) \in \mathscr{I}$, let first $\tilde{c}_0 := r_3$ for $r_3 = r_3(\mathfrak{p}_\star, c_{\mathfrak{S}}, \tilde{\kappa}_0) > 0$ as defined in (5.215). Set further $\tilde{c}_1 := 1/K'$ and $\tilde{c}_2 := \tilde{K}_1$, for $K' \,(\equiv K'|_{r=r_3}) = K_{r_3}(\mu_\star, K_{\mathfrak{S}})$ and $\tilde{K}_1 = \tilde{K}_1(r_3)$ both as defined right after display (5.215). The θ-related (first) assertion in (5.112) then follows from (5.216). Indeed, (5.216) gives the θ-identity in (5.112) with $\|E\| \leq \epsilon_r$, while the monotonicity of $[0, C_0 \vee \hat{r}_\varepsilon] \ni \tilde{r} \mapsto \epsilon_{\tilde{r}}$ implies that $\|E\| \leq \epsilon_{\tilde{\delta}} \leq \epsilon_{\hat{r}_\varepsilon} = \varepsilon$ for $\hat{r}_\varepsilon := \tilde{c}_1^{\alpha/2}(\varepsilon/(\tilde{c}_2 + \varepsilon))^{2/\alpha}$. The bound $\|M\| \leq \sqrt{d}\|A_{R_{f_*\mu}}\|_2 \leq \sqrt{d}L_{\mathfrak{p}_\star}\|A\|_2 =: \tilde{c}_3$ holds by ((5.202) and) (5.199) and (5.210).

We now prove (5.113), from which the $\mathscr{I}$-continuity of $\partial_{\hat{\Phi}}$ follows immediately. For this, fix any $\mathfrak{c} \equiv (\mu, f) \in \mathbb{B}_{\tilde{\delta}}(\mathfrak{p}_\star)$ as before and let $B \in \hat{\Phi}(f_*\mu)$ be arbitrary. Then by (5.112) there is $M \in \mathrm{M}_d$, in fact with $\kappa_2(M) \leq \sqrt{d}\hat{\sigma}^\star_{\tilde{c}_0}L_{\mathfrak{p}_\star}\kappa_2(A) =: \tilde{c}_4$ (see (5.200), followed by (5.90),(5.210)), s.t.

$$\eta := \sup\nolimits_{u\in\mathbb{R}^d\setminus\{0\}} \tfrac{|BAu - Mu|}{|Mu|} \,\leq\, \tilde{c}_4 \cdot \varepsilon \tag{5.114}$$

which follows by the same argument as in (5.107). Recalling (5.77), (5.75), (5.74), we have $f = A(\mathrm{I} + R)$ for $R \in C^1(\mathbb{R}^d; \mathbb{R}^d)$ with $\|D_R\|_\infty \leq \bar{d}(A, f) \leq \tilde{\delta}$, and since f is a bijection we can choose $v := -f^{-1}(0) \in \mathbb{R}^d$. Further $(\mathrm{I}+R)(u-v) = (\mathrm{I}+\tilde{R})(u)$ for $\tilde{R}(u) := R(u-v)-v$, so that

$$\tfrac{|B\circ f(u-v)-Mu|}{|Mu|} = \tfrac{|BA(\mathrm{I}+R)(u-v)-Mu|}{|Mu|} = \tfrac{|BAu-Mu+BA\tilde{R}(u)|}{|Mu|} \leq \eta + \tfrac{|BA\tilde{R}(u)|}{|Mu|}$$

for each $u \in \mathbb{R}^d \setminus \{0\}$. As to the last summand, note $BA = M(\mathrm{I}+E)$ (by (5.112)) and $\tilde{R}(0) = R(f^{-1}(0)) + f^{-1}(0) = A^{-1}(f(-v)) = 0$, so that for $\hat{R} := \tilde{R}\circ M^{-1}$ and each $u \in \mathbb{R}^d$,

$$|BA\tilde{R}(u)| \leq \|M(\mathrm{I}+E)\|_2 |\hat{R}(Mu)| \quad \text{and} \quad |\hat{R}(Mu)| \leq \|M^{-1}\| \|D_R\|_\infty |Mu|$$

where the last inequality uses (the chain rule and) the mean-value theorem. All combined,

$$\sup_{u\in D_\mu^{(v)}} \tfrac{|B\circ f(u-v)-Mu|}{|Mu|} \mathbb{1}_\times(u) \leq \tilde{c}_4\cdot\varepsilon + \kappa_2(M)(1+\varepsilon)\bar{d}(A,f) \leq \tilde{c}_4\big(\varepsilon + \bar{d}(A,f)(1+\varepsilon)\big)$$

from which the asserted $\partial_{\hat{\Phi}}$-inequality (5.113) follows via (5.20). The theorem's last assertion is a direct consequence of Remark 5.8.10 (ii) applied to (5.112) and (5.113). □

We emphasize that the proof of Theorem 5.6.14 is entirely constructive, and that each of the above constants c_i and $\tilde{c}_i$, while not fine-tuned for optimality, can be read off in explicit form; see Section 5.8.4.

5.7 Applications

The robustness of an ICA inversion map, as formulated in Definition 5.3.10 and established in Theorem 5.6.14, is clearly of central importance in any situation where the exact structural assumptions (5.18) of the ICA-model are violated or other sources of (non)systematic error, including approximations of the involved signals and their statistics, are unavoidable. This naturally includes virtually all practical applications, as for these the (exact) law of a signal is typically not available and its associated statistics can only be estimated from empirical data which, in turn, is frequently corrupted by noise or subject to other forms of uncertainty. In this section, we present a selection of three essentially independent exemplary use cases, cf. Example 1.1.2, to illustrate how our theory helps to conveniently attain a meaningful quantification of various practice-relevant robustness scenarios.

Throughout this section, let $(X,S,A)_{\mathscr{I}_*}$ be any fixed BSS-triple on (5.106) with $S = (S_t)$ some stochastic process in $\mathbb{R}^d$ (defined on a probability space $(\Omega, \mathscr{F}, \mathbb{P})$), and consider $\hat{\Phi}$ as in (5.109) but with $\tilde{\Xi}_1$ given by (5.111) for $b := \kappa_2(\bar{B}_{(A,S)})$ and any fixed $\Delta_\kappa > 0$.

5.7.1 Statistical Estimation

In practice we usually have access to neither the law $\mu_X \equiv \mathbb{P}_X$ of the observable (which is needed to compute the statistics (5.38)) nor to its continuous-time realisations, but instead only to time-discretized sample trajectories of X. The issue of time-discretisation is easily resolved, since a sequence of (random) vectors $X_{t_1}, \dots, X_{t_n}$ in $\mathbb{R}^d$ can be immediately identified with a (random) element of $\mathcal{C}_d$ via piecewise-linear interpolation, see Remark 5.4.4. The inaccessibility of the law, however, is more subtle and must be addressed by estimating μ_X from the available data. While μ_X is generally an infinite-dimensional object, our inversion procedure (5.109) only requires information on its low-dimensional projection (5.39) so that any viable estimate $\hat{\mu}$ of μ_X needs to be accurate enough only at the low-order levels

$$\langle \hat{\mu} \rangle_w \approx \langle \mu_X \rangle_w = \mathbb{E}\big[\mathfrak{sig}_w(X)\big] \quad \text{for} \quad |w| = 2, 3. \tag{5.115}$$

What all reasonable estimators $\hat{\mu} = (\hat{\mu}_n) \subset \mathfrak{D}$ of μ_X will have in common is the property of statistical *consistency*, i.e. that in the 'infinite data limit' $n \to \infty$ (cf. Section 4.4)

$$\langle \hat{\mu}_n \rangle_w = \int_{\mathcal{C}_d} \mathfrak{sig}_w(x)\, \hat{\mu}_n(\mathrm{d}x) \xrightarrow{n \to \infty} \langle \mu_X \rangle_w \quad \text{for each} \quad |w| = 2, 3, \tag{5.116}$$

almost surely or in probability; this convergence is but a topological specification of the informal accuracy requirement (5.115).

Remark 5.7.1. To understand (5.116), note that strictly speaking any sample-based approximation $(\hat{\mu}_n)$ of μ_X is in fact naturally a sequence of random measures, i.e. a sequence of $\mathfrak{D}$-valued random variables. This also pertains to the example estimators in Remark 5.7.2 below; for instance, its last example can be written more precisely as $\hat{\mu} = \hat{\mu}(X, n) : \Omega \ni \omega \mapsto \hat{\mu}_n(X(\omega)) \in \mathfrak{D}$. We suppress this additional complexity in our notation to avoid cluttering our exposition.

Remark 5.7.2 (Signature Estimators)**.** An 'optimal' choice of $\hat{\mu}$ generally depends on the data that is given, but for many applications useful estimates can be straightforward. For example, if X describes the temporal evolution of some d-dimensional health trajectory of a prototypical patient and our data consists of n independently sampled path-valued[31] observations $x_1, \dots, x_n \in \mathcal{C}_d$, corresponding to the trajectories of n different real patients, then the empirical measure $\hat{\mu} \equiv \hat{\mu}_n(x_1, \dots, x_n) = \frac{1}{n}\sum_{j=1}^{n} \delta_{x_j} \in \mathfrak{D}$ can be sufficient when n is large. For other applications a good choice of $\hat{\mu}$ may be more involved and require additional statistical assumptions on X in order for a feasible approximation of its law from its samples to be possible. A very common such case is when the data consists of a

[31] Upon piecewise-linear interpolation of discretely-measured data points, say.

single trajectory $x = (x_t)$ of X only; in this case, common mixing or ergodicity assumptions on X [or, equivalently, on S] often suffice to qualify the usual ergodic plug-in estimators $\hat{\mu} \equiv \hat{\mu}_n(x) = \frac{1}{n}\sum_{j=1}^{n} \delta_{x|_{((j-1)\tau, j\tau]}}$ ($\tau > 0$ large enough) as suitable approximators for μ_X; see Section 4.4 for details. ◆

Our continuity results in Section 5.6 provide explicit bounds to relate the convergence (5.116) to a controllable approximation of the inverse A^{-1}. This allows us to conveniently derive finite-sample results of the following type.

Notation 5.7.3 (Convergence Rate)**.** In line with classical theory, an estimator $\hat{\mu} \equiv \hat{\mu}_n$ of μ_X is said to have convergence rate $\tilde{\alpha}_n$ – for $(\tilde{\alpha}_n)$ some strictly decreasing null sequence – if $\gamma_n := \max_{|w|=2,3} |\langle\hat{\mu}_n\rangle_w - \langle\mu_X\rangle_w| = O_P(\tilde{\alpha}_n)$, which is to say that for any $\epsilon > 0$ there exist numbers $0 < m, M < \infty$ such that $\sup_{n\geq m} \mathbb{P}(\gamma_n \geq M\tilde{\alpha}_n) \leq \epsilon$, see e.g. [152, Sect. 2.2].

Proposition 5.7.4. *Let $\hat{\mu} \equiv (\hat{\mu}_n) \subset \dot{\mathfrak{D}}$ be a consistent estimator of μ_X with $\max_{|w|=2,3} |\langle\hat{\mu}_n\rangle_w - \langle\mu_X\rangle_w| = O_P(\tilde{\alpha}_n)$ for some convergence rate $\tilde{\alpha} \equiv (\tilde{\alpha}_n)$. Then for any $\varepsilon > 0$ and any $0 \leq q < 1$, there exists an explicitly computable – as per the below proof – threshold $n_0 = n_0(\varepsilon, q, \tilde{\alpha}, \hat{\mu}, A, \mu_S) \in \mathbb{N}$ such that for each $n \geq n_0$ the following holds with a probability of at least q:*

$$\textit{for every } \hat{\theta} \in \hat{\Phi}(\hat{\mu}_n) \textit{ there is } M \in \mathrm{M}_d \textit{ such that } \sup_{t\in\mathbb{I}} \frac{|\hat{\theta}X_t - MS_t|}{|MS_t|}\mathbb{1}_{\times}(S_t) \leq \varepsilon. \quad (5.117)$$

Proof. Let $B := A^{-1}$ and $\upsilon_n := B \cdot \hat{\mu}_n$. Lemma 5.4.7 then asserts the identities of 2- and 3-tensors $[\upsilon_n]^{(2)} := (\langle\upsilon_n\rangle_{ij}) = B^{\otimes 2}(\langle\hat{\mu}_n\rangle_{ij})$ and $[\upsilon_n]^{(3)} := (\langle\upsilon_n\rangle_{ijk}) = B^{\otimes 3}(\langle\hat{\mu}_n\rangle_{ijk})$, where $B^{\otimes 2} = B \otimes B$ and $B^{\otimes 3} = B \otimes B \otimes B$ are Kronecker powers of B (see also [54, Prop. 7.52]). Since the action of $B^{\otimes m}$ on $(\mathbb{R}^d)^{\otimes m}$ is continuous, we obtain by way of the continuous mapping theorem that the [in probability] convergence (5.116) implies that also $[\upsilon_n]^{(m)} = B^{\otimes m}[\hat{\mu}_n]^{(m)} \longrightarrow_{n\to\infty} B^{\otimes m}[\mu_X]^{(m)} = [\mu_S]^{(m)}$ [in probability] for $m = 2, 3$. In fact, it is easy to see that $\gamma_n := \max_{|w|=2,3} |\langle\upsilon_n\rangle_w - \langle\mu_S\rangle_w| = O_P(\beta_n)$ with $\beta_n = C \cdot \tilde{\alpha}_n$ for an explicit constant $C = C(B) > 0$.

Take now $\tilde{\varepsilon} > 0$, and set $\varepsilon := \tilde{\varepsilon}/\tilde{c}_4$ for $\tilde{c}_4 \equiv \tilde{c}_4(\mu_S, A) > 0$ as in Theorem 5.6.14. This theorem then provides an explicit radius $\tilde{\delta} = \tilde{\delta}(\varepsilon, \mu_S, A, \tilde{\kappa}_0) > 0$ such that (5.113) holds for $\mu_\star = \mu_S$, i.e. such that

$$\sup_{(\upsilon,A)\in\mathbb{B}_{\tilde{\delta}}(\mu_S,A)} \partial_{\hat{\Phi}}(\upsilon, A) \leq \varepsilon. \quad (5.118)$$

To connect this with (5.116) note that by definition (5.55) of $\delta(\cdot,\cdot)$ and continuity,[32] there is an $\eta > 0$ such that $\delta(\mu_S, \upsilon) < \tilde{\delta}$ for each $\upsilon \in \dot{\mathfrak{D}}$ with $\max_{|w|=2,3} |\langle\upsilon\rangle_w - \langle\mu_S\rangle_w| < \eta$. Now fix any $0 \leq q < 1$ and let $q' := 1 - q$. Since $\gamma_n = O_P(\beta_n)$, we can find $m_q, M_q > 0$ with

[32] More precisely, by the fact that $\delta(\mu_S, \upsilon) = \varphi_S([\upsilon]_{\mathfrak{c}})$ for $\varphi_S : \mathcal{V} \to \mathbb{R}_+$ continuous with $\varphi_S([\mu_S]_{\mathfrak{c}}) = 0$.

$\sup_{n\geq m_q}\mathbb{P}(\gamma_n \geq CM_q{\cdot}\tilde{\alpha}_n) \leq q'$. Due to $\tilde{\alpha}_n \downarrow 0$, we can choose a (smallest) $n_0 \geq m_q$ such that $\tilde{\alpha}_n \leq \eta/(CM_q)$ for each $n \geq n_0$. For this n_0 we have $\sup_{n\geq n_0}\mathbb{P}(\gamma_n \geq \eta) \leq \sup_{n\geq n_0}\mathbb{P}(\gamma_n \geq CM_q\tilde{\alpha}_n) \leq \sup_{n\geq m_q}\mathbb{P}(\gamma_n \geq CM_q\tilde{\alpha}_n) \leq q'$ and thus $\inf_{n\geq n_0}\mathbb{P}(\gamma_n < \eta) \geq 1 - q' = q$. Hence for each $n \geq n_0$, it holds with probability at least q that $\mathbb{d}((\upsilon_n, A), (\mu_S, A)) = \delta(\mu_S, \upsilon_n) < \tilde{\delta}$, i.e. that $(\upsilon_n, A) \in \mathbb{B}_{\tilde{\delta}}(\mu_S, A)$. Assertion (5.117) now follows by (5.118) and Lemma 5.3.9 upon noting that here, by (5.114), the spatial support in (5.23) can be replaced by the entirety of $\mathbb{R}^d$. □

5.7.2 Asynchronous Observations

Other than statistical estimation (5.116), another source of systematic error in the blind inversion of time-dependent signals are deviations from the instantaneity assumption (5.7) that are caused by an asynchronous time-discretizations of the components $X^1, \ldots, X^d$ of X, cf. e.g. [2] for the classical context of covariance estimation:

In practice, often the channel signals X^i are each observed by different sensors which take their measurements at some discrete sets of time-points that may be different from sensor to sensor. Empirical observation, e.g. [101, 133], demonstrates that this asynchronicity perturbs BSS-performances up to the point of dysfunctionality. As the following proposition shows, the causal topology (5.76) is coarse enough to capture this 'perturbation by asynchronicity' so that its effect can then be quantified and analysed via Theorem 5.6.14.

To this end, we may formalize the synchronicity infringement of (5.7) as follows: Write $X^i_{\mathcal{J}_i} := (X^i_t \mid t \in \mathcal{J}_i := \{0 \leq t^{(i)}_0 < t^{(i)}_1 < \ldots < t^{(i)}_{n_i-1} \leq 1\})$ for the sensor data of channel X^i, denoting $\hat{X}^i_{\mathcal{J}_i}$ for the piecewise-linear interpolation of this data. Formally, $\hat{X}^i_{\mathcal{J}_i} = \hat{\iota}_{n_i}(X^i_{\mathcal{J}_i})$ for $\hat{\iota}_{n_i}$ as in (5.35) for $d = 1$, see Rem. 5.4.4. We then say that the channels X^i are observed synchronously if $\mathcal{J}_1 = \mathcal{J}_2 = \ldots = \mathcal{J}_d$, and otherwise say that the $X^1, \ldots, X^d$ are observed asynchronously.

Writing $\|\mathcal{J}_i\| := \max_{\nu\in[n_i-1]}|t^{(i)}_\nu - t^{(i)}_{\nu-1}|$ for the mesh-size of $\mathcal{J}_i$, we can handle asynchronicity via Theorem 5.6.14 by noting that the interpolated observations $\hat{X}^*_{(\mathcal{J}_i)} := (\hat{X}^1_{\mathcal{J}_1}, \cdots, \hat{X}^d_{\mathcal{J}_d})$ of X over the channel-dependent sample times $(\mathcal{J}_i) \equiv (\mathcal{J}_1, \ldots, \mathcal{J}_d)$ converge to X in (5.76) as $\|(\mathcal{J}_i)\| := \max_i \|\mathcal{J}_i\| \to 0$. The symbol $\tilde{\varrho}$ below denotes the relative error from (5.21).

Proposition 5.7.5. *For each $i \in [d]$, let $(\mathcal{J}^n_i)_{n\in\mathbb{N}}$ be a sequence of dissections of $[0,1]$ whose mesh-size $\|\mathcal{J}^n_i\|$ converges to zero. Further, suppose $\mathbb{E}\|X\|^3_{1\text{-var}} < \infty$ and let $\hat{\mu}_n := \mathbb{P}_{\hat{X}^*_n} \in \dot{\mathfrak{D}}$ be the law of $\hat{X}^*_n := \hat{X}^*_{(\mathcal{J}^n_i)}$. Then for each $\varepsilon > 0$ there is an explicit $n_0 = n_0(\varepsilon, A, \mu_S, (\hat{\mu}_n)) \in \mathbb{N}$ such that for all $n \geq n_0$, the error bound $\operatorname{dist}_{\tilde{\varrho}}(\hat{\Phi}(\hat{\mu}_n) \cdot X, \lceil S \rceil_{\mathfrak{m}}) \leq \varepsilon$ holds almost surely.*

Proof. This follows from Theorem 5.6.14 and a few standard convergence arguments. Indeed: Denoting by $\hat{\pi}_n^i : \mathcal{C}_1^1 \to \mathcal{C}_1^1$ (cf. (5.28)) the linear operator that sends a $\mathcal{C}_1^1$-path $y : [0,1] \to \mathbb{R}$ to its $\hat{\mathcal{J}}_n^i$-piecewise-linear interpolation $\hat{y}_{\mathcal{J}_n^i} := \hat{\iota}_{|\mathcal{J}_n^i|}(\pi_{\mathcal{J}_n^i}(y))$ (cf. Remark 5.4.4), note that $\|\hat{\pi}_n^i(y)\|_{1\text{-var}} \leq \|y\|_{1\text{-var}}$ for all $y \in \mathcal{C}_1^1$ and each $n \in \mathbb{N}$, and further that $\lim_{n\to\infty} \hat{\pi}_n^i = \mathrm{id}$ pointwise [on $\mathcal{C}_1^1$] wrt. $\|\cdot\|_\infty$, see [54, Thm. 5.23]. Consequently, for each $x \in \mathcal{C}_d^1$ and $n \in \mathbb{N}$,

$$d^{-1/2}\|x - \hat{x}_n^*\|_\infty \leq \max_{i\in[d]} \left\|x^i - \pi_i(\hat{x}_n^*)\right\|_\infty = \max_{i\in[d]} \left\|x^i - \hat{\pi}_n^i(x^i)\right\|_\infty \longrightarrow 0 \quad (\text{as } n \to \infty)$$

and $\|\hat{x}_n^*\|_{1\text{-var}} \leq \sum_{i=1}^d \|\hat{\pi}_n^i(x^i)\|_{1\text{-var}} \leq \sum_{i=1}^d \|x^i\|_{1\text{-var}}$ (recall (5.32) for the first inequality). By [54, Lemma 5.27 (i)], the above implies that $\lim_{n\to\infty} \|x - \hat{x}_n^*\|_{p\text{-var}} = 0$ for each $x \in \mathcal{C}_d^1$, for any $p > 1$. Hence by the p-variation continuity of $\mathfrak{sig}(\cdot)$ (for $p \in [1,2)$) we have

$$\lim_{n\to\infty} \mathfrak{sig}_w(\hat{x}_n^*) \;=\; \mathfrak{sig}_w(x) \qquad (\text{each } |w| \leq 3)$$

for each $x \in \mathcal{C}_d^1$. Due to this and the $L_1(\mathbb{P}_X)$-domination $|\mathfrak{sig}_w(\hat{x}_n^*)| \lesssim \|\hat{x}_n^*\|_{1\text{-var}}^3 \leq d^3\|x\|_{1\text{-var}}^3$ (which holds by the signature estimate underlying (5.51)), we can apply dominated convergence to obtain (5.116), i.e. that $[\hat{\mu}_n]_\mathfrak{c} \to [\mu_X]_\mathfrak{c}$ and hence $\delta(\hat{\mu}_n, \mu_X) \to 0$ as $n \to \infty$. The claimed conclusion, which is equivalent to (5.117), now follows as in the proof of Proposition 5.7.4. ☐

5.7.3 Additive and Multiplicative Noise

Despite ever-increasing efforts towards more sophisticated measuring instruments and data-processing procedures, empirical observations such as experimental measurements or econometric records are never exact. Instead, any type of empirically rooted data is prone to unwanted deviations and corruptions whose effect needs to be understood as it may otherwise invalidate the (algorithmic) consistency of a given statistical method. One way to account for this are mathematical ‘noise models’, of which the two most established ones are *additive* and *multiplicative* noise affecting a signal $Y = (Y_t)_{t\in\mathbb{I}}$,

$$Y_t \;=\; Y_t + \eta_t \qquad \text{or}^{33} \qquad Y_t \;=\; \eta_t \cdot Y_t \qquad (t \in \mathbb{I}) \tag{5.119}$$

for some stochastic process $\eta = (\eta_t)_{t\in\mathbb{I}}$ in $\mathbb{R}^d$ referred to as the “noise”. For noisy BSS models [112], the noise η is typically subject to additional, often (semi)parametric and sometimes quite involved structural assumptions, see e.g. [12, 38, 89, 92] and the references therein.[34]

[33] The multiplication $\eta_t \cdot Y_t$ is understood componentwise, i.e. $\eta_t \cdot Y_t \equiv (\eta_t^1 Y_t^1, \cdots, \eta_t^d Y_t^d)^\intercal$ for each $t \in \mathbb{I}$.

[34] In order to avoid dealing with any identifiability issues at all, some authors impose rather restrictive (‘contrasting’) structural assumptions on S or η to ensure that the additive noise model (5.120) remains exactly invertible. Notable such assumptions include: time-dependence (S) vs time-independence (η) [35], non-Gaussianity (of S) vs Gaussianity (of η) [76], or time-variability of [auto]covariances (S) vs group-wise stationary confounding (η) [125]. Of course, these cases are all hand-picked exceptions and treatment of the general, model-free case (5.120) requires a more comprehensive analytical framework, see Proposition 5.7.6.

The generality of our premetric framework allows us to keep such assumptions to a minimum.

In the classical BSS-context (5.7) where f is linear, it is essentially irrelevant whether the additive noise (5.119) applies to $Y = X$ or $Y = S$ (or both) because

$$X + \eta \;=\; A(S + \eta') \qquad \text{for} \quad \eta' := A^{-1}\eta\,; \tag{5.120}$$

for the corruption with multiplicative noise, we may distinguish[35] the cases

$$X^{\eta}_t := \eta_t \cdot X_t \quad \text{ or } \quad S^{\eta}_t := \eta_t \cdot S_t \qquad (t \in \mathbb{I}). \tag{5.121}$$

We can apply Theorem 5.6.14 to quantify how these corruptions affect the accuracy of $(\mathscr{I}_\star, \hat{\Phi})$.

Proposition 5.7.6. *Let $\eta \in \mathfrak{D}$ be a noise process and consider the corruptions*

$$\tilde{X}^{(1)} := A(S + \eta) \quad \text{ and } \quad \tilde{X}^{(2)} := AS^{\eta} \quad \text{ and } \quad \tilde{X}^{(3)} := X^{\eta} \tag{5.122}$$

with $\mu^{(i)} := \mathbb{P}_{\tilde{X}^{(i)}}$ their respective laws. There are then explicit constants $c_j = c_j(S, A, \Delta_\kappa) > 0$ $(j = 0, \ldots, 3)$ with the following property: Given $\varepsilon > 0$ define $\beta_\varepsilon := c_0 \wedge \big(c_1(\varepsilon/(c_2 + \varepsilon))^2\big)$, then

$$\text{for each } \ \hat{\theta} \in \hat{\Phi}(\mu^{(i)}) \ \text{ there is } \ M \in \mathrm{M}_d \ \text{ with } \ \tfrac{\|\hat{\theta} - MA^{-1}\|}{\|MA^{-1}\|} \le \varepsilon \quad \text{and} \tag{5.123}$$

$$\max\big\{\partial_{\hat{\Phi}}(S + \eta, A), \partial_{\hat{\Phi}}(S^{\eta}, A), \partial_{\hat{\Phi}}(S, \eta A)^{*}\big\} \;\le\; \varepsilon \tag{5.124}$$

provided that S and the noise process η meet the following case-dependent assumptions:

- *for the case $i = 1$, suppose that η and S are independent and both mean-stationary and that η satisfies the growth condition $\sum_{\nu=0}^{d} \langle S\rangle_{\nu\nu}^{-1} \big\| N_S^{-1}[\eta]_\nu N_S^{-1} \big\|^2 \le \beta_\varepsilon^2$;*
- *for the case $i = 2$, suppose that S is mean-stationary and centered and that η is independent of S and satisfies the growth condition $\beta_2(S, \eta) \le \beta_\varepsilon^2$;*
- *for the case $i = 3$, suppose that η is independent from S, mean-stationary with $\mathbb{E}[\eta] \equiv \mathrm{I}$ and IC and satisfies the growth condition $\beta_3(X, \eta) \le \gamma_\varepsilon$ with β_3 from (5.230) and $\gamma_\varepsilon = \gamma_\varepsilon(\beta_\varepsilon, A) > 0$ as introduced directly thereafter;*

[35] Since A is a homoeomorphism these cases, too, are essentially equivalent, but see the proof of Prop. 5.7.6.

* Denoting $\eta A : [0,1] \times \mathbb{R}^d \to \mathbb{R}^d$, $(t, u) \mapsto \eta_t A u := \mathrm{ddiag}(\eta^1_t, \cdots, \eta^d_t) \cdot Au$ (a random time-dependent transformation), the last inequality in (5.124) ('$\partial_{\hat{\Phi}}(S, \eta A) \le \varepsilon$') is to be read as: $\partial_{\hat{\Phi}}(S, \eta A) \equiv \sup_{B \in \hat{\Phi}((\eta_t A S_t \,|\, t \in [0,1]))} \inf_{(M,v) \in \mathrm{M}_d \times \mathbb{R}^d} \sup_{t \in [0,1]} \sup_{u \in D^{(v)}_{S_t}} \frac{|B \circ (\eta_t A)(u - v) - Mu|}{|Mu|} \mathbb{1}_{\times}(Mu) \le \varepsilon$ with probab. one.

here, $\langle S\rangle_{00} := 1$ *and* $\beta_2(S,\eta) := \sum_{i=1}^{d}\langle S\rangle_{ii}^{-2}(\tilde{\alpha}_i^2 + \langle S\rangle_{ii}^{-1}\check{\alpha}_i^2)$ *with* $\tilde{\alpha}_i :=$ (5.227) *and* $\check{\alpha}_i :=$ (5.228)*.*

Proof. Thm. 5.6.14 provides constants $\tilde{c}_j = \tilde{c}_j(S,A) > 0$ ($j = 0,1,2$) and $\tilde{c}_4 = \tilde{c}_4(S,A,\Delta_\kappa) > 0$ such that for any $\varepsilon > 0$ and $r_{\tilde{\varepsilon}} := \tilde{c}_0 \wedge \left[\tilde{c}_1^2\left(\tilde{\varepsilon}/(\tilde{c}_2+\tilde{\varepsilon})\right)^2\right]$ with $\tilde{\varepsilon} := \varepsilon/\tilde{c}_4$, we have that

$$\text{if}\quad \delta(S,\tilde{\mu}) \le r_{\tilde{\varepsilon}} \quad\text{then:}\quad \text{for every}\ \hat{\theta}\in\Phi(A\tilde{\mu})\ \text{ there is }\ M\in\mathrm{M}_d\ \text{ with }\ \tfrac{\|\hat{\theta}-MA^{-1}\|}{\|MA^{-1}\|}\le\varepsilon,$$

where the above implication applies to all $\tilde{\mu}\in\dot{\mathfrak{D}}$. To see that this holds, simply combine (5.112) with the argument behind (5.203) $\Rightarrow$ (5.204) and the definition of $\tilde{c}_4$. Setting $c_j := \tilde{c}_j$ for $j = 0,1$ and $c_2 := \tilde{c}_2\tilde{c}_4$ yields $r_{\tilde{\varepsilon}} = \beta_\varepsilon$, for β_ε as defined above. Hence (5.123) follow if we have

$$\delta(S, A^{-1}\mu^{(i)}) \,\le\, \beta_\varepsilon \qquad \text{for each of the corruptions in }\ (5.122). \tag{5.125}$$

Moreover, from Theorem 5.6.14 (5.113) we further obtain that (5.125) also implies $\partial_{\hat{\Phi}}(v^{(i)},A) \le \varepsilon$ for $v^{(i)} := A^{-1}\mu^{(i)}$ and each $i = 1,2,3$, which proves (5.124) for the cases $i = 1,2$. To see that it also implies the case $\partial_{\hat{\Phi}}(S,\eta A) \le \varepsilon$ and hence (5.124) altogether, notice that (cf. footnote $*$)

$$\sup_{t\in[0,1]}\sup_{u\in D^{(v)}_{S_t}}\tfrac{|B\circ(\eta_t A)(u-v)-Mu|}{|Mu|}\mathbb{1}_\times \,\le\, \sup_{u\in D^{(v)}_{v^{(3)}}}\tfrac{|BA(u-v)-Mu|}{|Mu|}\mathbb{1}_\times \qquad \text{almost surely}$$

(for any $B\in\hat{\Phi}(\tilde{X}^{(3)})$ and each $(M,v)\in\mathrm{M}_d\times\mathbb{R}^d$) and hence $\partial_{\hat{\Phi}}(S,\eta A) \le \partial_{\hat{\Phi}}(v^{(3)},A)$ a.s. Let us first establish (5.125) for the additive-noise case $i = 1$, which the triple $(\tilde{X}^{(1)},A,\tilde{S})$, with $\tilde{S} := S+\eta$, represents in full generality (cf. (5.120)). For any $ijk\in[d]_3^\star$ and by multilinearity,

$$\langle\tilde{S}\rangle_{ij} \,=\, \mathbb{E}\left[\int_0^1\!\!\int_0^t \mathrm{d}\tilde{S}_s^i\,\mathrm{d}\tilde{S}_t^j\right] \,=\, \langle S\rangle_{ij} + \langle S^i,\eta^j\rangle + \langle\eta^i,S^j\rangle + \langle\eta\rangle_{ij}, \quad \text{and similarly} \tag{5.126}$$

$$\begin{aligned}\langle\tilde{S}\rangle_{ijk} \,=\, &\langle S\rangle_{ijk} + \langle S^i,S^j,\eta^k\rangle + \langle S^i,\eta^j,S^k\rangle + \langle S^i,\eta^j,\eta^k\rangle + \langle\eta^i,S^j,S^k\rangle \\ &+ \langle\eta^i,S^j,\eta^k\rangle + \langle\eta^i,\eta^j,S^k\rangle + \langle\eta\rangle_{ijk},\end{aligned} \tag{5.127}$$

with $\langle Y^i,\tilde{Y}^j\rangle := \mathbb{E}\big[\int_0^1\!\int_0^t \mathrm{d}Y_s^i\mathrm{d}\tilde{Y}_t^j\big]$ and $\langle Y^i,\tilde{Y}^j,\hat{Y}^k\rangle := \mathbb{E}\big[\int_0^1\!\int_0^t\!\int_0^s \mathrm{d}Y_r^i\mathrm{d}\tilde{Y}_s^j\hat{Y}_t^k\big]$ for $Y,\tilde{Y},\hat{Y}\in\{S,\eta\}$. Since $\eta,S\in\mathcal{C}_d^1$, we can use Fubini to evaluate the above statistics (5.126) and (5.127) to

$$\langle S^i,\eta^j\rangle = \int_0^1\!\!\int_0^t \mathbb{E}[\dot{S}_s^i\dot{\eta}_t^j]\mathrm{d}s\mathrm{d}t \qquad\text{and}\qquad \langle\eta^i,S^j,\eta^k\rangle = \int_0^1\!\!\int_0^t\!\!\int_0^s \mathbb{E}[\dot{\eta}_r^i\dot{S}_s^j\dot{\eta}_t^k]\mathrm{d}r\mathrm{d}s\mathrm{d}t \quad \text{etc.} \tag{5.128}$$

From the assumptions that η and S are independent and both mean-stationary, we obtain $\mathbb{E}[\dot{S}_s^i\dot{\eta}_t^j] = \mathbb{E}[\dot{S}_s^i]\mathbb{E}[\dot{\eta}_t^j] = 0$ and $\mathbb{E}[\dot{\eta}_r^i\dot{S}_s^j\dot{\eta}_t^k] = \mathbb{E}[\dot{S}_s^j]\mathbb{E}[\dot{\eta}_r^i\dot{\eta}_t^k] = 0$ etc. For this we used that $\mathbb{E}[\dot{S}_t^i] = 0 = \mathbb{E}[\dot{\eta}_t^j]$ for each $i,j\in[d]$ and $t\in[0,1]$, which follows by interchanging expectation and differentiation (as is permitted by the assumption that the derivatives $\dot{S}$ and $\dot{\eta}$ are

$L^1(\mathbb{P})$-dominated). This implies that the mixed statistics (5.128) vanish, so that we are left with

$$[\tilde{S}]_\nu \;=\; [S]_\nu + [\eta]_\nu \qquad \text{for each} \qquad \nu = 0, 1, \ldots, d.$$

Hence and from (5.57), we obtain that

$$\delta(S, \tilde{S})^2 \;=\; \sum_{\nu=0}^{d} \langle S\rangle_{\nu\nu}^{-1} \big\|N_S^{-1}[\eta]_\nu N_S^{-1}\big\|^2 \;\leq\; \beta_\varepsilon^2,$$

which proves (5.125) for the additive-noise case $i = 1$. The proof of (5.125) for the multiplicative-noise cases $i = 2, 3$ is similar but slightly more technical, see Section 5.8.5. □

Remark 5.7.7. The proposition shows that both the additive and the multiplicative noise case extend the noiseless idealisations $\eta \equiv 0$ and $\eta \equiv \mathrm{I}$ (resp.) continuously. Notice how the deviance bounds (5.124) provide a concise quantification of the inversion stability of $\hat{\Phi}$ against the respective ICA-violations (5.122). ◆

5.8 Proof of Theorems 5.6.10 and 5.6.14

This section contains some technical auxiliary results to complete the proofs of Theorem 5.6.10 and Theorem 5.6.14; the latter is done in Sections 5.8.3 and 5.8.4, respectively.

5.8.1 Some Technical Estimates

Lemma 5.8.1. *For any $p \in (1, 2)$ and $\eta > 0$, denote $\alpha := 1 - 1/p$ and $\gamma := 1/p - \alpha$.*

(i) *Let $\{Y, \tilde{Y}\} \subset \mathcal{C}_{d|\mathbb{P}} \cap \mathfrak{D}$ be core-integrable and such that $\|\tilde{Y} - Y\|_\infty \leq \eta$ with probability one. Then for each $w \in [d]^\star$ with $2 \leq |w| \leq 3$ we have*

$$\big|\langle\tilde{Y}\rangle_w - \langle Y\rangle_w\big| \;\leq\; \kappa_{p,|w|} \cdot \eta^\alpha \tag{5.129}$$

with $\kappa_{p,k} := (2 + k)c_p(2^{\beta_p}K_p + 1)$, for the constants $\beta_p := 2 + 1/p$ and $c_p := (1 - 2^{1-2/p})^{-1}$ and with $K_p := \max_{V\in\{\tilde{Y},Y\}} \mathbb{E}[\|V\|_{1\text{-var}}^{\beta_p}]$.

(ii) *For any $\tilde{\mu}, \mu \in \mathfrak{D}$ such that $\max_{|w|=2,3} \big|\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w\big| \leq \kappa\eta$ for some $\kappa \geq 0$, we have*

$$\big|\delta_{\perp\!\!\!\perp}(\tilde{\mu}) - \delta_{\perp\!\!\!\perp}(\mu)\big| \;\leq\; \varphi_\kappa(\mu, \tilde{\mu})\sqrt{\eta + \eta^2} \tag{5.130}$$

for the auxiliary function $\varphi_\kappa(\mu, \tilde{\mu}) := \sqrt{(3\vartheta_{\mu,\tilde{\mu}} + \kappa)\kappa \sum_{k=0}^{d}\sum_{i,j=1}^{d} \varphi_{ijk}(\mu, \tilde{\mu})}$ with

$$\varphi_{ijk}(\mu, \tilde{\mu}) := \frac{2\big(\langle\mu\rangle_{ijk} \vee \sqrt{\langle\tilde{\mu}\rangle_{ii}\langle\tilde{\mu}\rangle_{jj}\langle\mu\rangle_{kk}}\big)^2}{\big(\sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}\langle\tilde{\mu}\rangle_{kk}} \wedge \sqrt{\langle\tilde{\mu}\rangle_{ii}\langle\tilde{\mu}\rangle_{jj}\langle\mu\rangle_{kk}}\big)^4} \tag{5.131}$$

and $\vartheta_{\mu,\tilde{\mu}} := \max\{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}, \langle\mu\rangle_{ii}\langle\tilde{\mu}\rangle_{jj} \mid i, j \in [d]\}$.

(iii) *For any* $\mu, \tilde{\mu} \in \mathfrak{D}$ *with* $\delta(\mu, \tilde{\mu}) \leq \eta$*, we have*

$$\max_{|w|=2,3} \left|\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w\right| \;\leq\; \mathfrak{m}_\mu \eta \tag{5.132}$$

for $\mathfrak{m}_\mu := \max\{\sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}\langle\mu\rangle_{kk}} \mid i,j \in [d], k \in [d]_0\}$*; if further* μ *is orthogonal, then* $\delta_{\perp\!\!\perp}(\tilde{\mu}) \;\leq\; \hat{\varphi}(\mu,\tilde{\mu})\sqrt{\eta+\eta^2}$ *for* $\hat{\varphi}(\mu,\tilde{\mu})^2 := \max\{3\vartheta_{\mu,\tilde{\mu}}\mathfrak{m}_\mu, \mathfrak{m}_\mu^2\}\sum_{k=0}^d\sum_{i,j=1}^d \varphi_{ijk}(\mu,\tilde{\mu})$.

(iv) *Let* $R \in C^1(\mathbb{R}^d;\mathbb{R}^d)$ *with*[36] $\|D_R\|_\infty < \infty$*, and let* $\mu \in \mathfrak{D}$ *be core-integrable with* $\tilde{\mu} := (\mathrm{I}+R)_*\mu$*. Then, for each* $w \in [d]^\star$ *with* $|w| = 2,3$*:*

$$\left|\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w\right| \;\leq\; \tilde{K}_{p,|w|}\cdot\phi_{|w|}(\|D_R\|_\infty)\big(\|D_R\|_\infty\|R\|_\infty\big)^\alpha, \quad \textit{and} \tag{5.133}$$

$$\left|\delta_{\perp\!\!\perp}(\tilde{\mu}) - \delta_{\perp\!\!\perp}(\mu)\right| \;\leq\; \varphi_{\tilde{K}_{p,3}}(\mu,\tilde{\mu})\,\hat{\phi}_p(R)\cdot\phi_3(\|D_R\|_\infty)^{\frac{1}{2}}\big(\|D_R\|_\infty\|R\|_\infty\big)^{\frac{\alpha}{2}} \tag{5.134}$$

for $\tilde{K}_{p,k} := 2^\alpha c_p(1+k/2)(1+\mathbb{E}_\mu[\|x\|_{1\text{-var}}^\beta])$ *and* $\phi_k(u) := (1+u)^{k-1}u^\gamma$ *and with* $\varphi_{\tilde{K}_{p,3}}(\mu,\tilde{\mu})$ *as in* (ii) *but for* $\kappa := \tilde{K}_{p,3}$*, and* $\hat{\phi}_p(R) := \sqrt{1+\phi_3(\|D_R\|_\infty)(\|D_R\|_\infty\|R\|_\infty)^\alpha}$.

Proof. (i): Let $\{Y,\tilde{Y}\} \subset \mathfrak{D}$ be square-integrable with $\|\tilde{Y}-Y\|_\infty \leq \eta$ almost surely, and let $w \in [d]^{\leq m}$ be an arbitrary multiindex of length $|w| \leq m$. Take any $p \in (1,2)$. The classical interpolation inequality [54, Prop. 5.5 (i)] then implies that for $Z := \tilde{Y}-Y$ we have

$$\|Z\|_{p\text{-var}} \;\leq\; 2^\alpha\|Z\|_{1\text{-var}}^{1/p}\|Z\|_\infty^{(1-1/p)} \;\leq\; 2^\alpha\|Z\|_{1\text{-var}}^{1/p}\eta^\alpha \;\leq\; 2\,\Xi^{1/p}\cdot\eta^\alpha \tag{5.135}$$

where $\alpha := 1-1/p$ and $\Xi := \max\{\|Y\|_{1\text{-var}}, \|\tilde{Y}\|_{1\text{-var}}\}$. The local Lipschitz continuity of the signature transform (cf. [54, Proposition 7.63]) then implies that, for each $2 \leq |w| \leq m$,

$$\left|\langle\tilde{Y}\rangle_w - \langle Y\rangle_w\right| \;\leq\; \mathbb{E}\left|\mathfrak{sig}_w(\tilde{Y}) - \mathfrak{sig}_w(Y)\right| \;\leq\; (2+|w|)c_p\,\mathbb{E}\big[\Xi^{|w|-1+1/p}\big]\cdot\eta^\alpha \tag{5.136}$$

for $c_p := (1-2^{1-2/p})^{-1}$, as follows from the proof of [54, Prop. 7.63] but with the Young-Loeve estimates $|\tilde{Y}_{s,r}-Y_{s,r}| = |Z_{s,r}| \leq \|Z\|_{p\text{-var};[s,t]}$ and $|\int_s^t \tilde{Y}_{s,r}\,\mathrm{d}Z_r| \leq c_p\|\tilde{Y}\|_{p\text{-var};[s,t]}\|Z\|_{p\text{-var};[s,t]}$.

Indeed: For each $|w| \geq 2$ we have

$$\begin{aligned}\left|\mathfrak{sig}_w(\tilde{Y}) - \mathfrak{sig}_w(Y)\right| &= \left|\langle\pi_{|w|}\big[\mathfrak{sig}_{0,1}(\tilde{Y}) - \mathfrak{sig}_{0,1}(Y)\big], w\rangle\right| \leq \left\|\pi_{|w|}\big[\mathfrak{sig}_{0,1}(\tilde{Y}) - \mathfrak{sig}_{0,1}(Y)\big]\right\| \\ &= \left\|\int_0^1 \pi_{|w|-1}\big[\mathfrak{sig}_{0,r}(\tilde{Y}) - \mathfrak{sig}_{0,r}(Y)\big]\mathrm{d}\tilde{Y}_r + \int_0^1 \pi_{|w|-1}\big[\mathfrak{sig}_{0,r}(Y)\big]\mathrm{d}Z_r\right\|\end{aligned} \tag{5.137}$$

where $\|\cdot\| := |\cdot|_1$ is the projective tensor norm on the tensor powers of $\mathbb{R}^d$, cf. [105, Sect. 5.6]. Using that for any $x \in \mathcal{C}^1([0,1];\mathbb{R}^d)$ and $y \in \mathcal{C}^1([0,1];\mathbb{R}^{e\times d})$ we have

$$\int_s^\cdot y_{s,r}\,\mathrm{d}x_r \;\in\; \mathcal{C}^1([s,t];\mathbb{R}^e) \quad \text{and} \quad \left|\int_s^t y_{s,r}\,\mathrm{d}x_r\right| \;\leq\; c_p\|y\|_{p\text{-var};[s,t]}\|x\|_{p\text{-var};[s,t]}$$

[36] Here and in the following, we norm the Jacobian D_R of a map $R \in C^1(\mathbb{R}^d;\mathbb{R}^d)$ via $\|D_R\|_\infty := \sup_{x\in\mathbb{R}^d}\|D_R(x)\|_2$ where $\|\cdot\|_2$ is the Euclidean operator norm on $\mathbb{R}^{d\times d}$.

for each $0 \le s \le t \le 1$, see e.g. [104, Sect. 1.3] and [54, Prop. 6.4], we find from (5.137) that

$$\big|\mathfrak{sig}_w(\tilde{Y}) - \mathfrak{sig}_w(Y)\big| \;\le\; c_p(\|Y\|_{p\text{-var}} + \|\tilde{Y}\|_{p\text{-var}})\|Z\|_{p\text{-var}} \le 2c_p\Xi\|Z\|_{p\text{-var}} \quad \text{if} \quad |w| = 2, \quad \text{and}$$
$$\big|\mathfrak{sig}_w(\tilde{Y}) - \mathfrak{sig}_w(Y)\big| \;\le\; (2c_p\Xi\|Y\|_{1\text{-var}} + \frac{c_p}{2!}\|\tilde{Y}\|^2_{1\text{-var}})\|Z\|_{p\text{-var}} \;\le\; \tfrac{5}{2}c_p\Xi^2\|Z\|_{p\text{-var}} \quad \text{if} \quad |w| = 3,$$

where for the penultimate inequality we used the standard estimate $\|\pi_2[\mathfrak{sig}_{s,t}(\tilde{Y})]\|_{1\text{-var}} \le \frac{1}{2!}\|\tilde{Y}\|^2_{1\text{-var}}$ (see e.g. [104, Prop. 2.2]). Combined with (5.135), we obtain that for $|w| = 2, 3$,

$$\big|\langle\tilde{Y}\rangle_w - \langle Y\rangle_w\big| \;\le\; (1 + |w|/2)c_p\mathbb{E}\Big[\Xi^{|w|-1}\|Z\|_{p\text{-var}}\Big] \;\le\; (2 + |w|)c_p\mathbb{E}\Big[\Xi^{|w|-1+1/p}\Big]\cdot\eta^\alpha,$$

as claimed in (5.136). To finally arrive at (5.129), notice that $\Xi^{|w|-1+1/p} \le 1 + \Xi^{2+1/p} \le 1 + (\|\tilde{Y}\|_{1\text{-var}} + \|Y\|_{1\text{-var}})^\beta$ for each $|w| = 2, 3$ and $\beta := 2 + 1/p$ and hence, by Minkowski,

$$\mathbb{E}\Big[\Xi^{|w|-1+1/p}\Big] - 1 \;\le\; \Big[\mathbb{E}\big[(\tilde{U} + U)^\beta\big]^{1/\beta}\Big]^\beta \;\le\; \Big[\mathbb{E}\big[\tilde{U}^\beta\big]^{1/\beta} + \mathbb{E}\big[U^\beta\big]^{1/\beta}\Big]^\beta \;\le\; 2^\beta K$$

for $\tilde{U} := \|\tilde{Y}\|_{1\text{-var}}$ and $U := \|Y\|_{1\text{-var}}$, as asserted.

(ii): Let $\mu, \tilde{\mu} \in \mathfrak{D}$ and suppose, as required, that there is $\kappa \ge 0$ and $\eta > 0$ such that

$$\max_{|w|=2,3}\big|\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w\big| \le \kappa\eta. \tag{5.138}$$

The bound (5.130) then follows from (5.138) and the definitions (5.55) and (5.56), as shown next.

We use the representation (5.57) of δ for convenience. As is conventional for direct sums, we thereby write $C \equiv C_0 \boldsymbol{+} \ldots \boldsymbol{+} C_d$ for any $C \equiv (C_0, \cdots, C_d) \in \mathcal{V}$ (boldfaced '+'), and accordingly set $N_1 \cdot C \cdot N_2 \equiv \boldsymbol{\sum}_{\nu=0}^d N_1 C_\nu N_2 \equiv (N_1 C_\nu N_2)_{\nu=0}^d \in \mathcal{V}$ for any $N_1, N_2 \in \mathbb{R}^{d\times d}$. Then

$$\delta(\mu, \tilde{\mu}) \;=\; \Big\|N_\mu^{-1}(C_0 \boldsymbol{+} \boldsymbol{\sum}_{k=1}^d C_k)N_\mu^{-1}\Big\|_{\mathcal{V}} \qquad \text{with} \qquad C_\nu = \left(\frac{\langle\tilde{\mu}\rangle_{ij\nu} - \langle\mu\rangle_{ij\nu}}{\sqrt{\langle\mu\rangle_{\nu\nu}}}\right)_{ij}$$

where $\langle\mu\rangle_{ij0} := \langle\mu\rangle_{ij}$, $\langle\mu\rangle_{00} := 1$ and N_μ as in (5.57); we further set $C_{\mu\tilde{\mu}} := C_0 \boldsymbol{+} \boldsymbol{\sum}_{k=1}^d C_k$. The assumed core-integrability of $\{\mu, \tilde{\mu}\} \subset \mathfrak{D}$ ensures that all of these quantities exist.

The bound (5.130) is based on the simple triangle inequality

$$\begin{aligned} |\delta_\perp(\tilde{\mu}) - \delta_\perp(\mu)| \;&=\; \Big|\|N_{\tilde{\mu}}^{-1}C_{\tilde{\mu}_\star\tilde{\mu}}N_{\tilde{\mu}}^{-1}\|_{\mathcal{V}} - \|N_\mu^{-1}C_{\mu_\star\mu}N_\mu^{-1}\|_{\mathcal{V}}\Big| \\ &\le\; \Big\|N_{\tilde{\mu}}^{-1}\Big[C_{\tilde{\mu}_\star\tilde{\mu}}R - R^{-1}C_{\mu_\star\mu}\Big]N_\mu^{-1}\Big\|_{\mathcal{V}} \qquad \text{for} \quad R := N_{\tilde{\mu}}^{-1}N_\mu, \end{aligned} \tag{5.139}$$

where for a stream $\mu = (\mu^i) \in \mathfrak{D}$ we denote by $\mu_\star := \mu^1 \otimes \cdots \otimes \mu^d$ its associated product signal. By definition the matrices $R \equiv (R_{ij})$ are diagonal with $R_{ii} = \sqrt{\langle\mu\rangle_{ii}/\langle\tilde{\mu}\rangle_{ii}}$. Moreover, if $\tilde{\mu}, \mu$ are both mean-stationary, then so are their respective product signals $\tilde{\mu}_\star, \mu_\star$, whence Lemma 5.5.9 gives that the (flattened) tensors $C_{\mu_\star\mu} =: (C_{ijk})$ and $C_{\tilde{\mu}_\star\tilde{\mu}} =: (\tilde{C}_{ijk})$ read

$$C_{ij0} = (1 - \delta_{ij})\langle\mu\rangle_{ij} \quad \text{and} \quad C_{ijk} = \frac{(1 - \delta_{ijk})\langle\mu\rangle_{ijk}}{\sqrt{\langle\mu\rangle_{kk}}} \qquad (i, j, k \in [d]), \tag{5.140}$$

and likewise for $(\tilde{C}_{ijk})$. If $\tilde{\mu}$ or μ are not mean-stationary, then all of the above remains valid as stated but with (5.140) holding by definition of $\delta_{\perp\!\!\perp}(\cdot)$, see Remark 5.5.6. The flattened [along its 3rd dimension, i.e. embedded in $\mathcal{V}$] tensor $M \equiv (m_{ijk}) := N_{\tilde{\mu}}^{-1}(C_{\tilde{\mu}\star\tilde{\mu}}R - R^{-1}C_{\mu\star\mu})N_{\mu}^{-1}$ thus reads

$$m_{ij0} = \frac{C_{ij0}R_{jj} - \tilde{C}_{ij0}R_{ii}^{-1}}{\sqrt{\langle\tilde{\mu}\rangle_{ii}\langle\mu\rangle_{jj}}} = \left[\frac{\langle\mu\rangle_{ij}}{\sqrt{\langle\tilde{\mu}\rangle_{ii}\langle\tilde{\mu}\rangle_{jj}}} - \frac{\langle\tilde{\mu}\rangle_{ij}}{\sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}}}\right](1-\delta_{ij}), \qquad \text{and}$$

$$m_{ijk} = \frac{C_{ijk}R_{jj} - \tilde{C}_{ijk}R_{ii}^{-1}}{\sqrt{\langle\tilde{\mu}\rangle_{ii}\langle\mu\rangle_{jj}}} = \left[\frac{\langle\mu\rangle_{ijk}}{\sqrt{\langle\tilde{\mu}\rangle_{ii}\langle\tilde{\mu}\rangle_{jj}\langle\mu\rangle_{kk}}} - \frac{\langle\tilde{\mu}\rangle_{ijk}}{\sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}\langle\tilde{\mu}\rangle_{kk}}}\right](1-\delta_{ijk})$$

Writing $\langle\mu\rangle_{00} := 1$ and $\langle\mu\rangle_{ij0} := \langle\mu\rangle_{ij}$, as well as $\beta_{ijk} := \sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}\langle\tilde{\mu}\rangle_{kk}}$ and $\tilde{\beta}_{ijk} := \sqrt{\langle\tilde{\mu}\rangle_{ii}\langle\tilde{\mu}\rangle_{jj}\langle\mu\rangle_{kk}}$ and $\gamma_{ijk} := \min\{\beta_{ijk}, \tilde{\beta}_{ijk}\}$ and $\varsigma_{ijk} := \max\{|\langle\mu\rangle_{ijk}|, |\tilde{\beta}_{ijk}|\}$, we find that

$$|m_{ijk}| = \left|\frac{\langle\mu\rangle_{ijk}\beta_{ijk} - \langle\tilde{\mu}\rangle_{ijk}\tilde{\beta}_{ijk}}{\beta_{ijk}\tilde{\beta}_{ijk}}\right| \le \frac{\varsigma_{ijk}}{\gamma_{ijk}^2}\Big(|\beta_{ijk} - \tilde{\beta}_{ijk}| + |\langle\mu\rangle_{ijk} - \langle\tilde{\mu}\rangle_{ijk}|\Big). \tag{5.141}$$

Writing $\nu_j := \langle\mu\rangle_{ii}$ and $\tilde{\nu}_j := \langle\tilde{\mu}\rangle_{jj}$, for the first of the above bounding summands we find[37]

$$\begin{aligned}\big|\beta_{ijk} - \tilde{\beta}_{ijk}\big|^2 &\le \big|\nu_i\nu_j\tilde{\nu}_k - \tilde{\nu}_i\tilde{\nu}_j\nu_k\big| \le \nu_i\nu_j\big|\tilde{\nu}_k - \nu_k\big| + \nu_k\big|\nu_i\nu_j - \tilde{\nu}_i\tilde{\nu}_j\big| \\ &\le \nu_i\nu_j\big|\tilde{\nu}_k - \nu_k\big| + \nu_k\big(\nu_i|\nu_j - \tilde{\nu}_j| + \tilde{\nu}_j|\nu_i - \tilde{\nu}_i|\big).\end{aligned}$$

Hence after denoting $\vartheta := \max\{\nu_i\nu_j, \nu_i\tilde{\nu}_j \mid i,j \in [d]\}$ and by (5.138), we obtain that

$$\big|\beta_{ijk} - \tilde{\beta}_{ijk}\big| \le \sqrt{3\vartheta\kappa}\cdot\eta^{\frac{1}{2}}. \tag{5.142}$$

Defining $\xi_{ijk} := \sqrt{2}\varsigma_{ijk}/\gamma_{ijk}^2$, we can now combine (5.142) and (5.138) with (5.141) to arrive at

$$m_{ijk}^2 \le \xi_{ijk}^2\big(|\beta_{ijk} - \tilde{\beta}_{ijk}|^2 + |\langle\mu\rangle_{ijk} - \langle\tilde{\mu}\rangle_{ijk}|^2\big) \le K_1\xi_{ijk}^2\cdot(\eta + \eta^2)$$

for $K_1 := (3\vartheta + \kappa)\kappa$. Using (5.139), this allows us to conclude

$$\big|\delta_{\perp\!\!\perp}(\tilde{\mu}) - \delta_{\perp\!\!\perp}(\mu)\big|^2 \le \sum_{k=0}^{d}\sum_{i,j=1}^{d} m_{ijk}^2 \le K_2\cdot(\eta + \eta^2)$$

for $K_2 := K_1\sum_{k=0}^{d}\sum_{i,j=1}^{d}\xi_{ijk}^2$, which gives (5.130) as claimed.

(iii): Given $\delta(\mu,\tilde{\mu}) \le \eta$, we know from Def. 5.5.1 that for each $w \equiv ijk \in [d]^{\star}_{|\cdot|=2,3}$ we have

$$\big|\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w\big| \le \sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}\langle\mu\rangle_{kk}}\cdot\eta \le \mathfrak{m}_\mu\eta \tag{5.143}$$

[37] Recall that $|\sqrt{a} - \sqrt{b}\,| \le \sqrt{|a-b|}$ for any $a, b \ge 0$.

for $\mathfrak{m}_\mu = \max\{\sqrt{\langle\mu\rangle_{ii}\langle\mu\rangle_{jj}\langle\mu\rangle_{kk}} \mid i,j \in [d], k \in [d]_0\}$. Hence by the very same computations as in point (ii) (upon replacing assumption (5.138) by (5.143)) we find

$$\big|\delta_{\perp\!\!\!\perp}(\tilde{\mu}) - \delta_{\perp\!\!\!\perp}(\mu)\big|^2 \;\leq\; \max\{3\vartheta_{\mu,\tilde{\mu}}\mathfrak{m}_\mu, \mathfrak{m}_\mu^2\} \sum_{k=0}^{d}\sum_{i,j=1}^{d} \xi_{ijk}^2 \cdot (\eta + \eta^2) \tag{5.144}$$

with $\vartheta_{\mu,\tilde{\mu}}$ and ξ_{ijk} as defined previously. If $\delta_{\perp\!\!\!\perp}(\mu) = 0$ (cf. Lemma 5.5.9), this gives the claim.

(iv) : Let $w \in [d]^\star$ be an arbitrary multiindex. Then by the change-of-variables formula,

$$\begin{aligned} |\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w| \;&=\; \left|\int_{\mathcal{C}_d} \mathfrak{sig}_w(x + R(x))\,\mu(\mathrm{d}x) - \int_{\mathcal{C}_d} \mathfrak{sig}_w(x)\,\mu(\mathrm{d}x)\right| \\ &\leq\; \int_{\mathcal{C}_d} \big|\mathfrak{sig}_w(x + R(x)) - \mathfrak{sig}_w(x)\big|\,\mu(\mathrm{d}x). \end{aligned} \tag{5.145}$$

Now by definition of $\|\cdot\|_{1\text{-var}}$ we find (with the sup taken over all dissections (t_ν) of $[0,1]$)

$$\|R(x)\|_{1\text{-var}} = \sup_{(t_\nu)} \sum_\nu \big|R(x_{t_{\nu+1}}) - R(x_{t_\nu})\big| \;\leq\; \|D_R\|_\infty \|x\|_{1\text{-var}} \tag{5.146}$$

and hence – again by [54, Proposition 5.5 (i)], just as for (5.135) – that

$$\|R(x)\|_{p\text{-var}} \;\leq\; 2^\alpha \|D_R\|_\infty^{1/p} \|x\|_{1\text{-var}}^{1/p} \|R\|_\infty^\alpha \;=\; 2^\alpha \|x\|_{1\text{-var}}^{1/p} \|D_R\|_\infty^\gamma \cdot \big(\|D_R\|_\infty \|R\|_\infty\big)^\alpha. \tag{5.147}$$

for $\gamma := 1/p - \alpha > 0$. From the estimates of point (i) – with $\tilde{Y}(\omega)$ and $Y(\omega)$ replaced by $x + R(x)$ and x, respectively – combined with (5.146), we obtain for $|w| = 2, 3$ that

$$\big|\mathfrak{sig}_w(x + R(x)) - \mathfrak{sig}_w(x)\big| \;\leq\; (1 + |w|/2)c_p(1 + \|D_R\|_\infty)^{|w|-1} \|x\|_{1\text{-var}}^{|w|-1} \|R(x)\|_{p\text{-var}}. \tag{5.148}$$

Together with (5.145) and (5.147), the above combines to

$$|\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w| \;\leq\; \tilde{c}_{p,|w|}\mathbb{E}_\mu\big[\|x\|_{1\text{-var}}^{|w|-1+1/p}\big] \cdot \phi_{|w|}(\|D_R\|_\infty)\big(\|D_R\|_\infty \|R\|_\infty\big)^\alpha$$

for $\tilde{c}_{p,|w|} := 2^\alpha c_p(1 + |w|/2)$ and $\phi_k(u) := (1 + u)^{k-1}u^\gamma$. Since $\mathbb{E}_\mu\big[\|x\|_{1\text{-var}}^{|w|-1+1/p}\big] \leq 1 + \mathbb{E}_\mu\big[\|x\|_{1\text{-var}}^\beta\big] =: 1 + \tilde{K}_p$ for $\beta = 2 + 1/p$, we thus find for each $w \in [d]^\star$ with $|w| = 2, 3$ that

$$|\langle\tilde{\mu}\rangle_w - \langle\mu\rangle_w| \;\leq\; \tilde{c}_{p,|w|}(1 + \tilde{K}_p) \cdot \phi_{|w|}(\|D_R\|_\infty)\big(\|D_R\|_\infty \|R\|_\infty\big)^\alpha \tag{5.149}$$

as claimed. Given (5.149), the last inequality (5.134) then holds by (ii) for $\kappa := \tilde{c}_{p,3}(1+\tilde{K}_p)$ and $\eta := \phi_3(\|D_R\|_\infty)(\|D_R\|_\infty \|R\|_\infty)^\alpha$. □

5.8.2 Some Technical Lemmas

5.8.2.1 The Matrix C_μ is Generically Invertible

Lemma 5.8.2. *Let $C_\mu := \frac{1}{2}([\mu]_0 + [\mu]_0^\intercal)$ for a signal $\mu \in \mathfrak{D}$ with 0^{th}-coredinate $[\mu]_0$. Then $C_\mu \notin \mathrm{GL}_d$ only if $\mu_{0,1}(H) = 1$ for some hyperplane H in $\mathbb{R}^d$.*

Proof. Suppose that $C_\mu \notin \mathrm{GL}_d$, which implies that there is $v = (v_i) \in \mathbb{R}^d \backslash \{0\}$ with $C_\mu v = 0$. By recalling the definition (5.39) of $[\mu]_0$ and from a quick computation of the C_μ-defining sums of integrals (5.38), we see that $C_\mu = \frac{1}{2}\big(\mathbb{E}[\mu^i_{0,1} \cdot \mu^j_{0,1}]\big)_{ij}$. Therefore $0 = v^\intercal C_\mu v = \sum_{i,j=1}^d v_i C_\mu^{ij} v_j = \mathbb{E}[Z_v^2]$ for $Z_v := \sum_{i=1}^d \mu^i_{0,1} v_i = \mu_{0,1}^\intercal \cdot v$ ($:= \varphi_*(\mu)$ for $\varphi := \langle \pi_{0,1}(\cdot), v\rangle_2$, see (5.4) and (5.5)). This implies $Z_v = \delta_0$ and hence $\mu_{0,1}(H) = 1$ for the hyperplane $H := \{u \in \mathbb{R}^d \mid u^\intercal v = 0\}$. □

5.8.2.2 Hyperplanes, Inverses, and Contracted Tensors

The following auxiliary observations facilitate the analysis that underlies the proof of Theorem 5.6.10.

Lemma 5.8.3. *Let $\vartheta \in \mathbb{R}^{d\times d}$ be invertible and $(e_i)_{i\in[d]}$ the standard basis of $\mathbb{R}^d$. Then the set*

$$\{u \in \mathbb{R}^d \mid \exists\, i, j \in [d],\, i \neq j \,:\, \langle \vartheta^\intercal u, e_i\rangle = \kappa_{ij}\langle \vartheta^\intercal u, e_j\rangle\} \tag{5.150}$$

has Lebesgue measure zero for any $\kappa \equiv (\kappa_{ij}) \in \mathbb{R}^{d\times d}$.

Proof. For any fixed $\kappa \equiv (\kappa_{ij}) \in \mathbb{R}^{d\times d}$, consider the sets

$$H_{ij} := \{u = (u_\nu) \in \mathbb{R}^d \mid u_i = \kappa_{ij} \cdot u_j\}$$

for $i, j \in [d]$ with $i \neq j$. As $H_{ij} = \ker(\varphi_{ij})$ for the non-zero linear map $\varphi_{ij} : u \mapsto u_i - \kappa_{ij} \cdot u_j$, each set H_{ij} has codimension one (by the rank-nullity theorem) and hence is a Lebesgue nullset in $\mathbb{R}^d$. The union $H := \bigcup_{i,j\in[d], i\neq j} H_{ij}$ is thus also a Lebesgue nullset in $\mathbb{R}^d$, and hence so is its diffeomorphic image $(\vartheta^\intercal)^{-1}(H)$. The latter equals (5.150), which proves the claim. □

The following two lemmas apply for $\|\cdot\|$ any submultiplicative matrix norm.

Lemma 5.8.4. *Let $\Lambda \in \mathrm{GL}_d(\mathbb{R})$ and $\Delta \in \mathbb{R}^{d\times d}$ with $\epsilon := \|\Lambda^{-1}\Delta\| < 1$. Then $\Lambda + \Delta$ is invertible with $\big\|(\Lambda + \Delta)^{-1} - \Lambda^{-1}\big\| \leq \|\Lambda^{-1}\|\frac{\epsilon}{1-\epsilon}$.*

Proof. Abbreviating $\rho := -\Lambda^{-1}\Delta$, notice that the invertibility of $\Lambda + \Delta$ is due to $\Lambda + \Delta = \Lambda(\mathrm{I} - \rho)$ and the convergence of the Neumann series (e.g. [74, Ex. 5.6.P26]). In fact

$$(\Lambda + \Delta)^{-1} = (\mathrm{I} - \rho)^{-1}\Lambda^{-1} = (\mathrm{I} + \textstyle\sum_{n=1}^\infty \rho^n)\Lambda^{-1}$$

so that, by $\epsilon < 1$ and for $\eta := \|\Lambda^{-1}\|$, we have $\big\|(\Lambda+\Delta)^{-1} - \Lambda^{-1}\big\| \leq \eta \sum_{n=1}^\infty \|\rho\|^n \leq \frac{\eta\epsilon}{1-\epsilon}$. □

Lemma 5.8.5. *Let $C \equiv (C_0, \ldots, C_d) \in \mathcal{V}$ and denote $\Lambda_\nu := \mathrm{ddiag}[C_\nu]$ and $\underline{C} := (0, C_1, \ldots, C_d)$ as well as $\underline{\Lambda} := (0, \Lambda_1, \ldots, \Lambda_d)$. If $\Lambda_0 \in \mathrm{GL}_d(\mathbb{R})$ and $\delta_1 := \|\Lambda_0^{-1} \cdot (C_0 - \Lambda_0)\| < 1$, then C_0 is invertible and the contraction*

$$C_\odot^c := c_0 C_0^{-1} \cdot \textstyle\sum_{\nu=1}^d c_v C_\nu,$$

with $c \equiv (c_0, \underline{c}) \in \mathbb{R}^{1+d}$ arbitrary, is of the form

$$C_\odot^c = c_0 \Lambda_0^{-1} \Lambda_c + \Delta_c \quad \textit{with} \quad \|\Delta_c\| \leq \beta_c(\delta_1, \delta_2)$$

for the diagonal matrix $\Lambda_c := \sum_{\nu=1}^d c_\nu \Lambda_\nu$, the number $\delta_2 := \|\underline{C} - \underline{\Lambda}\|_{\mathcal{V}}$, and the bounding function $\beta_c(a, b) := |c_0|\eta \left[\|\Lambda_c\| + 1 + \gamma_{a,b}^c\right] \gamma_{a,b}^c$ with $\eta := \|\Lambda_0^{-1}\|$ and $\gamma_{a,b}^c := \max\left\{\frac{a}{1-a}, b|\underline{c}|\right\}$.

Proof. Fix any $c \equiv (c_0, \underline{c}) \in \mathbb{R}^{1+d}$. Then by definition, $C_\odot^c = c_0 \tilde{D}_0 \cdot D_1$ for the factors $\tilde{D}_0 := C_0^{-1}$ and $D_1 := \sum_{\nu=1}^d c_\nu C_\nu$, where $D_1 = \Lambda_c + \Delta_1$ with $\|\Delta_1\| \leq |\underline{c}|\delta_2$ by construction. Since $\delta_1 < 1$, Lemma 5.8.4 implies that $\tilde{D}_0 = \Lambda_0^{-1} + \tilde{\Delta}_0$ with $\|\tilde{\Delta}_0\| \leq \eta\tilde{\delta}_1$ for $\tilde{\delta}_1 := \frac{\delta_1}{1-\delta_1}$. From this we obtain that the 'off-diagonal' $\Delta_c := C_\odot^c - c_0\Lambda_0^{-1}\Lambda_c = c_0(\tilde{\Delta}_0\Lambda_c + \Lambda_0^{-1}\Delta_1 + \tilde{\Delta}_0\Delta_1)$ can be bounded by

$$\begin{aligned}
|c_0|^{-1}\|\Delta_c\| &\leq \|\tilde{\Delta}_0\|\|\Lambda_c\| + \eta\|\Delta_1\| + \|\tilde{\Delta}_0\|\|\Delta_1\| \\
&\leq \eta\tilde{\delta}_1\|\Lambda_c\| + \eta\tilde{\delta}_2 + \eta\tilde{\delta}_1\tilde{\delta}_2 \leq \eta\left[\|\Lambda_c\| + 1 + \bar{\delta}\right] \cdot \bar{\delta}
\end{aligned}$$

for $\tilde{\delta}_2 := |\underline{c}|\delta_2$ and $\bar{\delta} := \max\{\tilde{\delta}_1, \tilde{\delta}_2\}$, as desired. □

5.8.2.3 Eigengaps

Given a vector $w \equiv (w_1, \ldots, w_d) \in \mathbb{C}^d$, we call the *gap of w* the number

$$\gamma[w] := \min_{i,j \in [d] \,:\, i \neq j} |w_i - w_j| \qquad (d \geq 2)$$

and refer to $\ell(w) := \|w - (w_{i_0}, \ldots, w_{i_0})\|_2$, with $i_0 = \min\left[\arg\min_{i \in [d]} |w_i|\right]$, as the *spreadlength of w*.

Lemma 5.8.6. *Over all vectors in $\mathbb{R}^d$ of fixed spreadlength $\ell > 0$, the gap is maximised by any spreadlength-ℓ vector $w^* \equiv (w_i^*) \in \mathbb{R}^d$ such that $|w_{i+1}^* - w_i^*| = |w_{j+1}^* - w_j^*|$ for $i, j \in [d-1]$. In particular: for $k_d := \sqrt{\frac{1}{6}(d-1)d(2d-1)}$,*

$$\max_{w \in \mathbb{R}^d \,:\, \ell(w) = \ell} \gamma[w] = \ell / k_d. \tag{5.151}$$

Proof. Let $w = (w_1, \ldots, w_d) \in \mathbb{R}^d$ be any fixed vector of spreadlength $\ell(w) =: \ell > 0$. Since the gap γ is invariant under permutations and constant offsets (i.e., $\gamma[w] = \gamma[\tau(w) + (c, \ldots, c)]$ for any $c \in \mathbb{R}$ and $\tau \in S_d$), we may assume that $w_1 = 0$ and $w_\nu \leq w_{\nu+1}$ for all

$\nu \in [d-1]$ wlog. Suppose now that $|w_{i+1} - w_i| = |w_{j+1} - w_j| =: \delta$ for each $i, j \in [d-1]$. Then $w_2 = |w_2 - w_1| = \delta$ and likewise $w_\nu = \sum_{j=1}^{\nu-1}(w_{j+1} - w_j) = (\nu - 1)\delta$ for each $\nu \in [d]$, implying that

$$\ell^2 = \ell(w)^2 = \|w\|_2^2 = \delta^2 \textstyle\sum_{j=0}^{d-1} j^2 = \delta^2 k_d^2$$

and thus $\delta = \frac{\ell}{k_d}$. Also, since by the above assumption $|w_i - w_j| \geq |w_{j+1} - w_j| = \delta$ for any $i, j \in [d]$ with $i \neq j$, we find that $\gamma[w] = \delta$. Assuming now that there is $\tilde{w} \equiv (\tilde{w}_i) \in \mathbb{R}^d$ (with $0 = \tilde{w}_1 \leq \tilde{w}_2 \leq \ldots \leq \tilde{w}_d$ (wlog) as above) of spreadlength $\ell(\tilde{w}) = \ell$ such that $\gamma[\tilde{w}] > \delta \equiv \ell / k_d$, we deduce that

$$\ell^2 \;=\; \ell(\tilde{w})^2 = \sum_{i=1}^{d} \tilde{w}_i^2 = \sum_{i=1}^{d} \left[\sum_{j=1}^{i-1} |\tilde{w}_{j+1} - \tilde{w}_j| \right]^2 > \sum_{i=1}^{d} \left[(i-1)\delta\right]^2 = \delta^2 k_d^2 \;=\; \ell^2.$$

As this is a contradiction, (5.151) holds as claimed. $\square$

In the next two subsections, all eigenvalues are counted with their respective multiplicity, and $\kappa(B) := \|B\|\|B^{-1}\|$ denotes the condition number (for $\|\cdot\|$ a given norm) of a matrix B.

5.8.2.4 Perturbation Bounds on Non-Degenerate Eigenspaces

As an important ingredient to our robustness analysis, we cite the following bounds for simple eigenvalues and -vectors of perturbed (not necessarily symmetric) matrices.

Lemma 5.8.7. *Let $C, E \in \mathbb{C}^{d \times d}$ and $\tilde{C} := C + E$. Suppose that the eigenvalues $\lambda_1, \ldots, \lambda_d$ of C are pairwise distinct, let $v_1, \ldots, v_d$ be their associated unit-norm eigenvectors (i.e., $\|v_i\|_2 = 1$ for all $i \in [d]$) and $(u_1, \lambda_1), \ldots, (u_d, \lambda_d)$ be the associated 'relatively-normed' left-eigenpairs of C (i.e., $u_i^\intercal \cdot C = \lambda_i u_i^\intercal$ with $u_i^\intercal v_i = 1$, $i \in [d]$).*[38] *Provided*

$$\begin{aligned} \|E\|_{\mathrm{F}} \;&<\; \tfrac{1}{2} \min_{i \in [d]} \frac{s_i}{\kappa_i} \qquad \textit{with} \\ s_i := \sigma_{d-1}([\mathrm{I}_d - v_i v_i^\intercal]\,(\lambda_i \mathrm{I}_d - C)) \quad &\textit{and} \quad \kappa_i := \|u_i\|_2 + \sqrt{\|u_i\|_2^2 - 1}, \end{aligned} \tag{5.152}$$

where $\sigma_{d-1}(\cdot)$ denotes the 2^{nd}-smallest singular value of its argument, it holds that

$$\begin{aligned} &\forall\, i \in [d] \;:\; \exists\, \mathrm{EV} \;\; \tilde{v}_{\sigma(i)} \;\textit{of}\; \tilde{C} \quad \textit{such that} \\ &u_i^\intercal \tilde{v}_{\sigma(i)} = 1 \quad \textit{and} \quad \|\tilde{v}_{\sigma(i)} - v_i\|_2 \;\leq\; \frac{2\kappa_i}{s_i}\|E\|_{\mathrm{F}} \end{aligned} \tag{5.153}$$

for $\sigma : [d] \to [d]$ some assignment, where each associated eigenpair $(\tilde{\lambda}_j, \tilde{v}_j)$ of $\tilde{C}$ satisfies

$$|\tilde{\lambda}_{\sigma(i)} - \lambda_i| \;\leq\; |u_i^\intercal E v_i + \ell_E^i| \quad \textit{with} \quad |\ell_E^i| \;\leq\; \frac{2\|u_i\|_2 \kappa_i}{s_i}\|E\|_{\mathrm{F}}^2. \tag{5.154}$$

Proof. This is immediate from – and in fact basically a quote of – [88, Theorem 4.1 (ii)]. $\square$

With the above, we can formulate and prove the main result of Section 5.6.

[38] In the context of Lemma 5.8.7, $v^\intercal := \bar{v}^\intercal$ denotes the complex conjugate of a vector $v \in \mathbb{C}^d$.

5.8.2.5 Continuity and Bounds of Some Matrix Functions

Lemma 5.8.8. *The map* $\lambda : \mathrm{Sym}_d \to \mathbb{R}^d$, $\mathfrak{a} \mapsto (\lambda_1(\mathfrak{a}), \cdots, \lambda_d(\mathfrak{a}))$ *which assigns a symmetric* $d \times d$*-matrix to the vector of its eigenvalues (listed in descending order, counting multiplicities) and the 'inverse square root' function* $\mathcal{R} : \mathrm{Sym}_d^+ \to \mathrm{Sym}_d^+$, $\mathfrak{a} \mapsto \mathfrak{a}^{-1/2}$, *are both continuous. Furthermore, for* $\|\cdot\|$ *any unitarily-invariant and submultiplicative matrix norm,*

$$\left\|\mathcal{R}(\mathfrak{a}_1)^{-1} - \mathcal{R}(\mathfrak{a}_2)^{-1}\right\| \leq \left[\lambda_d^{1/2}(\mathfrak{a}_1) + \lambda_d^{1/2}(\mathfrak{a}_2)\right]^{-1} \cdot \|\mathfrak{a}_1 - \mathfrak{a}_2\| \tag{5.155}$$

for each $\mathfrak{a}_1, \mathfrak{a}_2 \in \mathrm{Sym}_d^+$. *In particular, for any fixed* $A \in \mathrm{GL}_d$ *and* $\mathcal{S}(\mathfrak{a}) := \frac{1}{2}(\mathfrak{a} + \mathfrak{a}^\intercal)$, *both of*

$$\text{the maps} \quad \mathbb{R}^{d\times d} \ni \mathfrak{a} \longmapsto \lambda(A \cdot \mathcal{S}(\mathfrak{a}) \cdot A^\intercal) \in \mathbb{R}^d \qquad \text{and} \tag{5.156}$$

$$\mathrm{Pos}_d \ni \mathfrak{a} \longmapsto \mathcal{R}(A \cdot \mathcal{S}(\mathfrak{a}) \cdot A^\intercal)^{-1} \in \mathbb{R}^{d\times d} \quad \text{are continuous.} \tag{5.157}$$

Moreover, for every $\mathfrak{a} \equiv (\tilde{a}_{ij}) \in \mathbb{R}^{d\times d}$ *with* $\min_{i\in[d]} |\tilde{a}_{ii}| > 0$ *and* $\mathcal{C}_\mathfrak{a} := A\mathcal{S}(\mathfrak{a})A^\intercal \in \mathrm{Sym}_d^+$ *and* $N_\mathfrak{a} := \mathrm{ddiag}(|\tilde{a}_{11}|, \cdots, |\tilde{a}_{dd}|)^{1/2}$, *and* $A_{\mathcal{R}(\mathcal{C}_\mathfrak{a})}$ *and* $\bar{A}_{\mathcal{R}(\mathcal{C}_\mathfrak{a})}$ *defined analogous to above, we have*

$$\|A^{-1}_{\mathcal{R}(C_\mathfrak{a})}\|_\mathrm{F} \leq \|A^{-1}\|_\mathrm{F} \max_{i\in[d]} \lambda_i^{1/2}(C_\mathfrak{a}) =: \beta_1(\mathfrak{a}; A), \tag{5.158}$$

$$\kappa_2(A_{\mathcal{R}(C_\mathfrak{a})}) \leq \frac{2\det(\mathrm{ddiag}[\lambda(C_\mathfrak{a})]^{1/2})\|A\|_\mathrm{F}^d}{\det(A)\big(\sqrt{d}\min_{i\in[d]} \lambda_i^{1/2}(C_\mathfrak{a})\big)^d} =: \beta_2(\mathfrak{a}; A). \tag{5.159}$$

Proof. These are simple consequences of classical facts and inequalities from matrix analysis.

Indeed: The continuity of the map λ (which is well-defined as the elements of its (closed) domain Sym_d are all diagonalisable with real eigenvalues) is simply a [specialised] restatement of the well-known fact that the spectrum of a matrix depends continuously on its coefficients, see e.g. [74, Appendix D]. The (continuity-implying; recall that matrix inversion is (locally Lipschitz) continuous) inequality (5.155) is precisely the Ando-Hemmen inequality [72, Theorem 6.2]. As to the remaining assertions: Note that (5.158) is immediate from the definition of A_R, the representation (5.90) and the standard inequality: $\|\mathfrak{a}\mathfrak{b}\|_\mathrm{F} \leq \|\mathfrak{a}\|_\mathrm{F}\|\mathfrak{b}\|_2$ for each $\mathfrak{a}, \mathfrak{b} \in \mathbb{R}^{d\times d}$. The latter combined with the bound in [61] gives (5.159). □

5.8.2.6 Continuity of Extremal Functions

Lemma 5.8.9. *Let* $\|\cdot\|$ *be any norm on* $\mathbb{R}^n$, *and* $\varphi : \mathbb{R}^n \to \mathbb{R}$ *be a continuous function. Let further* $\bar{B}_r := \{x \in \mathbb{R}^n \mid \|x\| \leq r\}$ *be the zero-centered closed* $\|\cdot\|$*-ball of radius* $r \geq 0$. *Then the functions* $\phi, \psi : [0, \infty) \to \mathbb{R}$ *given by* $\phi(r) := \min_{y\in\bar{B}_r} \varphi(y)$ *and* $\psi(r) := \max_{y\in\bar{B}_r} \varphi(y)$, *respectively, are both continuous.*

Proof. Since $\psi(r) = -\min_{y\in\bar{B}_r}(-\varphi(y))$ for each $r \geq 0$, it suffices to show continuity of ϕ.

Because the open rays are a subbase of the Euclidean (i.e., order) topology on $\mathbb{R}$, the continuity of ϕ follows if, for each $a \in \mathbb{R}$, the sets $\ell_a := \{r \mid \phi(r) < a\}$ and $\rho_a := \{r \mid \phi(r) > a\}$ are both open in $[0,\infty)$. To see this for ℓ_a, fix any $s \in \ell_a$. Then $m_s := \min_{y\in\bar{B}_s}\varphi(y) < a$ and, since φ is continuous, $\varphi(x_s) = m_s$ for some $x_s \in \bar{B}_s$. Denoting by $B_u(z)$ the open z-centered $\|\cdot\|$-ball of radius u, the continuity of φ further implies that there is $\delta_s > 0$ such that $\varphi(y) < a$ for each $y \in B_{\delta_s}(x_s)$, and hence $\phi(s') < a$ for all $s' \in (\|x_s\| - \delta_s/2, s] \cap [0,\infty)$. Since trivially $\phi(s') < a$ for all $s' \geq s$, we thus found $\phi|_{J_s} < a$ for the open (in $[0,\infty)$) set $J_s := (\|x_s\| - \delta_s/2, \infty) \cap [0,\infty) \ni s$. Hence $J_s \subset \ell_a$, which shows that ℓ_a is open.

The openness of ρ_a follows similarly: Fix any $s \in \rho_a$. Then $\phi(s) > a$ and, hence, clearly also $\phi(s') > a$ for each $s' \leq s$. Moreover: $\min_{y\in\partial B_s}\varphi(y) \geq \phi(s) > a$, from which the continuity of φ implies that for each $z \in \partial B_s$ there is $\delta_z > 0$ such that: $\varphi(y) > a$ for each $y \in B_{\delta_z}(z)$. Hence $\varphi|_{G_s} > a$ for the open superset $G_s := \bigcup_{z\in\partial B_s} B_{\delta_z}(z) \cup B_s$ of $\bar{B}_s$. Now since $\bar{B}_s$ is compact and G_s^{c} (i.e. the complement of G_s) is closed, the distance $\hat{\delta}_s := d(\bar{B}_s, G_s^{\mathrm{c}}) = \inf_{u\in\bar{B}_s, v\in G_s^{\mathrm{c}}}\|u - v\|$ is positive, which implies that $\bar{B}_{s+\delta} \subset G_s$ for any $0 \leq \delta < \hat{\delta}_s$. Hence $\phi|_{J'_s} > a$ for the open (in $[0,\infty)$) superset $J'_s := (-\infty, s + \hat{\delta}_s) \cap [0,\infty)$ of s, which shows that ρ_a is open, as desired. □

5.8.3 Proof of Theorem 5.6.10

This subsection shows assertion (5.105) and thus completes the proof of Theorem 5.6.10.

For $A \equiv (a_1|\cdots|a_d) \in \mathrm{GL}_d$ and $\zeta \equiv (\zeta^1, \ldots, \zeta^d) \in \mathfrak{D}$, suppose

$$\chi = A\cdot\zeta \quad \text{and} \quad \varepsilon \equiv \delta_{\perp\!\!\!\perp}(\zeta) < \varepsilon_0 \tag{5.160}$$

for ε_0 as in (5.104), as required. Upon rescaling ζ by N_ζ^{-1} and noting that for $\tilde{\zeta} := N_\zeta^{-1}\cdot\zeta$ we have $N_{\tilde{\zeta}} = \mathrm{I}$ and $\delta_{\perp\!\!\!\perp}(\tilde{\zeta}) = \delta_{\perp\!\!\!\perp}(\zeta)$, whence (5.160) still holds for (ζ, A) and $(\tilde{\zeta}, \tilde{A} := AN_\zeta)$ replaced, we obtain for each $\nu \in [d]$, see Lemma 5.6.5 and (5.93), that

$$\begin{aligned} \mathfrak{s}_0 &:= \mathfrak{x}_0(A^{-1}) = \mathfrak{x}_0(\tilde{A}^{-1}) = N_{\tilde{\zeta}}^{-1}[\tilde{\zeta}]_0 N_{\tilde{\zeta}}^{-1} = [\tilde{\zeta}]_0 \qquad \text{and} \\ \mathfrak{s}_\nu &:= \mathfrak{x}_\nu(A^{-1}) = \mathfrak{x}_\nu(\tilde{A}^{-1}) = [\tilde{\zeta}]_\nu/\sqrt{\langle\tilde{\zeta}\rangle_{\nu\nu}} = [\tilde{\zeta}]_\nu. \end{aligned} \tag{5.161}$$

Moreover, for the whitening matrix $R \equiv R_{(\chi)}$ from (5.89) and the preprocessed observable $\bar{\chi} := R\cdot\chi$, we find that for each $\nu \in [d]_0$ and every $\theta \in \Xi_1$,

$$N_{\theta\cdot\bar{\chi}} = \mathrm{I} \quad \text{and hence} \quad \mathfrak{x}_\nu(\theta R) = [\theta R\cdot\chi]_\nu. \tag{5.162}$$

Indeed: A direct computation of the iterated integrals (5.38) yields that for each $\mu \in \mathfrak{D}$, the matrix $C_\mu \equiv \frac{1}{2}([\mu]_0 + [\mu]_0^\intercal)$ is identical to the classical moment matrix $\big(\frac{1}{2}\mathbb{E}[\mu_{0,1}^i\mu_{0,1}^j]\big)$.

Hence and since $C_{\bar{\chi}} = \mathrm{I}$ by choice of $R_{(\chi)}$, we have for any $\theta \equiv (\theta_1|\cdots|\theta_d)^\intercal \in \Xi_1$ that

$$\begin{aligned} \langle \theta \cdot \bar{\chi} \rangle_{ii} \;&=\; \tfrac{1}{2}\mathbb{E}\Big[\big(\textstyle\sum_{j=1}^d \theta_{ij} \bar{\chi}^j_{0,1}\big)^2\Big] \;=\; \tfrac{1}{2}\displaystyle\sum_{j=1}^d \theta_{ij}^2 \mathbb{E}[(\bar{\chi}^j_{0,1})^2] + \sum_{j<k} \theta_{ij}\theta_{ik}\mathbb{E}[\bar{\chi}^j_{0,1}\bar{\chi}^k_{0,1}] \\ &=\; \sum_{j=1}^d \theta_{ij}^2 \;=\; \|\theta_i\|_2^2 \;=\; 1 \qquad \text{for each } i \in [d], \end{aligned}$$

whence $N_{\theta \cdot \bar{\chi}} = \mathrm{I}$ and hence $\mathfrak{x}_\nu(\theta R)_\nu = [\theta \cdot \chi]_\nu$ for all $\nu \in [d]_0$, as claimed. From (5.160) we find

$$\mathfrak{s}_\nu = \Lambda^\zeta_\nu + \Delta^\zeta_\nu \quad \text{with} \quad \Lambda^\zeta_\nu \text{ diagonal and } \|\Delta^\zeta_\nu\| \le \varepsilon \tag{5.163}$$

for each $\nu \in [d]_0$. Indeed: By definition (5.57) and Lemma 5.5.9 resp. Remark 5.5.6, we have

$$\varepsilon_0^2 \;>\; \varepsilon^2 = \delta_\perp(\zeta)^2 = \delta_\perp(\tilde{\zeta})^2 \;=\; \|\mathfrak{s}_0 - \Lambda^\zeta_0\|^2 + \sum_{\nu=1}^d \|\mathfrak{s}_\nu - \Lambda^\zeta_\nu\|^2 \tag{5.164}$$

for the diagonal matrices

$$\Lambda^\zeta_0 \;:=\; \mathrm{ddiag}([\tilde{\zeta}]_0) \;=\; N^2_{\tilde{\zeta}} \;=\; \mathrm{I} \quad \text{and} \quad \Lambda^\zeta_\nu \;:=\; \big(\langle\tilde{\zeta}\rangle_{\nu\nu\nu} \cdot \delta_{ij\nu}\big)_{ij}\,; \tag{5.165}$$

setting $\Delta^\zeta_\nu := \mathfrak{s}_\nu - \Lambda^\zeta_\nu$ then gives (5.163) as asserted. Now from (5.160), the relations $\mathfrak{x}_\nu(A^{-1}) = \mathfrak{s}_\nu$ from (5.161) and the definition (5.95) of the contrast ϕ_χ and its scale invariance (5.96), we obtain

$$\min_{\theta \in \Xi_1} \phi_\chi(\theta) \;\le\; \phi_\chi(\bar{B}_R) = \phi_\chi(B_R) \;=\; \sum_{\nu=0}^d \|\mathfrak{s}_\nu\|^2_\times \;\le\; \sum_{\nu=0}^d \|\Delta^\zeta_\nu\|^2 \;\le\; \varepsilon^2\,, \tag{5.166}$$

where $B_R := (RA)^{-1} =: (b_1|\cdots|b_d)^\intercal$ and $\bar{B}_R := \mathrm{ddiag}(|b_1|,\ldots,|b_d|)^{-1} \cdot B_R \in \Xi_1$ scales the rows of B_R to unit length. Let now $\theta_\star \in \Xi_1$be arbitrary with

$$\phi_\chi(\theta_\star) \;\le\; \varepsilon^2\,, \qquad \text{and set} \quad \tilde{\theta}_\star := \theta_\star R \tag{5.167}$$

(note that $\theta_\star$ exists by (5.166) and the compactness of Ξ_1). By construction, cf. (5.162), we have

$$\mathfrak{x}_\nu(\tilde{\theta}_\star) = \big[\tilde{\theta}_\star \cdot \chi\big]_\nu \quad \text{for all } \nu \in [d]_0, \tag{5.168}$$

and from (5.167) and the definition of ϕ_χ we obtain for each $\nu \in [d]_0$ that

$$\|\Delta^\star_\nu\| \;\le\; \varepsilon \quad \text{with} \quad \Delta^\star_\nu \;:=\; \mathfrak{x}_\nu(\tilde{\theta}_\star) - \Lambda^\star_\nu \;\text{ and }\; \Lambda^\star_\nu := \mathrm{ddiag}(\mathfrak{x}_\nu(\tilde{\theta}_\star)).$$

Hence by (5.168) and the affine equivariance (5.88) of the signature coredinates, we have that

$$\begin{aligned} \tilde{\theta}_\star \cdot [\chi]_0 \cdot \tilde{\theta}_\star^\intercal \;&=\; \Lambda^\star_0 + \Delta^\star_0 \quad \text{and} \\ \tilde{\theta}_\star \cdot \Big[\textstyle\sum_{j=1}^d \tilde{\theta}_{\nu j}[\chi]_j\Big] \cdot \tilde{\theta}_\star^\intercal \;&=\; \Lambda^\star_\nu + \Delta^\star_\nu \quad \text{for all } \nu \in [d], \\ \text{as well as} \quad \|(\Lambda^\star_0)^{-1}\Delta^\star_0\| \;&=\; \|\Delta^\star_0\| \quad (\text{since } \Lambda^\star_0 = N^2_{\tilde{\theta}_\star \cdot \chi} = \mathrm{I}). \end{aligned} \tag{5.169}$$

Next we connect the relations (5.169) to the source statistics (5.161) via the coredinate identities (5.88) which hold by (5.160): Using (5.88) to substitute the matrices $[\chi]_\nu$ in (5.169), we obtain

$$\vartheta \cdot \diamondsuit_\nu^\zeta \cdot \vartheta^\intercal \;=\; \diamondsuit_\nu^\star \quad (\nu \in [d]_0) \quad \text{ with } \quad \vartheta := \tilde{\theta}_\star A \tag{5.170}$$

and where the $\diamondsuit_\nu^\zeta$, $\diamondsuit_\nu^\star$ are defined as the coredinate-based statistics

$$\diamondsuit_0^\zeta := [\zeta]_0 \quad \text{and} \quad \diamondsuit_\nu^\zeta := \textstyle\sum_{\gamma=1}^d \vartheta_{\nu\gamma}[\zeta]_\gamma \quad (\nu \in [d]), \quad \text{and} \quad \diamondsuit_\nu^\star := \Lambda_\nu^\star + \Delta_\nu^\star \quad (\nu \in [d]_0).$$

For an arbitrary vector $c \equiv (c_0, \underline{c}) \in \mathbb{R}_\times \times \mathbb{R}^d$, denoting $\underline{c} \equiv (\tilde{c}_1, \ldots, \tilde{c}_d)$, the congruences (5.170) can then be contracted to the 'condensed' relations

$$\vartheta \cdot c_0 \diamondsuit_0^\zeta \cdot \vartheta^\intercal \;=\; c_0 \diamondsuit_0^\star \quad \text{ and } \quad \vartheta \cdot \diamondsuit_{\underline{c}}^\zeta \cdot \vartheta^\intercal \;=\; \diamondsuit_{\underline{c}}^\star \tag{5.171}$$

which involve the $\underline{c}$-weighted linear combinations

$$\diamondsuit_{\underline{c}}^\zeta := \textstyle\sum_{\nu=1}^d \tilde{c}_\nu \diamondsuit_\nu^\zeta \quad \text{ and } \quad \diamondsuit_{\underline{c}}^\star := \textstyle\sum_{\nu=1}^d \tilde{c}_\nu \diamondsuit_\nu^\star.$$

As a consequence (upon 'left-multiplying the inverse of the first identity in (5.171) with the second identity in (5.171)') we obtain the $\underline{c}$-weighted *conjugate* relation

$$\vartheta^{-\intercal} \cdot \hat{\diamondsuit}_c^\zeta \cdot \vartheta^\intercal \;=\; \hat{\diamondsuit}_c^\star, \qquad \forall\, c \in \mathbb{R}_\times \times \mathbb{R}^d, \tag{5.172}$$

for the composite statistics $\hat{\diamondsuit}_c^\zeta := \tilde{c}_0 \big[\diamondsuit_0^\zeta\big]^{-1} \cdot \diamondsuit_{\underline{c}}^\zeta$ and $\hat{\diamondsuit}_c^\star := \tilde{c}_0 \big[\diamondsuit_0^\star\big]^{-1} \cdot \diamondsuit_{\underline{c}}^\star$, with $\tilde{c}_0 := c_0^{-1}$. By Lemma 5.8.5 (cf. (5.163), (5.169) and that $\varepsilon_0 < 1$), these composite statistics take the form

$$\hat{\diamondsuit}_c^\zeta \;=\; \hat{\Lambda}_c^\zeta \,+\, \hat{\Delta}_c^\zeta \quad \text{ and } \quad \hat{\diamondsuit}_c^\star \;=\; \hat{\Lambda}_c^\star \,+\, \hat{\Delta}_c^\star$$

for the diagonal matrices, notationally distinguished via $\eta \in \{\zeta, \star\}$,

$$\hat{\Lambda}_c^\eta \;:=\; \tilde{c}_0 \big[\Lambda_0^\eta\big]^{-1} \textstyle\sum_{\nu=1}^d \tilde{c}_\nu \bar{\Lambda}_\nu^\eta \qquad \text{with} \qquad \bar{\Lambda}_\nu^\star := \Lambda_\nu^\star \quad \text{ and } \quad \bar{\Lambda}_\nu^\zeta := \textstyle\sum_{\gamma=1}^d \vartheta_{\nu\gamma} \Lambda_\gamma^\zeta, \tag{5.173}$$

and off-diagonals $\hat{\Delta}_c^\eta \equiv \hat{\diamondsuit}_c^\eta - \hat{\Lambda}_c^\eta$ which are bounded by

$$\begin{aligned} \|\hat{\Delta}_c^\zeta\| \;&\leq\; |\tilde{c}_0| \Big[\|\Lambda_c^\zeta\|_2 + 1 + \gamma_{\epsilon_0^\zeta, \epsilon_1^\zeta}^c\Big] \cdot \gamma_{\epsilon_0^\zeta, \epsilon_1^\zeta}^c \quad \text{and} \\ \|\hat{\Delta}_c^\star\| \;&\leq\; |\tilde{c}_0| \Big[\|\Lambda_c^\star\|_2 + 1 + \gamma_{\epsilon_0^\star, \epsilon_1^\star}^c\Big] \cdot \gamma_{\epsilon_0^\star, \epsilon_1^\star}^c \end{aligned} \tag{5.174}$$

for $\Lambda_c^\zeta := \sum_{\nu=1}^d c_\nu \bar{\Lambda}_\nu^\zeta \; (= c_0 \hat{\Lambda}_c^\zeta)$ and $\Lambda_c^\star := c_0 \hat{\Lambda}_c^\star$ and $\gamma_{a,b}^c := \max\{\frac{a}{1-a}, |\underline{c}| \cdot b\}$, and the scalars

$$\epsilon_0^\eta := \|\Delta_0^\eta\| \quad \text{ and } \quad \epsilon_1^\eta := \|(0, \Delta_1^\eta, \ldots, \Delta_d^\eta)\|_{\mathcal{V}}.$$

Note that by (5.164) and (5.167) we have

$$(\epsilon_0^\eta)^2 \,+\, (\epsilon_1^\eta)^2 \;\leq\; \varepsilon^2 \quad \text{ for both } \; \eta = \zeta \; \text{ and } \; \eta = \star. \tag{5.175}$$

Let us now distinguish two cases.

First suppose that $\varepsilon = 0$, i.e. that $\delta_{\perp\!\!\!\perp}(\zeta) = 0$, meaning that the coredinates $([\zeta]_\nu)_{\nu\in[d]_0}$ of the source $\zeta \equiv (\zeta^1, \dots, \zeta^d)$ coincide with the coredinates of the product stream $\zeta^1 \otimes \cdots \otimes \zeta^d$. In this case (cf. (5.173),(5.174),(5.175)) the matrices $\hat{\Diamond}_c^\star$ and $\hat{\Diamond}_c^\zeta$ are both diagonal, the latter reading

$$\hat{\Diamond}_c^\zeta = \hat{\Lambda}_c^\zeta = \tilde{c}_0 \sum_{\gamma,\nu=1}^{d} \vartheta_{\nu\gamma}\tilde{c}_\nu \Lambda_\gamma^\zeta = \mathrm{ddiag}(\lambda_1^c, \lambda_2^c, \cdots, \lambda_d^c) \tag{5.176}$$

for c-dependent eigenvalues $\lambda_1^c, \dots, \lambda_d^c$ given by

$$\lambda_i^c = \alpha_i(c) \cdot \left(\langle\zeta\rangle_{iii}/\langle\zeta\rangle_{ii}^{3/2}\right) \quad \text{for} \quad \alpha_i(c) := \tilde{c}_0 \vartheta_i^\intercal \cdot \underline{c} \tag{5.177}$$

with $\vartheta =: (\vartheta_1|\cdots|\vartheta_d)$. Now since ϑ is invertible, Lemma 5.8.3 implies that the set

$$\mathcal{N} := \{c \in \mathbb{R}^{1+d} \mid \exists\, i, j \in [d],\, i \neq j \,:\, \lambda_i^c = \lambda_j^c\} \tag{5.178}$$

has Lebesgue measure zero, as follows from applying said lemma for the coefficient matrix

$$(\kappa_{ij}) := \left(\tfrac{\langle\zeta\rangle_{jjj}\langle\zeta\rangle_{ii}^{3/2}}{\langle\zeta\rangle_{jj}^{3/2}\langle\zeta\rangle_{iii}} \cdot (1 - \delta_{ii_0})\right) \tag{5.179}$$

with $i_0 \in [d]$ being the unique index at which $\zeta_{i_0 i_0 i_0} = 0$. (By assumption on ζ, the index i_0 is unique if it exists; otherwise, i.e. if none of the $\langle\zeta\rangle_{iii}$ vanish, set $(\kappa_{ij}) := \left(\tfrac{\langle\zeta\rangle_{jjj}\langle\zeta\rangle_{ii}^{3/2}}{\langle\zeta\rangle_{jj}^{3/2}\langle\zeta\rangle_{iii}}\right)$.) Since the complement of (5.178) is thus certainly non-empty, we can find a $c_\star \in \mathbb{R}_\times \times \mathbb{R}^d$ such that the eigenvalues $\lambda_1^{c_\star}, \dots, \lambda_d^{c_\star}$ of $\hat{\Diamond}_c^\zeta$ in (5.176) — which, by matrix conjugation (5.172), coincide with those of $\hat{\Diamond}_{c_\star}^\star$ — are pairwise distinct and hence all simple. As $\hat{\Diamond}_c^\zeta$ and $\hat{\Diamond}_{c_\star}^\star$ are both diagonal matrices, this allows us to conclude $\vartheta^\intercal \in \mathrm{M}_d$, whence

$$\vartheta \in \mathrm{M}_d \qquad \text{and hence} \qquad \tilde{\theta}_\star \in \mathrm{M}_d \cdot A^{-1}$$

as claimed in the theorem.

Next we consider the case $\varepsilon > 0$, i.e. the case where the components ζ^i of the source signal $\zeta = (\zeta^1, \dots, \zeta^d)$ are allowed to be statistically dependent. The identity (5.172) then reads

$$\tilde{\vartheta}^{-1} \cdot \left[\hat{\Lambda}_c^\zeta + E_c\right] \cdot \tilde{\vartheta} = \hat{\Lambda}_c^\star \quad \text{for} \quad \tilde{\vartheta} := \vartheta^\intercal, \tag{5.180}$$

with $\hat{\Lambda}_c^\zeta$ and $\hat{\Lambda}_c^\star$ as in (5.173) and for the off-diagonal matrix

$$E_c := \hat{\Delta}_c^\zeta - \tilde{\vartheta}\hat{\Delta}_c^\star\tilde{\vartheta}^{-1}. \tag{5.181}$$

Using (5.174), we observe that the 'perturbation' E_c can be controlled by

$$\|E_c\| \leq \|\hat{\Delta}_c^\zeta\| + \kappa_2(\tilde{\vartheta})\|\hat{\Delta}_c^\star\| \leq |\tilde{c}_0|\left[\kappa_c \cdot \hat{\varphi}_{\underline{c}}(\varepsilon) + (1 + \kappa_0)\hat{\varphi}_{\underline{c}}(\varepsilon)^2\right] \tag{5.182}$$

for the bounding function

$$\hat{\varphi}_{\underline{c}}(a) \;:=\; \max\left\{\tfrac{a}{1-a}, |\underline{c}|\cdot a\right\} \tag{5.183}$$

and the c-dependent constant

$$\kappa_c := \|\tilde{\vartheta}\|_{\varsigma}|\underline{c}| + 1 + \kappa_2(\tilde{\vartheta})\,[\xi d|\underline{c}| + 1] \le \|A_R\|_{\varsigma}\sqrt{d}|\underline{c}| + 1 + \kappa_0\kappa_2(A_R)\,[\xi d|\underline{c}| + 1]$$

for $\xi := \|(0, [\bar{\chi}]_1, \ldots, [\bar{\chi}]_d)\|_{\mathcal{V}}$ and $\varsigma := \sqrt{\sum_{i=1}^d \langle\zeta\rangle_{iii}^2/\langle\zeta\rangle_{ii}^3}$ and $A_R := RA$. Indeed: Writing ϑ_i for the i^{th} column of ϑ, Cauchy-Schwarz applied to (5.173), combined with (5.165), yields that

$$\|\Lambda_c^{\zeta}\| \;\le\; |\underline{c}|\|(0, (\bar{\Lambda}_\nu^{\zeta})_{\nu=1}^d)\|_{\mathcal{V}} \;=\; |\underline{c}|\sqrt{\textstyle\sum_{i=1}^d \|\vartheta_i\|_2^2\langle\tilde{\zeta}\rangle_{iii}^2} \;\le\; |\underline{c}|\|\tilde{\vartheta}\|_{\varsigma},$$

while the inclusion $\theta_\star \equiv (\theta_{ij}^\star) \in \Xi_1$ (unit rows) implies $\sum_{\alpha,\beta,\gamma=1}^d \left(\theta_{i\alpha}^\star\theta_{i\beta}^\star\theta_{\nu\gamma}^\star\right)^2 = 1$ and hence

$$\begin{aligned}\|\Lambda_c^\star\| \;&\le\; |\underline{c}|\|(0, (\mathrm{ddiag}[\mathfrak{r}_\nu(\theta_\star R)])_{\nu=1}^d)\|_{\mathcal{V}} \;=\; |\underline{c}|\sqrt{\textstyle\sum_{\nu,i=1}^d \left[\sum_{\alpha,\beta,\gamma=1}^d \theta_{i\alpha}^\star\theta_{i\beta}^\star\theta_{\nu\gamma}^\star\langle\bar{\chi}\rangle_{\alpha\beta\gamma}\right]^2}\\ &\le\; |\underline{c}|\sqrt{d^2\textstyle\sum_{\alpha,\beta,\gamma=1}^d\langle\bar{\chi}\rangle_{\alpha\beta\gamma}^2} = d|\underline{c}|\xi.\end{aligned}$$

The (second) inequality in (5.182) hence follows from (5.174) and $\max_{\eta\in\{\zeta,\star\}}\gamma^{c}_{\epsilon_0^\eta,\epsilon_1^\eta} \le \hat{\varphi}_{\underline{c}}(\varepsilon)$, where the latter holds by (5.175) and monotonicity. With the perturbation (5.181) of the matrix $\hat{\Lambda}_c^\zeta$ (cf. (5.176)) under control, we can now apply Lemma 5.8.7 to find contraction parameters $c = (c_0, \underline{c})$ which via (5.180) imply that $\tilde{\vartheta}$ is close to an element of M_d. To this end, recall from (5.176) and (5.177) (and $N_{\tilde{\zeta}} = \mathrm{I}$) that the diagonal entries $\underline{\lambda}_c^{\vartheta} := (\lambda_1^c, \ldots, \lambda_d^c)$ of $\hat{\Lambda}_c^\zeta$ read

$$\underline{\lambda}_c^{\vartheta} \;=\; \tfrac{1}{c_0}H_\vartheta\cdot\underline{c} \quad\text{ for }\quad H_\vartheta := \mathrm{ddiag}(\langle\tilde{\zeta}\rangle_{111}, \ldots, \langle\tilde{\zeta}\rangle_{ddd})\cdot\tilde{\vartheta},$$

where, upon permuting $\tilde{\zeta}^1, \ldots, \tilde{\zeta}^d$ if necessary, we may assume that the 'vanishing' index i_0 from (5.179) reads $i_0 = 1$. Now for a given $\alpha > 0$, which will be specified below, define

$$\underline{\lambda}_\alpha \;:=\; (0, \alpha/k_d, 2\alpha/k_d, \ldots, (d-1)\alpha/k_d) \equiv (\lambda_i^\alpha)_{i=1}^d \tag{5.184}$$

for the constant k_d defined in Lemma 5.8.6, and set further $c_\alpha := (1, \underline{c}_\alpha)$ with

$$\underline{c}_\alpha \;:=\; \tilde{\vartheta}^{-1}\mathrm{ddiag}(0, \langle\tilde{\zeta}\rangle_{222}^{-1}, \ldots, \langle\tilde{\zeta}\rangle_{ddd}^{-1})\cdot\underline{\lambda}_\alpha. \tag{5.185}$$

Clearly $|\underline{\lambda}_\alpha| = \alpha$ and $\underline{\lambda}_\alpha = \underline{\lambda}_{c_\alpha}^{\vartheta}$, and by Lemma 5.8.6 the matrix $\hat{\Lambda}_{c_\alpha}^\zeta$ is such that the minimal distance between any two of its eigenvalues (the latter are the entries of $\underline{\lambda}_\alpha$) is maximal among any matrix in $\{\hat{\Lambda}_c^\zeta \mid c \in \mathbb{R}_\times\times\mathbb{R}^d \;:\; |\mathrm{diag}(\hat{\Lambda}_c^\zeta)| = \alpha\}$. To simplify the following estimates, we choose α such that $\hat{\varphi}_{\underline{c}_\alpha}(\varepsilon)$ from (5.183) is dominated by its second argument, i.e. such that $\hat{\varphi}_{\underline{c}_\alpha}(\varepsilon) = |\underline{c}_\alpha|\varepsilon$. Clearly this holds for

$$\alpha > 0 \quad\text{ such that }\quad |\underline{c}_\alpha| \;=\; \frac{1}{1-\varepsilon}. \tag{5.186}$$

Note that an according choice of α is certainly possible since, cf. (5.185),

$$|\underline{c}_\alpha| = \alpha \cdot \kappa_\zeta^\vartheta \quad \text{for} \quad \kappa_\zeta^\vartheta := |\tilde{\vartheta}^{-1}[D_\zeta \underline{\lambda}_1]| > 0, \tag{5.187}$$

where $D_\zeta := \operatorname{ddiag}\big[(0, \langle\tilde{\zeta}\rangle_{222}^{-1}, \ldots, \langle\tilde{\zeta}\rangle_{ddd}^{-1})\big]$ and the positivity of κ_ζ^ϑ follows from the invertibility of $\tilde{\vartheta}^{-1}$ (as thus $\ker(\tilde{\vartheta}^{-1}) = \{0\} \not\ni D_\zeta \cdot \underline{\lambda}_1$). In respect of (5.182) and (5.186), we may bound $E_\alpha := E_{c_\alpha}$ by

$$\begin{aligned} \|E_\alpha\| \;&\le\; \kappa_{c_\alpha}|\underline{c}_\alpha| \cdot \varepsilon + (1+\kappa_0)|\underline{c}_\alpha|^2 \cdot \varepsilon^2 \\ &=\; \frac{\varepsilon}{(1-\varepsilon)^2}\Big[1 + \|\tilde{\vartheta}\|\varsigma + \kappa_2(\tilde{\vartheta})\big[\xi d + (1-\varepsilon)\big] + \kappa_0\varepsilon\Big] \\ &<\; \Big[1 + \kappa_0 + \|\tilde{\vartheta}\|\varsigma + (1+\xi d)\kappa_2(\tilde{\vartheta})\Big] \cdot \frac{\varepsilon}{(1-\varepsilon)^2}. \end{aligned} \tag{5.188}$$

Recalling $A_R \equiv RA$ and $\|\tilde{\vartheta}\| = \|\vartheta\| \le \sqrt{d}\|A_R\|$ (as $\vartheta = \theta_\star A_R$ with $\theta_\star \in \Xi_1$), we find that

$$K_\vartheta \;:=\; 1 + \kappa_0 + \|\tilde{\vartheta}\|\varsigma + (1+\xi d)\kappa_2(\tilde{\vartheta}) \tag{5.189}$$

$$\le\; 1 + \sqrt{d}\|A_R\|\varsigma + \big(1 + (1+\xi d)\kappa_2(A_R)\big)\kappa_0 \;=:\; \bar{K}_0 \tag{5.190}$$

since also $\kappa_2(\tilde{\vartheta}) \le \kappa_0\kappa_2(A_R)$ by the assumed condition number bound on $\theta_\star$. Next we tie in Lemma 5.8.7. Using the terminology of said lemma, note that by the above choice (5.185) of c_α the eigenvalues $\lambda_i := \lambda_i^{c_\alpha}$, $i \in [d]$, of the diagonal matrix $C := \hat{\Lambda}_{c_\alpha}^\zeta$ are pairwise distinct with associated unit-norm eigenvectors given by $(v_i)_{i\in[d]} := (e_i)_{i\in[d]}$, the standard basis of $\mathbb{R}^d$. Accordingly we have (in the notation of Lem. 5.8.7) $(u_i)_{i\in[d]} = (e_i^\intercal)_{i\in[d]}$ and, thus, $\kappa_i = 1$ and

$$s_i \;=\; \sigma_{d-1}\big((\delta_{\mu\nu} - \delta_{i\mu\nu})_{\mu\nu} \cdot (C - \lambda_i \mathrm{I})\big) \;=\; \min_{j\in[d]\setminus\{i\}} \big|\lambda_j - \lambda_i\big| = \alpha/k_d \tag{5.191}$$

for each $i \in [d]$ (cf. (5.184)). We observe that E_α ($\equiv E$) thus meets the hypothesis (5.152) of Lemma 5.8.7. Indeed: Notice first that the stronger (than (5.152)) inequality

$$\|E_\alpha\| \;<\; \gamma^{-1}\min_{i\in[d]}\frac{s_i}{\kappa_i} \;=\; \frac{\alpha}{\gamma k_d}\,, \qquad \gamma := 1 + \sqrt{5}, \tag{5.192}$$

follows – due to (5.187), (5.188) and (5.189) – from the (yet to be proved) inequality

$$\frac{\varepsilon}{1-\varepsilon} \;<\; \big(\gamma k_d \kappa_\zeta^\vartheta K_\vartheta\big)^{-1} =: q_\vartheta\,, \tag{5.193}$$

which in turn is equivalent to $\varepsilon < q_\vartheta/(1+q_\vartheta)$. But since by definition, cf. (5.187) and (5.190),

$$\kappa_\zeta^\vartheta \cdot K_\vartheta \;\le\; \kappa_\zeta^\vartheta(1+\kappa_0) + \kappa_2(\tilde{\vartheta})\varsigma\varsigma_1 + (1+\xi d)\kappa_\zeta^\vartheta\kappa_2(\tilde{\vartheta})$$

for the source-dependent constant $\varsigma_1 := \|D_\zeta \cdot \underline{\lambda}_1\| = k_d^{-1}\sqrt{\sum_{i=1}^d \frac{(i-1)^2}{\langle\tilde{\zeta}\rangle_{iii}^2}}$ and with

$$\kappa(\tilde{\vartheta}) \le \kappa_0\kappa_2(A_R) \quad \text{as well as} \quad \kappa_\zeta^\vartheta \;\le\; \|\tilde{\vartheta}^{-1}\|\varsigma_1 \;\le\; \frac{\kappa_0\|A_R^{-1}\|}{\sqrt{d}}\varsigma_1$$

(since $\|\tilde{\vartheta}^{-1}\| = \|\vartheta^{-1}\| = \|A_R^{-1}\|\|\theta_\star^{-1}\|$ and $\|\theta_\star^{-1}\| \leq \kappa_0/\|\theta_\star\| = \kappa_0/\sqrt{d}$), we find that

$$\kappa_\zeta^\vartheta K_\vartheta \;\leq\; r_0 := \kappa_0\varsigma_1\left[\frac{\|A_R^{-1}\|}{\sqrt{d}}(1+\kappa_0+(1+\xi d)\kappa_0\kappa_2(A_R)) + \kappa_2(A_R)\varsigma\right] \tag{5.194}$$

and hence $q_\vartheta \;\geq\; \big(\gamma k_d r_0\big)^{-1} =: q_0$ and thus

$$\frac{q_\vartheta}{1+q_\vartheta} \;\geq\; \frac{q_0}{1+q_0}. \tag{5.195}$$

But since $\varepsilon < \varepsilon_0 = q_0/(1+q_0)$ by assumption, we conclude that (5.195), and hence (5.193) and thus (5.192) and with it (5.152) in particular, hold as desired. This guarantees that Lemma 5.8.7 applies to the above matrix $C \equiv \hat{\Lambda}^\zeta_{c_\alpha}$ and its perturbation $\tilde{C} := C + E_\alpha$ (cf. (5.180)). Denote by $\tilde{\lambda}_1, \ldots, \tilde{\lambda}_d$ the eigenvalues of $\tilde{C}$ enumerated in the order (from left to right) in which they appear on the diagonal matrix $\hat{\Lambda}^\star_{c_\alpha}$ in (5.180). Since, in further consequence of (5.191) and (5.192),

$$\|E_\alpha\| + \frac{2\|u_i\|_2\kappa_i}{s_i}\|E_\alpha\|^2 \;<\; \frac{(2+\gamma)\alpha}{\gamma^2 k_d} = \frac{\alpha}{2k_d}\,,$$

Lemma 5.8.7 (cf. (5.154)) then implies that there is an assignment $\sigma : [d] \to [d]$ such that

$$\big|\tilde{\lambda}_{\sigma(i)} - \lambda_i\big| \;<\; \frac{\alpha}{2k_d} \qquad \text{for each } i \in [d]. \tag{5.196}$$

It follows that σ is in fact a permutation and the eigenvalues $\tilde{\lambda}_1, \ldots, \tilde{\lambda}_d$ of $\tilde{C}$ (all real) are pairwise distinct. Indeed: assuming otherwise that there are $i, j \in [d]$ with $i \neq j$ such that $\tilde{\lambda}_{\sigma(i)} = \tilde{\lambda}_{\sigma(j)}$ implies, due to (5.196), that

$$|\lambda_i - \lambda_j| \;\leq\; \big|\lambda_i - \tilde{\lambda}_{\sigma(i)}\big| + \big|\tilde{\lambda}_{\sigma(j)} - \lambda_j\big| \;<\; \alpha/k_d\,,$$

contradicting $\min_{i,j\in[d],\, i\neq j}|\lambda_i - \lambda_j| = \alpha/k_d$ (recall (5.184)). Hence for each $i \in [d]$, the eigenspace V_i of $\tilde{C}$ associated to $\tilde{\lambda}_i$ is of real dimension one, namely $V_i = \langle\tilde{\vartheta}_i\rangle_{\mathbb{R}}$ for $\tilde{\vartheta}_i$ the i^{th} column of $\tilde{\vartheta}$. Thus by Lemma 5.8.7 (5.153) we obtain that,[39] for each $i \in [d]$ and $\sigma \in S_d$ as above,

$$\big\|\beta_i^{-1}\cdot\tilde{\vartheta}_{\sigma(i)} - e_i\big\|_2 \;\leq\; \frac{2k_d}{\alpha}\|E_\alpha\| \quad \text{ for }\;\; \beta_i := \tilde{\vartheta}_{\sigma(i)i}. \tag{5.197}$$

Using (5.187) and (5.188), the right-hand side in (5.197) can be estimated to

$$\frac{2k_d}{\alpha}\|E_\alpha\| \;\leq\; r_1\frac{\varepsilon}{1-\varepsilon} \qquad \text{for} \qquad r_1 := 2k_d\kappa_\zeta^\vartheta K_\vartheta\,, \tag{5.198}$$

whence in partic. $\frac{2k_d}{\alpha}\|E_\alpha\| < 2/\gamma < 1$ by (5.193). As for an estimate of the β_i, notice first that

$$\beta_i \;=\; \max_{\nu\in[d]}\big|\tilde{\vartheta}_{\nu i}\big| \;\equiv\; \|\tilde{\vartheta}_i\|_\infty \qquad \text{for all } i \in [d],$$

[39] Since for every eigenpair $(\tilde{\lambda}_j, \tilde{w}_j)$ of $\tilde{C}$ we must have $\tilde{w}_j = q\cdot\tilde{\vartheta}_j$, $q \in \mathbb{R}_\times$, as $\tilde{\lambda}_j$ is simple, the eigenvector $\tilde{v}_j \equiv (\tilde{v}_j^1, \ldots, \tilde{v}_j^d)$ to $\tilde{\lambda}_j$ which has been singled out in (5.153) via $(\tilde{v}_j^i =) \, u_i^\intercal\tilde{v}_j = 1$ reads $\tilde{v}_j = \tilde{\vartheta}_{ij}^{-1}\cdot\tilde{\vartheta}_j$.

as otherwise $\max_{\nu\neq\sigma(i)} |\beta_i^{-1} \cdot \tilde{\vartheta}_{\nu i}| \geq 1$ and thus $\|\beta_i^{-1} \cdot \tilde{\vartheta}_{\sigma(i)} - e_i\|_2 \geq 1$ in contradiction to $\frac{2k_d}{\alpha}\|E_\alpha\| < 1$. Hence for each $i \in [d]$,

$$\beta_i \;\geq\; d^{-1/2}\|\tilde{\vartheta}_i\|_2 \;\geq\; \frac{\|A_R\|_2}{\kappa_2(A_R)\sqrt{d}}$$

where the first inequality is due to the equivalence of maximum and Euclidean norm and the last inequality follows from $\|\tilde{\vartheta}_i\|_2 = \|A_R^\intercal \cdot \theta_i\|_2 \;\geq\; \sigma_d(A_R) = \|A_R\|_2/\kappa_2(A_R)$, where we used that θ_i, which denotes the i^{th} row of $\theta_\star$, has Euclidean norm 1. Further,

$$\beta_i \;\leq\; \|\tilde{\vartheta}_i\|_2 = \|A_R^\intercal \cdot \theta_i\|_2 \;\leq\; \|A_R\|_2 \tag{5.199}$$

for each $i \in [d]$. Thus, the condition number of $\Lambda := \mathrm{ddiag}[\beta_1^{-1}, \ldots, \beta_d^{-1}]$ is bounded by

$$\kappa_2(\Lambda) \;\leq\; \sqrt{d}\kappa_2(A_R). \tag{5.200}$$

Combining (5.197) and (5.198), we obtain

$$\left\|\tilde{\vartheta} \cdot \tilde{M} - \mathrm{I}\right\| \;\leq\; r_2\frac{\varepsilon}{1-\varepsilon} \qquad \text{for} \qquad r_2 := d \cdot r_1, \tag{5.201}$$

where $\tilde{M} := P_\sigma^\intercal \cdot \Lambda \in \mathrm{M}_d$ for Λ as above and $P_\sigma = (\delta_{\sigma(i)j})$ the permutation matrix associated to σ. Hence for $M := \tilde{M}^{-\intercal} \in \mathrm{M}_d$ and $E := (\tilde{\vartheta}\tilde{M} - \mathrm{I})^\intercal$ we get

$$\vartheta \;=\; M(\mathrm{I} + E) \qquad \text{with} \qquad \|E\| \leq \tfrac{r_2\varepsilon}{1-\varepsilon} \tag{5.202}$$

and with $r_2 \;\leq\; 2dk_d r_0$, for r_0 as in (5.194); this proves (5.105). To see the last inequality in (5.105), we note first that since $\vartheta = \tilde{\theta}_\star A$ by definition, equation (5.202) yields

$$\tilde{\theta}_\star \;=\; M \cdot A^{-1} + ME \cdot A^{-1}. \tag{5.203}$$

In particular, we thus have obtained the relative-error bound

$$\frac{\|\tilde{\theta}_\star - MA^{-1}\|}{\|MA^{-1}\|} \;\leq\; \frac{\|MEM^{-1}\|_2\|MA^{-1}\|}{\|MA^{-1}\|} \;\overset{(5.200)}{\leq}\; r_3\frac{\varepsilon}{1-\varepsilon} \;=:\; \varepsilon' \tag{5.204}$$

where $r_3 := \sqrt{d}\kappa_2(A_R)r_2$ (recall (5.200) and (5.201)). This concludes proof. ◻

5.8.4 Proof of Theorem 5.6.14: The Continuity of $\mathfrak{I}$

Proof. By Remark 5.5.13 and because $(E, \mathbb{d})$ is a premetric space and $(F, \mathfrak{d})$ is a metric space, the map $\mathfrak{I}$ is continuous on $\mathscr{I}$ if (and only if; see Remark 5.5.3) for each $\mathfrak{p}_\star \in \mathscr{I}$ we have that:

$$\forall\, \varepsilon > 0 \,:\, \exists\, r > 0 \,:\, \forall\, \mathfrak{p} \in E \,:\; \text{if} \quad \mathbb{d}(\mathfrak{p}_\star, \mathfrak{p}) < r \quad \text{then} \quad \mathfrak{d}(\mathfrak{I}(\mathfrak{p}_\star), \mathfrak{I}(\mathfrak{p})) < \varepsilon. \tag{5.205}$$

To prove (5.205), fix any $\mathfrak{p}_\star \equiv (\mu_\star, A) \in \mathscr{I}$ and let $\varepsilon > 0$ be arbitrary.

Let further $0 < r < r_0$, for r_0 as in Lemma 5.5.14, and take any $\mathfrak{p} \equiv (\mu, f)$ in E with $\mathbb{d}(\mathfrak{p}_\star, \mathfrak{p}) < r$. By Lemma 5.5.14, this $\mathfrak{p}$ can then be assigned a unique signal μ_ζ in $\mathfrak{D}$ such that

$$\mu_\chi \equiv f_*\mu = A\mu_\zeta \quad \text{with} \quad \delta_{\perp\!\!\!\perp}(\mu_\zeta) \leq K' r^{\alpha/2} \quad \text{and} \quad \big\|[\mu_\zeta]_\mathfrak{c} - [\mu_\star]_\mathfrak{c}\big\|_\mathcal{V} \leq \hat{K} r^\alpha \tag{5.206}$$

for (explicit) constants $K' \equiv K_r := K_r(\mu_\star, K_\mathfrak{S}) > 0$, increasing in r, and $\hat{K} = \hat{K}(\mu_\star, K_\mathfrak{S}) > 0$.

By Section 5.6.1, the triple (μ_χ, μ_ζ, A) can be associated the identifiability-related parameters

$$\varepsilon(\mu_\zeta) := \frac{q(\mu_\zeta)}{1 + q(\mu_\zeta)} \quad \text{with} \quad q(\mu_\zeta) := \big(\gamma k_d p(\mu_\zeta)\big)^{-1} \tag{5.207}$$

for $\gamma = 1 + \sqrt{5}$, $k_d = \sqrt{\frac{d}{6}(d-1)(2d-1)}$ and the function $p(\nu) = p(\nu; \mathfrak{p}_\star)$ $(\nu \in \mathfrak{D})$ given by

$$p(\nu) = \tilde{\kappa}_0 \varsigma_1(\nu) \left[\frac{\|A_{R_\nu}^{-1}\|}{\sqrt{d}} \big(1 + \tilde{\kappa}_0 + (1 + \xi(\nu) d)\tilde{\kappa}_0 \kappa_2(A_{R_\nu})\big) + \kappa_2(A_{R_\nu})\varsigma(\nu) \right]$$

for $A_{R_\nu} := R_\nu A$ and the whitening matrix $R_\nu := \mathcal{R}(\mathcal{C}_{[\nu]_0})$ as in Lemma 5.8.8, writing $\mathcal{C}_\mathfrak{a} := A\mathcal{S}(\mathfrak{a})A^\intercal$ for the A fixed above, and with the auxiliary functions

$$\varsigma_1(\nu) := k_d^{-1} \sqrt{\textstyle\sum_{i=1}^d \Big(\frac{(i-1)\langle\nu\rangle_{ii}^{3/2}}{\langle\nu\rangle_{iii}}\Big)^2} \quad \text{and} \quad \varsigma(\nu) := \sqrt{\textstyle\sum_{i=1}^d \langle\nu\rangle_{iii}^2 / \langle\nu\rangle_{ii}^3} \tag{5.208}$$

and $\xi(\nu) := \sqrt{\sum_{j=1}^d \|[A_{R_\nu} \cdot \nu]_j\|^2}$ and $\tilde{\kappa}_0 := b + \Delta_\kappa$, cf. (5.111). Recall further from Remark 5.6.13 that there is a radius $r_* \equiv r_*(\mathfrak{p}_\star, \Delta_\kappa) > 0$ such that, for $\bar{B}_{(A,\nu)}$ as defined in that remark,

$$\bar{B}_{(A,\nu)} \in \tilde{\Xi}_1 \quad \text{if} \quad \big\|[\nu]_0 - [\mu_\star]_0\big\| \leq r_*. \tag{5.209}$$

Using the bound in (5.206) on the distance between $[\mu_\zeta]_\mathfrak{c}$ and $[\mu_\star]_\mathfrak{c}$, we see that the coefficients $(5.208)|_{\nu=\mu_\zeta}$ are well-defined – i.e. that $\langle\mu_\zeta\rangle_{ii} > 0$ and $\langle\mu_\zeta\rangle_{iii} > 0$ for all $i \in [d]$ – if

$$r \;<\; \big[(\rho_0 \wedge \varrho_0)/\hat{K}\big]^{1/\alpha} =: r_1, \qquad \text{where} \qquad \varrho_0 := \min_{i\geq 2} \langle\mu_\star\rangle_{iii}\,;$$

note that $\varrho_0 = \varrho_0(\mu_\star) > 0$ by definition of $\mathscr{I}$. As the denominators in (5.208) are bounded away from zero on $\{\nu \in \mathfrak{D} \mid \|[\nu]_\mathfrak{c} - [\mu_\star]_\mathfrak{c}\|_\mathcal{V} \leq \mathfrak{q}\} =: \tilde{B}_\mathfrak{q}$ for any $\mathfrak{q} < \rho_0 \wedge \varrho_0$, we see that the coefficient bounds $\mathfrak{K}_1(\mathfrak{q}) := \sup_{\nu\in\tilde{B}_\mathfrak{q}} \varsigma_1(\nu)$ and $\mathfrak{K}_2(\mathfrak{q}) := \sup_{\nu\in\tilde{B}_\mathfrak{q}} \varsigma(\nu)$ are both finite. Further, since the eigenvalues $(\tilde{\lambda}_i(\mathfrak{a}; A)) := \lambda(A\mathcal{S}(\mathfrak{a})A^\intercal)$ of the symmetrized matrix $\mathcal{C}_\mathfrak{a} = A\mathcal{S}(\mathfrak{a})A^\intercal$ depend continuously on $\mathfrak{a} \in \mathbb{R}^{d\times d}$ (Lemma 5.8.8), so does their minimum $\underline{\tilde{\sigma}}(\mathfrak{a}) := \min_{i\in[d]} \tilde{\lambda}_i(\mathfrak{a}; A)$. Hence and since $\sigma_\star \equiv \underline{\tilde{\sigma}}([\mu_\star]_0) > 0$ [recall that $\mathcal{C}_{[\mu_\star]_0}$ is congruent to $\mathcal{S}([\mu_\star]_0) \in \mathrm{Pos}_d$ (the Pos_d-inclusion holds as $C_{\mu_\star}$ is invertible), cf. (5.91)], we have $\tilde{\rho}_0 := \sup\{s \geq 0 \mid \min_{\{\mathfrak{a}\in\mathbb{R}^{d\times d}\,:\,\|\mathfrak{a}-[\mu_\star]_0\|\leq s\}} \underline{\tilde{\sigma}}(\mathfrak{a}) \geq \sigma_\star/2\} > 0$. This implies, via (5.90), that

$$\sup_{\mathfrak{a}\in\mathfrak{B}_{\tilde{\rho}_0}} \|\mathcal{R}(\mathcal{C}_\mathfrak{a})\| \leq \sqrt{2d/\sigma_\star} =: L_{\mathfrak{p}_\star} \quad \text{for} \quad \mathfrak{B}_{\tilde{\rho}_0} \equiv \mathfrak{B}_{\tilde{\rho}_0}([\mu_\star]_0) := \{\mathfrak{a} \in \mathbb{R}^{d\times d} \mid \|\mathfrak{a} - [\mu_\star]_0\| \leq \tilde{\rho}_0\}$$

(notice that $\{\mathcal{C}_\mathfrak{a} \mid \mathfrak{a} \in \mathfrak{B}_{\tilde{\rho}_0}\} \subset \mathrm{Sym}_d^+$ by definition of $\tilde{\rho}_0$). Hence it is clear that $\mathfrak{K}_3 := \sup\{\xi(\nu) \mid \nu \in \mathfrak{D} \,:\, \|[\nu]_\mathfrak{c} - [\mu_\star]_\mathfrak{c}\|_\mathcal{V} \leq \tilde{\rho}_0\}$ is finite, and combined with (5.78) we find

$$\sup\nolimits_{(\tilde{\mu},\tilde{f})\in\mathbb{B}_r^\star} \|\mathcal{R}(\mathcal{C}_{[\tilde{\mu}_\zeta]_0})\| \;\leq\; L_{\mathfrak{p}_\star} \;\text{ provided that }\; r \;<\; r_0 \wedge r_1 \wedge (\tilde{\rho}_0/\hat{K})^{1/\alpha} =: r_2'. \tag{5.210}$$

Here, the functional assignment [as used in (5.206) and (5.210)]

$$\mathbb{B}_r^\star \ni (\tilde{\mu}, \tilde{f}) \;\longmapsto\; \tilde{\mu}_\zeta \in \tilde{B}_{\hat{K}r^\alpha} \tag{5.211}$$

is the one asserted by Lemma 5.5.14, cf. (5.206). Assuming from now on that

$$r \;<\; r_2 \;:=\; r_2' \wedge (r_*/\hat{K})^{1/\alpha},$$

we just saw that, for each $\nu \in \tilde{B}_{\mathfrak{q}_r}$ with $\mathfrak{q}_r := \hat{K}r^\alpha$, we have both (5.209) and the domination

$$|p(\nu)| \;\leq\; \tilde{\kappa}_0\mathfrak{K}_1(\mathfrak{q}_r)\bigg[\frac{\beta_1([\nu]_0)}{\sqrt{d}}(1 + \tilde{\kappa}_0 + (1 + \mathfrak{K}_3 d)\tilde{\kappa}_0\beta_2([\nu]_0)) + \beta_2([\nu]_0)\mathfrak{K}_2(\mathfrak{q}_r)\bigg]$$

where $\beta_j \equiv \beta_j(\,\cdot\,; A)$ $(j = 1, 2)$ are the bounding functions from Lemma 5.8.8 (5.158)–(5.159). The above is equivalent (cf. the proof of Lemma 5.5.7) to the inequality

$$p(\cdot) \leq \vartheta_r \circ [\,\cdot\,]_0 \quad \text{ for } \quad \vartheta_r := \tilde{\kappa}_0\mathfrak{K}_1(\mathfrak{q}_r)\bigg[\frac{\beta_1}{\sqrt{d}}(1 + \tilde{\kappa}_0 + (1 + \mathfrak{K}_3 d)\tilde{\kappa}_0\beta_2) + \beta_2\mathfrak{K}_2(\mathfrak{q}_r)\bigg]$$

which holds pointwise on $\tilde{B}_{\mathfrak{q}_r}$ and where $[\,\cdot\,]_0 : \mathfrak{D} \to \mathbb{R}^{d\times d}$ is the map that sends a signal to its ground-level coredinate (cf. (5.39)). Note that while β_1 is continuous, the function β_2 is a fraction of the form $\beta_2 = \hat{\beta}/z$, with $\hat{\beta}$ the enumerator and z the denominator in its definition (5.159), where $\hat{\beta}$ and z are each continuous on $\mathbb{R}^{d\times d}$ (Lemma 5.8.8). Note further that by definition of r_2, see (5.210) [and recall (5.90)], the denominator z is bounded away from zero on $[\tilde{B}_{\mathfrak{q}_r}]_0 \equiv \{[\nu]_0 \mid \nu \in \tilde{B}_{\mathfrak{q}_r}\} =: \mathfrak{A}_r$, that is $\mathfrak{z}(r) := \inf\{z(\mathfrak{a}) \mid \mathfrak{a} \in \mathfrak{A}_r\} > 0$. Hence $\vartheta_r|_{\mathfrak{A}_r} \leq \hat{\vartheta}_r|_{\mathfrak{A}_r}$, and for $\hat{\beta}_r := \mathfrak{z}(r)^{-1}\hat{\beta}$ we consequently have the $\tilde{B}_{\mathfrak{q}_r}$-pointwise domination

$$p(\cdot) \leq \hat{\vartheta}_r \circ [\,\cdot\,]_0 \;\text{ for }\; \hat{\vartheta}_r := \tilde{\kappa}_0\mathfrak{K}_1(\mathfrak{q}_r)\bigg[\frac{\beta_1(r)}{\sqrt{d}}(1 + \tilde{\kappa}_0 + (1 + \mathfrak{K}_3 d)\tilde{\kappa}_0\hat{\beta}_r) + \hat{\beta}_r\mathfrak{K}_2(\mathfrak{q}_r)\bigg]. \tag{5.212}$$

The majorant $\hat{\vartheta}_r$ is continuous and strictly positive on the ball $\mathfrak{B}_{\tilde{\rho}_0} \supseteq \mathfrak{A}_r$ defined above.

In particular [recall $\mathfrak{q}_r \leq \tilde{\rho}_0$], for the function $\psi_r(\mathfrak{a}, \mathfrak{a}_1, \ldots, \mathfrak{a}_d) := \hat{\vartheta}_r(\mathfrak{a})$ we have

$$\psi_r \in C(\bar{B}_{\mathfrak{q}_r}^\mathcal{V}; \mathbb{R}_{>0}) \quad \text{on} \quad \bar{B}_{\mathfrak{q}_r}^\mathcal{V} \equiv \bar{B}_{\mathfrak{q}_r}^\mathcal{V}([\mu_\star]_\mathfrak{c}) := \{\tilde{\mathfrak{a}} \in \mathcal{V} \mid \|\tilde{\mathfrak{a}} - [\mu_\star]_\mathfrak{c}\|_\mathcal{V} \leq r\},$$

and by the above $p(\cdot)$-domination (5.212), recalling (5.211) (see also (5.206)), we obtain

$$p(\tilde{\mu}_\zeta) \;\le\; \psi_r([\tilde{\mu}_\zeta]_{\mathfrak{c}}) \quad \text{for each} \;\; (\tilde{\mu}, \tilde{f}) \in \mathbb{B}_r^\star. \tag{5.213}$$

Setting $\psi := \psi_{\tilde{r}_2}$ for $\tilde{r}_2 := 0.95 \cdot r_2$ and introducing the function $\tau : [0, \tilde{r}_2] \to (0, \infty)$ defined by

$$\tau(r) := \min\Big\{Q\big((\gamma k_d \psi(\tilde{\mathfrak{a}}))^{-1}\big) \mid \tilde{\mathfrak{a}} \in \bar{B}^{\mathcal{V}}_{\mathfrak{q}_r}([\mu_\star]_{\mathfrak{c}})\Big\} \quad \text{for} \quad Q(u) := \frac{u}{1+u},$$

we find for any $\mu_\zeta \in \tilde{B}_{\mathfrak{q}_r}$ as in (5.213) [recall (5.211)] that, in consequence of (5.207) and (5.213),

$$\varepsilon(\mu_\zeta) = Q\big(q(\mu_\zeta)\big) \;\ge\; Q\big((\gamma k_d \psi([\mu_\zeta]_{\mathfrak{c}}))^{-1}\big) \;\ge\; \tau(r), \quad \text{for each} \;\; r \in [0, \tilde{r}_2]. \tag{5.214}$$

Now since τ is continuous (Lemma 5.8.9) with $\tau(0) = Q\big((\gamma k_d \psi([\mu_\star]_{\mathfrak{c}}))^{-1}\big) > 0$, there will be a positive radius up to which τ is greater than the continuous majorant $\upsilon : [0, r_0] \ni r \mapsto K' r^{\alpha/2}$ of $\delta_{\perp\!\!\!\perp}(\mu_\zeta)$ (cf. (5.206)). This implies, cf. again (5.206), that we must have

$$\begin{aligned} r_3 &:= \inf\big\{0 \le r \le \tilde{r}_2 \mid \tau(r) - \sup\nolimits_{(\mu,f)\in\mathbb{B}_r^\star} \delta_{\perp\!\!\!\perp}(\mu_\zeta) \le 0\big\} \\ &\ge \inf\big\{0 \le r \le \tilde{r}_2 \mid \tau(r) - \upsilon(r) \le 0\big\} \;>\; 0. \end{aligned} \tag{5.215}$$

Let now $0 \le r \le r_3$, and set $\tilde{K}_1 := 2dk_d \max\{\psi(\tilde{\mathfrak{a}}) \mid \tilde{\mathfrak{a}} \in \bar{B}^{\mathcal{V}}_{\mathfrak{q}_{r_3}}\}$ and fix $K' \equiv K_{r_3}\,(= K'|_{r=r_3})$. Then for any $(\mu, f) \in \mathbb{B}_r(\mu_\star, A)$ with (5.211)-associated signal $\mu_\zeta \in \bar{B}^{\mathcal{V}}_{\mathfrak{q}_{r_3}}$, we have $\delta_{\perp\!\!\!\perp}(\mu_\zeta) \le \tau(r) \le \varepsilon(\mu_\zeta) < 1$ [by (5.215) and (5.214)] and hence, by Theorem 5.6.10 (5.105) [which by (5.209) is fully applicable], obtain that

$$\begin{aligned} &\forall\, \theta_{\mathfrak{p}} \in \Phi(f_*\mu) \;:\; \text{there is} \;\; M = M(\theta_{\mathfrak{p}}) \in \mathrm{M}_d \;\text{ and }\; E = E(\theta_{\mathfrak{p}}) \in \mathbb{R}^{d\times d} \;\;\text{ such that} \\ &\qquad \theta_{\mathfrak{p}} = M(\mathrm{I} + E)A^{-1} \;\;\text{with}\;\; \|E\| \le \epsilon_r \quad \text{for} \;\; \epsilon_r := \tilde{K}_1 \tfrac{\upsilon(r)}{1-\upsilon(r)} \end{aligned} \tag{5.216}$$

(see also (5.202) in the proof of Theorem 5.6.10 for reference). Clearly, the error bound ϵ_r is strictly increasing in r (recall that $\upsilon(r) < 1$) with $\lim_{r\to0+} \epsilon_r = 0$.

Next, denote the constants $K_2 := \kappa_2(A) L_{\mathfrak{p}_\star}$ and $\tilde{\gamma}_\star := \min_{i\in[d]} \big|e_i^\intercal \cdot A^{-1} \mathcal{C}^{1/2}_{[\mu_\star]_0}\big|$ and consider the auxiliary function $\eta : [0, r_3] \to \mathbb{R}_+$ given by

$$\eta(r) := \max\big\{\tilde{\gamma}_\star^{-1} \|A^{-1}(\mathcal{R}(\mathcal{C}_{\mathfrak{a}})^{-1} - \mathcal{R}_\star^{-1})\| \mid (\mathfrak{a}, \mathfrak{a}_1, \cdots, \mathfrak{a}_d) \in \bar{B}^{\mathcal{V}}_{\mathfrak{q}_r}\big\} \tag{5.217}$$

for $\mathcal{R}_\star := \mathcal{R}(\mathcal{C}_{[\mu_\star]_0})$. By (5.157) and Lemma 5.8.9, the function η is continuous with $\eta_i(0) = 0$. Let further $\Lambda_{\mathfrak{a}} := \mathrm{ddiag}_{i=1,\ldots,d}\big(|e_i^\intercal \cdot A^{-1}\mathcal{R}(\mathcal{C}_{\mathfrak{a}})^{-1}|\big)$ and set $\Lambda_{(\mu,f)} := \Lambda_{[\mu_\zeta]_0}$ via (5.211) (cf. (5.206)).

For the (arbitrary) $\varepsilon > 0$ fixed in the beginning of this proof, cf. (5.205), choose the largest

$$0 < r_4 \le r_3 \quad \text{such that} \quad \max\nolimits_{r\in[0,r_4)}\big(\sqrt{2}\eta(r) + (\|A^{-1}\|_2 \hat{\sigma}_r + 1)\epsilon_r\big) < \varepsilon / K_2, \tag{5.218}$$

where $\hat{\sigma}^\star : [0, r_3] \ni r \mapsto \hat{\sigma}^\star_r := \max_{\mathfrak{a}\in\mathfrak{B}_{\mathfrak{q}_r}} \sqrt{\max \lambda(\mathcal{C}_\mathfrak{a})} \in \mathbb{R}_+$ (which is [monotonically increasing and] continuous by (5.156) and Lemma 5.8.9). The proof concludes (recall (5.205)) if we can show

$$\mathfrak{d}(\mathfrak{I}(\mathfrak{p}_\star), \mathfrak{I}(\mathfrak{p})) \leq \varepsilon, \quad \text{for each } \ \mathfrak{p} \in \mathbb{B}_{r_4}(\mathfrak{p}_\star). \tag{5.219}$$

And indeed: Writing $\tilde{\Xi}_\mathcal{R} := \tilde{\Xi}_1 \cdot \mathcal{R}$, we know from Theorem 5.6.10 (applicable by def. of $\mathscr{I}$) that

$$\mathfrak{I}(\mathfrak{p}_\star) = \mathrm{M}_d \cdot A^{-1} \cap \tilde{\Xi}_{\mathcal{R}_\star} = \{P\Lambda_{\mathfrak{p}_\star}^{-1} A^{-1} \mid P \in \mathrm{P}_d^\pm\} =: \mathcal{A}_\star,$$

where the second identity can be easily verified from the above definitions (recalling (5.98) and the unit-row conditions (5.97)). To see that (5.219) holds, take any $\mathfrak{p} \equiv (\tilde{\mu}, \tilde{f}) \in \mathbb{B}_{r_4}(\mathfrak{p}_\star)$ and let us first fix any $\mathfrak{a}_\star \equiv P_{\mathfrak{a}_\star}\Lambda_{\mathfrak{p}_\star}^{-1}A^{-1} \in \mathcal{A}_\star$, denoting $\mathcal{R}_{(\tilde{\mu},\tilde{f})} := \mathcal{R}(\mathcal{C}_{[\tilde{\mu}_\varsigma]_0})$ (cf. (5.211)) for convenience. Since $\mathfrak{I}(\mathfrak{p})$ is non-empty (5.99), there exists some $\theta_\mathfrak{p} \in \tilde{\Xi}_{\mathcal{R}_\mathfrak{p}}$ such that $\theta_\mathfrak{p} = M_\mathfrak{p}(\mathrm{I} + E_\mathfrak{p})A^{-1}$ for $M_\mathfrak{p} \equiv P_\mathfrak{p}\tilde{\Lambda}_{(\mathfrak{p})} \in \mathrm{M}_d$ (with $P_\mathfrak{p} \in \mathrm{P}_d^\pm$ and $\tilde{\Lambda}_{(\mathfrak{p})} \equiv \mathrm{ddiag}(m_\mathfrak{p}^1, \ldots, m_\mathfrak{p}^d) > 0$) and $\|E_\mathfrak{p}\| \leq \epsilon_{r_4}$, see (5.216). By the $\mathrm{P}_d^\pm$-invariance of $\phi_{\tilde{f}_\star\tilde{\mu}}$ and $\tilde{\Xi}_{\mathcal{R}_\mathfrak{p}}$ (cf. (5.96) and (5.111)) we find that also $\theta_\mathfrak{p}^\star := (P_{\mathfrak{a}_\star}P_\mathfrak{p}^{-1})\theta_\mathfrak{p} \in \hat{\mathfrak{I}}(\mathfrak{p})$. Hence (and since $\mathrm{P}_d^\pm \subset \mathrm{O}_d$ and $\|A_1A_2\| \leq \|A_1\|\|A_2\|_2, \forall\, A_1, A_2 \in \mathbb{R}^{d\times d}$)

$$d(\mathfrak{a}_\star, \mathfrak{I}(\mathfrak{p})) \;\leq\; \|\mathfrak{a}_\star - \theta_\mathfrak{p}^\star\| \;\leq\; \|A^{-1}\|_2 \|\Lambda_{\mathfrak{p}_\star}^{-1} - \tilde{\Lambda}_{(\mathfrak{p})}(\mathrm{I} + E_\mathfrak{p})\|. \tag{5.220}$$

Since $\theta_\mathfrak{p}^\star \in \tilde{\Xi}_{\mathcal{R}_\mathfrak{p}}$ and hence $\theta_\mathfrak{p}^\star\mathcal{R}_p^{-1} = P_{\mathfrak{a}_\star}\tilde{\Lambda}_{(\mathfrak{p})}(\mathrm{I} + E_\mathfrak{p})B_\mathfrak{p} \in \tilde{\Xi}_1$ for $B_\mathfrak{p} := (\mathcal{R}_\mathfrak{p}A)^{-1}$, we have

$$1 = |e_i^\intercal \cdot \theta_\mathfrak{p}^\star \mathcal{R}_\mathfrak{p}^{-1}| = m_\mathfrak{p}^{\sigma(i)} \left| e_{\sigma(i)}^\intercal \cdot (B_\mathfrak{p} + E_\mathfrak{p}B_\mathfrak{p}) \right| \qquad \text{for each } \ i \in [d] \tag{5.221}$$

by the unit-rows requirement (5.97); here, σ is the permutation underlying $P_{\mathfrak{a}_\star}$. Consequently,

$$\begin{aligned}
\|\Lambda_{\mathfrak{p}_\star}^{-1} - \tilde{\Lambda}_{(\mathfrak{p})}(\mathrm{I} + E_\mathfrak{p})\| \;&\leq\; \sqrt{\sum_{i=1}^d \left| \frac{1}{|e_i^\intercal \cdot B_{\mathfrak{p}_\star}|} - m_\mathfrak{p}^i \right|^2} + \|A\|_2 L_{\mathfrak{p}_\star}\epsilon_{r_4} \\
&\overset{(5.221)}{\leq} \sqrt{2} \sqrt{\sum_{i=1}^d \left[\frac{m_\mathfrak{p}^i |e_i^\intercal (B_\mathfrak{p} - B_{\mathfrak{p}_\star})|}{|e_i^\intercal \cdot B_{\mathfrak{p}_\star}|} \right]^2 + \left(m_\mathfrak{p}^i |e_i^\intercal E_\mathfrak{p} B_\mathfrak{p}|\right)^2} + \|A\|_2 L_{\mathfrak{p}_\star}\epsilon_{r_4} \\
&\leq\; \frac{\sqrt{2}\|A\|_2 L_{\mathfrak{p}_\star} \|B_\mathfrak{p} - B_{\mathfrak{p}_\star}\|}{\min_{i\in[d]} |e_i^\intercal \cdot B_{\mathfrak{p}_\star}|} + \|A\|_2 \Big(\|A^{-1}\|_2\, \hat{\sigma}^\star_{r_4} + 1 \Big) L_{\mathfrak{p}_\star}\epsilon_{r_4}\,;
\end{aligned}$$

the first inequality is due to $\|\tilde{\Lambda}_{(\mathfrak{p})}E_\mathfrak{p}\| \leq \|\tilde{\Lambda}_{(\mathfrak{p})}\|_2\epsilon_r$ and $\|\tilde{\Lambda}_{(\mathfrak{p})}\|_2 \leq \|A\|_2\|\mathcal{R}_\mathfrak{p}\|$ (recall (5.199)), and for the last inequality we further used $\|B_\mathfrak{p}\|_2 \leq \|A^{-1}\|_2\|\mathcal{R}_\mathfrak{p}^{-1}\|_2 \leq \|A^{-1}\|_2\hat{\sigma}^\star_r$ (cf. (5.90)). Combined with (5.220) and (5.218) (also recalling (5.211)), the above implies that

$$d(\mathfrak{a}_\star, \mathfrak{I}(\mathfrak{p})) \;\leq\; K_2\Big(\sqrt{2}\eta(r_4) + (\|A^{-1}\|_2\, \hat{\sigma}^\star_{r_4} + 1)\epsilon_{r_4}\Big) \;\leq\; \varepsilon$$

and hence, since $\mathfrak{a}_\star \in \mathcal{A}_\star$ was arbitrary, $\sup_{\mathfrak{a}\in\hat{\mathfrak{I}}(\mathfrak{p}_\star)} d(\mathfrak{a}, \mathfrak{I}(\mathfrak{p})) \leq \varepsilon$. That, for any $\mathfrak{p} \in \mathbb{B}_{r_4}(\mathfrak{p}_\star)$, also $\sup_{\mathfrak{b}\in\hat{\mathfrak{I}}(\mathfrak{p})} d(\mathfrak{I}(\mathfrak{p}_\star), \mathfrak{b}) \leq \varepsilon$ follows similarly. Indeed: Take an arbitrary $\mathfrak{b} \in \mathfrak{I}(\mathfrak{p})$, say

$\mathfrak{b} \equiv m_{\mathfrak{b}}(\mathrm{I}+E_{\mathfrak{b}})A^{-1} \in \tilde{\Xi}_{\mathcal{R}_{\mathfrak{p}}}$ (cf. (5.216)) with $\|E_{\mathfrak{b}}\| \leq \epsilon_{r_4}$ and $m_{\mathfrak{b}} \equiv P_\sigma \tilde{\Lambda}_{(\mathfrak{b})} \in \mathrm{M}_d$ for $P_\sigma \in \mathrm{P}_d^{\pm}$ and $\tilde{\Lambda}_{(\mathfrak{b})} \equiv \mathrm{ddiag}[m_{\mathfrak{b}}^1, \ldots, m_{\mathfrak{b}}^d] > 0$. As above, unit-rows again implies (5.221) (but with $(\theta_{\mathfrak{p}}^\star, E_p)$ replaced by $(\mathfrak{b}, E_{\mathfrak{b}})$) so that for $\tilde{\mathfrak{a}} := P_\sigma \Lambda_{\mathfrak{p}_\star}^{-1} A^{-1} \in \mathfrak{I}(\mathfrak{p}_\star)$ we obtain as above that

$$d(\mathfrak{I}(\mathfrak{p}_\star), \mathfrak{b}) \;\leq\; \|\tilde{\mathfrak{a}} - \mathfrak{b}\| \;\leq\; \|A^{-1}\|_2 \|\Lambda_{\mathfrak{p}_\star}^{-1} - \tilde{\Lambda}_{(\mathfrak{b})}(\mathrm{I} + E_{\mathfrak{b}})\| \;\leq\; \varepsilon.$$

Hence $\sup_{\mathfrak{b} \in \mathfrak{I}(\mathfrak{p})} d(\mathfrak{I}(\mathfrak{p}_\star), \mathfrak{b}) \leq \varepsilon$ (as $\mathfrak{b} \in \hat{\mathfrak{I}}(\mathfrak{p})$ was arbitrary), implying (5.219) as desired. □

Remark 5.8.10. (i) For simplicity of constants, we can use (5.155) to majorize (5.217) via

$$\eta(r) \;\leq\; \kappa_2^2(A) L_{\mathfrak{p}_\star} \frac{\|A\|_2}{\tilde{\gamma}_\star \sigma_\star} \sup_{\mathfrak{a} \in \mathfrak{B}_{\mathfrak{q}_r}} \|\mathfrak{a} - [\mu_\star]_0\| \;\leq\; \check{K}_1 r^\alpha$$

for the ($\mathfrak{p}_\star$-dependent) constant $\check{K}_1 := \kappa_2^2(A) L_{\mathfrak{p}_\star} \frac{\|A\|_2}{\tilde{\gamma}_\star \sigma_\star} \hat{K}$; let also $\check{K}_2 := \|A^{-1}\|_2 \hat{\sigma}_{r_3} + 1$. Accordingly, for any given $\varepsilon > 0$ we may then instead of (5.218) choose $\hat{r} = \hat{r}(\varepsilon; \mathfrak{p}_\star)$ as

$$\hat{r} \;:=\; \sup\left\{0 \leq r \leq r_3 \;\middle|\; \max_{s \in [0,r)} (\sqrt{2}\check{K}_1 s^\alpha + \check{K}_2 \epsilon_s) \leq \varepsilon / (\kappa_2(A) L_{\mathfrak{p}_\star})\right\} \;\in (0, r_4]$$

to ensure that the continuity-relation $(5.205)|_{r=\hat{r}}$ holds.

(ii) If for any fixed $A \in \mathrm{GL}_d$ one is interested in the continuity at $\mathscr{I}$ of the map

$$\mathfrak{I}_A \;:\; (\mathfrak{D}, \delta) \ni \mu \;\longmapsto\; \hat{\Phi}(A \cdot \mu) \in (F, \mathfrak{d}), \tag{5.222}$$

i.e. the (partial) continuity of the solution map $\mathfrak{I}(\,\cdot\,, A)$ on the sectionalised identifiability conditions $\mathscr{I}_A := \tilde{\mathscr{I}} \cap (\mathfrak{D} \times \{A\})$ of $E' = \mathfrak{D} \times \mathrm{GL}_d$, then the above proof of Theorem 5.6.14 shows that this continuity holds without imposing any additional integrability assumptions on $\mathfrak{D}$ (cf. (5.50)). More precisely, the continuity of (5.222) admits a simplified quantification via Remark 5.5.15: Operating on the identifiability domain $E_A := \mathfrak{D} \times \{A\}$ allows us to replace the above r_0 by ρ_0/ρ_1 and then copy the argumentation for the proof of Theorem 5.6.14 but from the adapted premise

$$\mu_\chi \equiv A_* \mu \quad \text{with} \quad \delta_\perp(\mu) \leq L_r \sqrt{r} \quad \text{and} \quad \left\|[\mu]_{\mathfrak{c}} - [\mu_\star]_{\mathfrak{c}}\right\|_{\mathcal{V}} \;\leq\; \hat{L} r, \tag{5.206'}$$

which for any $(\mu, A) \in E_A$ with $\delta(\mu_\star, \mu) < r \leq \rho_0/\rho_1$ holds by Remark 5.5.15 (5.82). Working from (5.206'), the above proof of Thm. 5.6.14 goes through verbatim but with

$$\text{the (global}^{40}\text{) replacements: } K' \leftrightarrow L_r,\; \hat{K} \leftrightarrow \hat{L},\; \alpha \leftrightarrow 1,$$

which results in eased moduli of continuity build from (5.53)-independent constants.

[40] So that $r_1 = r_1(\rho_0, \varrho_0, \hat{K}, \alpha)$ is replaced by $r_1(\rho_0, \varrho_0, \hat{L}, 1)$ etc.

5.8.5 Proof of Proposition 5.7.6 (Cases $i = 2, 3$)

Let us now prove the assertion (5.125) for the cases $i = 2, 3$ of multiplicative noise.

Proof of Proposition 5.7.6. We first consider the case $i = 2$, that is the triple $(\tilde{X}^{(2)}, A, S^\eta)$ for $\tilde{X}^{(2)}$ as in (5.122) and S^η as in (5.121).

Then, since Lebesgue-Stieltjes integration obeys the product rule $\int_a^b \varphi \,\mathrm{d}(g_1 g_2) = \int_a^b \varphi g_1 \,\mathrm{d}g_2 + \int_a^b \varphi g_2 \,\mathrm{d}g_1$ for any $\varphi, g_1, g_2 \in \mathcal{C}_1^1$ and every $0 \leq a \leq b \leq 1$, we have

$$\begin{aligned} \langle S^\eta \rangle_{ij} \;&=\; \langle \eta^i, \eta^j \,|\, S^i, S^j \rangle + \langle S^i, \eta^j \,|\, \eta^i, S^j \rangle + \langle \eta^i, S^j \,|\, S^i, \eta^j \rangle + \langle S^i, S^j \,|\, \eta^i, \eta^j \rangle \quad \text{and} \\ \langle S^\eta \rangle_{ijk} \;&=\; \langle \eta^i, \eta^j, \eta^k \,|\, S^i, S^j, S^k \rangle + \langle \eta^i, \eta^j, S^k \,|\, S^i, S^j, \eta^k \rangle + \langle S^i, \eta^j, \eta^k \,|\, \eta^i, S^j, S^k \rangle \\ &\quad + \langle S^i, \eta^j, S^k \,|\, \eta^i, S^j, \eta^k \rangle + \langle \eta^i, S^j, \eta^k \,|\, S^i, \eta^j, S^k \rangle + \langle \eta^i, S^j, S^k \,|\, S^i, \eta^j, \eta^k \rangle \\ &\quad + \langle S^i, S^j, \eta^k \,|\, \eta^i, \eta^j, S^k \rangle + \langle S^i, S^j, S^k \,|\, \eta^i, \eta^j, \eta^k \rangle \end{aligned} \tag{5.223}$$

where for $Y, \tilde{Y}, \hat{Y}, \check{Y}, Z, \tilde{Z} \in \{S, \eta\}$ we denote $\langle Y^i, \tilde{Y}^j \,|\, \hat{Y}^i, \check{Y}^j \rangle := \mathbb{E}\big[\int_0^1\!\int_0^t Y_s^i \tilde{Y}_t^j \,\mathrm{d}\hat{Y}_s^i \mathrm{d}\check{Y}_t^j\big]$ and

$$\langle Y^i, \tilde{Y}^j, Z^k \,|\, \hat{Y}^i, \check{Y}^j, \tilde{Z}^k \rangle := \mathbb{E}\left[\int_0^1\!\!\int_0^t\!\!\int_0^s Y_r^i \tilde{Y}_s^j Z_t^k \,\mathrm{d}\hat{Y}_r^i \mathrm{d}\check{Y}_s^j \mathrm{d}\tilde{Z}_t^k\right]. \tag{5.224}$$

Evaluating the statistics (5.223) sumand-wise in a similar way as before, we find

$$\langle \eta^i, \eta^j \,|\, S^i, S^j \rangle = \int_{\Delta_2} \mathbb{E}[\eta_s^i \eta_t^j] \mathbb{E}[\dot{S}_s^i \dot{S}_t^j] \,\mathrm{d}^2\boldsymbol{t} \quad \text{and} \quad \langle S^i, \eta^j \,|\, \eta^i, S^j \rangle = \int_{\Delta_2} \mathbb{E}[\dot{\eta}_s^i \eta_t^j] \mathbb{E}[S_s^i \dot{S}_t^j] \,\mathrm{d}^2\boldsymbol{t} \quad \text{etc.,}$$

whence and by the fact that S is a mean-stationarity IC-source with $\mathbb{E}[S] = 0$, which implies that $\mathbb{E}[S_s^i \dot{S}_t^j] = \mathbb{E}[S_s^i]\mathbb{E}[\dot{S}_t^j] = 0 = \mathbb{E}[S_s^i S_t^j]$ whenever $i \neq j$, we obtain

$$\langle S^\eta \rangle_{ij} = \left[\int_{\Delta_2} \phi^i_{(S,\eta)}(\boldsymbol{t}) \,\mathrm{d}^2\boldsymbol{t}\right] \cdot \delta_{ij} \qquad \text{for each} \;\; ij \in [d]_2^\star, \tag{5.225}$$

where $\phi^i_{(S,\eta)}(s,t) := \mathbb{E}[\eta_s^i \eta_t^i]\mathbb{E}[\dot{S}_s^i \dot{S}_t^i] + \mathbb{E}[\dot{\eta}_s^i \eta_t^i]\mathbb{E}[S_s^i \dot{S}_t^i] + \mathbb{E}[\eta_s^i \dot{\eta}_t^i]\mathbb{E}[\dot{S}_s^i S_t^i] + \mathbb{E}[\dot{\eta}_s^i \dot{\eta}_t^i]\mathbb{E}[S_s^i S_t^i]$. The third-order statistics (5.224) behave analogously, which leads to

$$\langle S^\eta \rangle_{ijk} = \left[\int_{\Delta_3} \psi^i_{(S,\eta)}(\boldsymbol{t}) \,\mathrm{d}^3\boldsymbol{t}\right] \cdot \delta_{ijk} \qquad \text{for each} \;\; ijk \in [d]_3^\star, \tag{5.226}$$

where this time the integrand is the function

$$\begin{aligned} \psi^i_{(S,\eta)}(r,s,t) := \;& \mathbb{E}[\eta_r^i \eta_s^i \eta_t^i]\mathbb{E}[\dot{S}_r^i \dot{S}_s^i \dot{S}_t^i] + \mathbb{E}[\eta_r^i \eta_s^i \dot{\eta}_t^i]\mathbb{E}[\dot{S}_r^i \dot{S}_s^i S_t^i] + \mathbb{E}[\dot{\eta}_s^i \eta_s^i \eta_t^i]\mathbb{E}[S_r^i \dot{S}_s^i \dot{S}_t^i] \\ & + \mathbb{E}[\dot{\eta}_r^i \eta_s^i \dot{\eta}_t^i]\mathbb{E}[S_r^i \dot{S}_s^i S_t^i] + \mathbb{E}[\eta_r^i \dot{\eta}_s^i \eta_t^i]\mathbb{E}[\dot{S}_r^i S_s^i \dot{S}_t^i] + \mathbb{E}[\eta_r^i \dot{\eta}_s^i \dot{\eta}_t^i]\mathbb{E}[\dot{S}_r^i S_s^i S_t^i] \\ & + \mathbb{E}[\dot{\eta}_r^i \dot{\eta}_s^i \eta_t^i]\mathbb{E}[S_r^i S_s^i \dot{S}_t^i] + \mathbb{E}[\dot{\eta}_r^i \dot{\eta}_s^i \dot{\eta}_t^i]\mathbb{E}[S_r^i S_s^i S_t^i]. \end{aligned}$$

Enumerating the sumands $\phi_j^i \equiv \phi_{\ell,j}^i \cdot \phi_{r,j}^i$ which constitute $\phi^i_{(S,\eta)}$ — resp. the sumands $\psi_j^i \equiv \psi_{\ell,j}^i \cdot \psi_{r,j}^i$ which constitute $\psi^i_{(S,\eta)}$ — in the order in which they are displayed above (so that e.g. $\phi_2^i = \mathbb{E}[\dot\eta_s^i\eta_t^i]\mathbb{E}[S_s^i\dot S_t^i]$ and $(\phi_{\ell,2}^i,\phi_{r,2}^i) = (\mathbb{E}[\dot\eta_s^i\eta_t^i],\mathbb{E}[S_s^i\dot S_t^i])$, etc.), we from (5.225) and by Hölder's inequality obtain that, for each $i\in[d]$,

$$\begin{aligned}|\langle S^\eta\rangle_{ii} - \langle S\rangle_{ii}| \;&=\; \left|\int_{\Delta_2}(\phi_{\ell,1}^i(\boldsymbol{t})-1)\phi_{r,1}^i(\boldsymbol{t}) + \textstyle\sum_{j=2}^4\phi_j^i(\boldsymbol{t})\,\mathrm{d}^2\boldsymbol{t}\right|\\ &\le\; \left[\int_{\Delta_2}(\phi_{\ell,1}^i(\boldsymbol{t})-1)^2\mathrm{d}^2\boldsymbol{t}\right]^{1/2}\cdot\tilde c_{i,1} + \sum_{j=2}^4\left[\int_{\Delta_2}(\phi_{\ell,j}^i(\boldsymbol{t}))^2\mathrm{d}^2\boldsymbol{t}\right]^{1/2}\cdot\tilde c_{i,j}\end{aligned}\tag{5.227}$$

for the source-dependent constants $\tilde c_{i,j} := \big(\int_{\Delta_2}\phi_{r,j}^i(\boldsymbol{t})^2\,\mathrm{d}^2\boldsymbol{t}\big)^{1/2}$. Likewise, for each $i\in[d]$,

$$|\langle S^\eta\rangle_{iii} - \langle S\rangle_{iii}| \;\le\; \left[\int_{\Delta_3}(\psi_{\ell,1}^i(\boldsymbol{t})-1)^2\mathrm{d}^3\boldsymbol{t}\right]^{1/2}\cdot\check c_{i,1} + \sum_{j=2}^8\left[\int_{\Delta_3}(\psi_{\ell,j}^i(\boldsymbol{t}))^2\mathrm{d}^3\boldsymbol{t}\right]^{1/2}\cdot\check c_{i,j}\tag{5.228}$$

for the source-dependent constants $\check c_{i,j} := \big(\int_{\Delta_3}\psi_{r,j}^i(\boldsymbol{t})^2\,\mathrm{d}^3\boldsymbol{t}\big)^{1/2}$. For convenience, let us abbreviate the expression in line (5.227) by $\langle\tilde b_i,\tilde c_i\rangle \equiv \sum_{j=1}^4\tilde b_{i,j}\cdot\tilde c_{i,j}$ for $\tilde b_i := (\tilde b_{i,1},\tilde b_{i,2},\tilde b_{i,3},\tilde b_{i,4})$ and $\tilde c_i := (\tilde c_{i,1},\ldots,\tilde c_{i,4})$ and $\tilde b_{i,j} := \big(\int_{\Delta_2}(\phi_{\ell,j}^i(\boldsymbol{t})-\delta_{1j})^2\,\mathrm{d}^2\boldsymbol{t}\big)^{1/2}$, and likewise abbreviate the right-hand-side in (5.228) by $\langle\check b_i,\check c_i\rangle \equiv \sum_{j=1}^8\check b_{i,j}\cdot\check c_{i,j}$ for $\check b_i := (\check b_{i,1},\ldots,\check b_{i,8})$ and $\check c_i := (\check c_{i,1},\ldots,\check c_{i,8})$ and $\check b_{i,j} := \big(\int_{\Delta_3}(\psi_{\ell,j}^i(\boldsymbol{t})-\delta_{1j})^2\,\mathrm{d}^3\boldsymbol{t}\big)^{1/2}$. Then by (5.225), (5.226), (5.227), (5.228) and recalling (5.57),

$$\begin{aligned}\delta(S,S^\eta)^2 \;&=\; \|N_S^{-1}([S^\eta]_0-[S]_0)N_S^{-1}\|^2 + \sum_{\nu=1}^d\langle S\rangle_{\nu\nu}^{-1}\|N_S^{-1}([S^\eta]_\nu-[S]_\nu)N_S^{-1}\|^2\\ &=\; \sum_{i=1}^d\big(\langle S\rangle_{ii}^{-1}|\langle S^\eta\rangle_{ii}-\langle S\rangle_{ii}|\big)^2 + \sum_{\nu=1}^d\langle S\rangle_{\nu\nu}^{-1}\big(\langle S\rangle_{\nu\nu}^{-1}|\langle S^\eta\rangle_{\nu\nu\nu}-\langle S\rangle_{\nu\nu\nu}|\big)^2\\ &\le\; \sum_{i=1}^d\gamma_i^{-2}\big(\langle\tilde b_i,\tilde c_i\rangle^2 + \gamma_i^{-1}\langle\check b_i,\check c_i\rangle^2\big) =: \beta_2(S,\eta), \qquad \text{for}\quad \gamma_i := \langle S\rangle_{ii}.\end{aligned}$$

Since $\beta(S,\eta)\le\beta_\varepsilon^2$ by assumption on η, we obtain (5.125) for the noise case $i=2$ as desired.

The case $i=3$ follows (almost) analogously: As a consequence of the product rule, the statistics $\langle X^\eta\rangle_{ij}$ and $\langle X^\eta\rangle_{ijk}$ are also of the form (5.223) but with S replaced by X. Hence by the same arguments as above we can conclude that for each $ijk\in[d]_3^\star$,

$$\begin{aligned}\langle X^\eta\rangle_{ij} \;&=\; \langle\eta^i,\eta^j\,|\,X^i,X^j\rangle + \langle X^i,X^j\,|\,\eta^i,\eta^j\rangle \qquad\text{and}\\ \langle X^\eta\rangle_{ijk} \;&=\; \langle\eta^i,\eta^j,\eta^k\,|\,X^i,X^j,X^k\rangle + \langle X^i,\eta^j,X^k\,|\,\eta^i,X^j,\eta^k\rangle + \langle\eta^i,X^j,X^k\,|\,X^i,\eta^j,\eta^k\rangle\\ &\quad + \langle X^i,X^j,\eta^k\,|\,\eta^i,\eta^j,X^k\rangle + \langle X^i,X^j,X^k\,|\,\eta^i,\eta^j,\eta^k\rangle\end{aligned}\tag{5.229}$$

where we used that η is IC, independent from X and mean-stationary. Further, using $\mathbb{E}[\eta]\equiv\mathrm{I}$,

$$\begin{aligned}\|[X^\eta]_0-[X]_0\|^2 \;&=\; \sum_{i=1}^d|\langle\eta^i,\eta^i\,|\,X^i,X^i\rangle - \langle X\rangle_{ii} + \langle X^i,X^i\,|\,\eta^i,\eta^i\rangle|^2 \;=\; \sum_{i=1}^d\left|\int_{\Delta_2}\xi_i(\boldsymbol{t})\,\mathrm{d}^2\boldsymbol{t}\right|^2,\\ \|[X^\eta]_\nu-[X]_\nu\|^2 \;&=\; \tilde\Delta_{\nu\nu\nu}^2 + \textstyle\sum_{i\in[d]\setminus\{\nu\}}\Big(\tilde\Delta_{i\nu\nu}^2+\tilde\Delta_{\nu i\nu}^2\Big) \;=\; \left|\displaystyle\int_{\Delta_3}\Xi_\nu(\boldsymbol{t})\,\mathrm{d}^3\boldsymbol{t}\right|^2 + \displaystyle\sum_{i\neq\nu}(\tilde\Xi_{i,\nu}^{(1)}+\tilde\Xi_{i,\nu}^{(2)})\end{aligned}$$

for $\nu = 1, \dots, d$ and $\tilde{\Delta}_{ijk} := \big|\langle X^\eta\rangle_{ijk} - \langle X\rangle_{ijk}\big|$ and with the (5.229)-derived auxiliary quantities

$$\xi_i(s,t) := (\mathbb{E}[\eta^i_s\eta^i_t] - 1)\mathbb{E}[\dot{X}^i_s\dot{X}^i_t] + \mathbb{E}[\dot{\eta}^i_s\dot{\eta}^i_t]\mathbb{E}[X^i_sX^i_t] \quad \text{and}$$
$$\Xi_\nu(r,s,t) := (\mathbb{E}[\eta^\nu_r\eta^\nu_s\eta^\nu_t] - 1)\mathbb{E}[\dot{X}^\nu_r\dot{X}^\nu_s\dot{X}^\nu_t] + \mathbb{E}[\dot{\eta}^\nu_r\dot{\eta}^\nu_s\dot{\eta}^\nu_t]\mathbb{E}[X^\nu_rX^\nu_sX^\nu_t] + \mathcal{R}_\nu$$

for $R_\nu := \mathbb{E}[\dot{\eta}^\nu_r\eta^\nu_s\dot{\eta}^\nu_t]\mathbb{E}[X^\nu_r\dot{X}^\nu_sX^\nu_t] + \mathbb{E}[\eta^\nu_r\dot{\eta}^\nu_s\dot{\eta}^\nu_t]\mathbb{E}[\dot{X}^\nu_rX^\nu_sX^\nu_t] + \mathbb{E}[\dot{\eta}^\nu_r\dot{\eta}^\nu_s\eta^\nu_t]\mathbb{E}[X^\nu_rX^\nu_s\dot{X}^\nu_t]$ and $\tilde{\Xi}^{(1)}_{i,\nu} := \Big|\int_{\Delta_3}(\mathbb{E}[\eta^\nu_s\eta^\nu_t] - 1)\mathbb{E}[\dot{X}^i_r\dot{X}^\nu_s\dot{X}^\nu_t] + \mathbb{E}[\dot{\eta}^\nu_s\dot{\eta}^\nu_t]\mathbb{E}[\dot{X}^i_rX^\nu_sX^\nu_t]\,\mathrm{d}^3\boldsymbol{t}\Big|^2$, with $\tilde{\Xi}^{(2)}_{i,\nu}$ defined likewise.

Abbreviating $\hat{\xi}_i := \Big|\int_{\Delta_2}\xi_i(\boldsymbol{t})\,\mathrm{d}^2\boldsymbol{t}\Big|^2$ and $\hat{\Xi}_\nu := \Big|\int_{\Delta_3}\Xi_\nu(\boldsymbol{t})\,\mathrm{d}^3\boldsymbol{t}\Big|^2$, we thus have the control

$$\|[X^\eta]_{\mathfrak{c}} - [X]_{\mathfrak{c}}\|^2 \;=\; \sum_{i=1}^d(\hat{\xi}_i + \hat{\Xi}_i)^2 + \sum_{\nu=1}^d\sum_{i\neq\nu}(\tilde{\Xi}^{(1)}_{i,\nu} + \tilde{\Xi}^{(2)}_{i,\nu}) =: \beta_3(X,\eta). \tag{5.230}$$

Now by continuity of the A^{-1}-induced tensor action on $\mathcal{V}$, cf. the proof of Prop. 5.7.4 [first half], there is an explicitly computable threshold $\gamma_\varepsilon \equiv \gamma_\varepsilon(\beta_\varepsilon, A) > 0$ small enough such that

$$\beta_3(X,\eta) \;\leq\; \gamma_\varepsilon \quad \text{implies} \quad \delta(S, A^{-1}X^\eta) \;\leq\; \beta_\varepsilon.$$

This shows (5.125) for the case $i = 3$, as desired. □

5.9 Extensions and Outlook

Let us briefly comment on how the robustness analysis of this chapter may be extended to accommodate the alternative base assumption that the hidden mixing relation between source and observable may be nonlinear. The central update for this would concern the property (5.41) of linear equivariance, which was one of the core properties of the signature on which our blind inversion algorithm was based (Section 5.6.1): Is there a similarly structured way in which also *nonlinear* transformations of a signal are encoded within the signature representation of the signal itself? For polynomial nonlinearities the answer is positive and in fact a direct generalisation of (5.41), as we show in Proposition 5.9.9. This 'nonlinear equivariance' of the signature allows for a straightforward extension of the low-order robustness topology from Section 5.5 to higher-order signature moments, where the order grows with the (polynomial) 'degree of nonlinearity' of the assumed mixing relation between source and observable. Using that the identities (5.70) persist up to any order (at least for stationary IC sources), we may derive appropriately diagonalising tensor contractions (similar to (5.39)) of higher-order signature moments of the source, see Section 5.9.4, that we could then attempt to arrange in such a way as to explicitly characterise (up to monomial ambiguities, cf. Def. 3.23) accurately inverting nonlinear retransformations of the observable by way of the (expected) signature relation (5.245) (in generalisation of its linear reference relation (5.239), for which this blueprint was first conceived and implemented in Section 5.6). This would then result in an explicitly computable procedure for the exact blind inversion of polynomial mixtures that admits explicit robustness guarantees similarly to Theorems 5.6.10 and 5.6.14. Potential applications could include the stability analysis of invariant measures for random dynamical systems that are obtained from polynomial function iterations. Further explorations of this are left for future research.

5.9.1 Signature Representation of Nonlinear Functions on Paths

In this section, we use Lemma 2.2.4 to infer a duality-based embedding into $\mathbb{R}[[d]]$ ('coordinatisation') of polynomial, and potentially more general continuous functions on paths. This embedding will help to locate the information contained in (the signature of) a nonlinearly transformed path within the signature of the (untransformed) path itself, and it provides a new graded discretisation of a given transformation on $\mathbb{R}^d$ which, in particular, gives new and precise ways to study and quantify the action that even extremely subtle nonlinearities have on any given path.

5.9.2 Elementary Signature-Calculus

The nonlinear constitution of the signature vector (2.13) and the coordinate space (2.21) housing it, mandates to complement the basic vector space operations of addition and scalar multiplication by some elementary nonlinear extensions.

Recall to this end that the real vector space $\mathfrak{V} := \mathbb{R}[d]$ of all word-polynomials over the alphabet $[d] = \{1, \ldots, d\}$ is graded (2.22) with the closed[41] subspaces

$$V_k = \left\{ q \in \mathfrak{V} \;\middle|\; q = \sum\nolimits_{|w|=k} c_w \cdot w \right\} \qquad \text{and} \qquad V_{[k]} = \bigoplus_{j=0}^{k} V_j \,,$$

that is, V_k is the set of all homogeneous word-polynomials of length $k \in \mathbb{N}_0$ (with $V_0 = \mathbb{R}$) and $V_{[k]}$ is the set of all word-polynomials of degree less than k, cf. Section 2.3.3. Recall further that $\mathfrak{V}^*$, the algebraic dual of $\mathfrak{V}$, can be identified with $\mathbb{R}[[d]] =: \overline{\mathfrak{V}}$, the space of formal power series over $[d]$, with $\langle \cdot, \cdot \rangle : \mathfrak{V}^* \times \mathfrak{V} \to \mathbb{R}$ their dual pairing (Rem. 2.3.2 (iii)).

For simplicity of notation, we find it convenient to express the nonlinear interrelations between the signature moments (2.18) in the classical language of differential forms.

Notation 5.9.1 (Iterated Integration via Differential Forms)**.** For $U \subseteq \mathbb{R}^n$ open, recall that any absolutely continuous function $h : U \to \mathbb{R}$ induces an (a.e.-defined) differential 1-form

$$\mathrm{d}h := \partial_1 h \cdot \mathrm{d}u_1 \;+\; \cdots \;+\; \partial_n h \cdot \mathrm{d}u_n \;\; \in C_{\mathrm{a.e}}(U; TU),$$

[41] Wrt. the $[d]$-adic topology, cf. Section 2.3.4.

and any collection $h_1, \dots, h_k \in C(U)$ of such functions yields a differential k-form

$$\mathrm{d}h^{1\cdots k} \;:=\; \mathrm{d}h_1 \wedge \cdots \wedge \mathrm{d}h_k = \sum_{i_1,\dots,i_k=1}^{n} \partial_{i_1}h_1 \cdots \partial_{i_k}h_k \cdot \mathrm{d}u_{i_1} \wedge \cdots \wedge \mathrm{d}u_{i_k} \;\in\; C_{\text{a.e.}}(U; \wedge^k T^\star U)$$

where $\bigwedge^\bullet T^\star U \equiv \bigoplus_{p=0}^n \bigwedge^p T^* U$ denotes the exterior algebra over TU with wedge product $\wedge$. By linear extension, we set $\mathrm{d}h^q := \sum_{w\in q} c_w \cdot \mathrm{d}h^w$ for each $q \equiv \sum_{w\in q} c_w \cdot w \in \mathfrak{V}$.[42]

For the particular case $U = (0,1)$ and $x \in C([0,1];\mathbb{R}^d)$ absolutely continuous with (absolutely continuous) components $x^1, \dots, x^d \in \mathcal{C}_1$, we have $\mathrm{d}x^i = \dot{x}^i \cdot \mathrm{d}t$ and hence

$$\langle \mathfrak{sig}(x),\, i_1 \cdots i_k \rangle = \int_{\Delta_k} \mathrm{d}x^{i_1\cdots i_k} = \int_{\Delta_k} \underbrace{\dot{x}^{i_1}_{s_1} \cdots \dot{x}^{i_k}_{s_k}}_{=:\, \dot{x}^{i_1\cdots i_k}} \cdot \mathrm{d}s_1 \wedge \cdots \wedge \mathrm{d}s_k \tag{5.231}$$

for each $i_1 \cdots i_k \in [d]^*_k$, where the integral domain reads $\Delta_k = \Delta^k_{0,1}$ for the simplices

$$\Delta^n_{s,t} \;:=\; \{(u_1, \dots, u_n) \in \mathbb{R}^n \mid s \le u_1 \le u_2 \le \cdots \le u_n \le t\} \tag{5.232}$$

defined for $n \in \mathbb{N}$ and $s, t \in \mathbb{R}$. For $U \subseteq \mathbb{R}^k$ an open superset of $\mathrm{int}(\Delta^k_{0,1})$ and $\eta : U \to \bigwedge^k T^* U$ a measurable k-form, define (on the space of all measurable k-forms, by linear extension) *the (s,t)-iterated integral of η* as

$$\sigma_{s,t}(\eta) \;:=\; \int_{\Delta^k_{s,t}} \eta \qquad\qquad (0 \le s < t \le 1),$$

so that in particular $\mathfrak{sig}_w(x) = \sigma_{0,1}(\mathrm{d}x^w)$ and

$$\mathfrak{sig}_{s,t}(x) := \mathfrak{sig}(x\big|_{|s,t]}) \;=\; \epsilon \;+ \sum\nolimits_{w\in[d]^*} \sigma_{s,t}(\mathrm{d}x^w) \cdot w. \tag{5.233}$$

Finally, for any collection $(A_i)_{i\in I} \subset \mathbb{R}^k$ of measurable sets, I some index set, we write

$$\bigcup_{i\in I} A_i \;=:\; \boxed{\odot}_{i\in I} A_i \quad :\Longleftrightarrow \quad [\forall\, i \ne j \;:\; A_i \cap A_j \quad \text{is a Lebesgue nullset}]$$

if the sets A_i which form the union are disjoint up to a Lebesgue nullset.

The characteristic relations between the signature functionals (2.13) are based on the defining properties of the total order $(\le)$ on $\mathbb{R}$ which in turn manifest in elementary relations between the simplices (5.232). Although the following two results are well-known, we include our own proofs to fix notation.

Lemma 5.9.2 (Product and Iteration of Simplices)**.** *Given $k, \ell \in \mathbb{N}$, split each $u \equiv (u_1, \dots, u_{k+\ell}) \in [0,1]^{k+\ell}$ according to $u = (u^{(1)}, u^{(2)})$ with $u^{(1)} := (u_1, \dots, u_k)$ and $u^{(2)} :=$*

[42] For each $q \in \mathfrak{V}$, we write $u \in q$ if u is a monomial of q (i.e., $u \in q \;:\Leftrightarrow\; [u \in [d]^*$ and $\langle q, u\rangle \ne 0]$).

$(u_{k+1}, \dots, u_{k+\ell})$. *Then*

$$\Delta^{k+\ell}_{s,t} \;=\; \bigsqcup_{u^{(2)}\in\Delta^{\ell}_{s,t}} \Delta^{k}_{s,u_{k+1}} \times \{u^{(2)}\} \qquad \text{and} \tag{5.234}$$

$$\Delta^{k}_{s,t} \times \Delta^{\ell}_{s,t} \;=\; \bigsqcup\!\!\!\!\odot_{\tau\in S_{k,\ell}} \tau(\Delta^{k+\ell}_{s,t}) \qquad (s,t\in[0,1]). \tag{5.235}$$

In particular, for each integrable function $\varphi \equiv \varphi(u^{(1)}, u^{(2)}) : \Delta^{k+\ell}_{s,t} \to \mathbb{R}$ *it holds that*

$$\int_{\Delta^{k+\ell}_{s,t}} \varphi(u)\,\mathrm{d}u \;=\; \int_{\Delta^{\ell}_{s,t}} \int_{\Delta^{k}_{s,u_{k+1}}} \varphi(r, u^{(2)})\,\mathrm{d}r\,\mathrm{d}u^{(2)}. \tag{5.236}$$

Proof. Equation (5.234) is clear by (5.232), and the identity (5.236) follows from (5.234) by way of Fubini's theorem. Indeed, Fubini yields that

$$\begin{aligned}
\int_{\Delta^{k+\ell}_{s,t}} \varphi(u)\,\mathrm{d}u &= \int_{[0,1]^\ell} \left[\int_{[0,1]^k} \mathbb{1}_{\Delta^{k+\ell}_{s,t}}(r, u^{(2)}) \cdot \varphi(r, u^{(2)})\,\mathrm{d}r \right] \mathrm{d}u^{(2)} \\
&= \int_{\Delta^{\ell}_{s,t}} \left[\int_{[0,1]^k} \mathbb{1}_{\Delta^{k+\ell}_{s,t}}(r, u^{(2)}) \cdot \varphi(r, u^{(2)})\,\mathrm{d}r \right] \mathrm{d}u^{(2)} \\
&\overset{(5.234)}{=} \int_{\Delta^{\ell}_{s,t}} \left[\int_{[0,1]^k} \mathbb{1}_{\Delta^{k}_{s,u_{k+1}}}(r) \cdot \varphi(r, u^{(2)})\,\mathrm{d}r \right] \mathrm{d}u^{(2)} \\
&= \int_{\Delta^{\ell}_{s,t}} \int_{\Delta^{k}_{s,u_{k+1}}} \varphi(r, u^{(2)})\,\mathrm{d}r\,\mathrm{d}u^{(2)}
\end{aligned}$$

as claimed. Since $\Delta^{k}_{s,t} \times \Delta^{\ell}_{s,t} = \{(s_1, \dots, s_k, t_1, \dots, t_\ell) \mid (s \le s_1 \le \dots \le s_k \le t) \text{ and } (s \le t_1 \le \dots \le t_\ell \le t)\}$ and $\tau(\Delta^{k+\ell}_{s,t}) = \{(u_1, \dots, u_{k+\ell}) \mid s \le u_{\tau(1)} \le \dots \le u_{\tau(k+\ell)} \le t\}$ for each $\tau \in S_{k+\ell}$, the identity (5.235) is immediate by the fact that the order $(\le)$ on $\mathbb{R}$ is total. ☐

Although the following properties of the signature coordinates are well-known, let us include a proof for completeness.

Proposition 5.9.3 (Calculus of Iterated Path-Integrals)**.** *For* $x = (x^i)_{i\in[d]} \in \mathcal{C}_d$ *a path whose components* $x^i \in C_1$ *are absolutely continuous, let* $\mathrm{d}x^w := \mathrm{d}x^{i_1} \wedge \dots \wedge \mathrm{d}x^{i_k}$ $(w \equiv i_1 \cdots i_k \in [d]^*_k)$ *be the measurable k-form introduced in Remark 5.9.1. Then for any* $s, t \in [0,1]$ *and* $u, v \in [d]^*$ *it holds:*[43]

(i) $$\sigma_{s,t}(\sigma_{s,\cdot}(\mathrm{d}x^u) \cdot \mathrm{d}x^v) \;=\; \sigma_{s,t}(\mathrm{d}x^{u*v})\,;$$

(ii) $$\sigma_{s,t}(\mathrm{d}x^u) \cdot \sigma_{s,t}(\mathrm{d}x^v) \;=\; \sigma_{s,t}(\mathrm{d}x^{u \sqcup\!\sqcup v})\,;$$

(iii) $$\sum_{u' * u'' = u} \sigma_{s,r}(\mathrm{d}x^{u'}) \cdot \sigma_{r,t}(\mathrm{d}x^{u''}) \;=\; \sigma_{s,t}(\mathrm{d}x^u), \qquad \forall\, r \in (s,t).$$

[43] Explaining an abuse of notation, note that the argument differential on the LHS of statement (i) is understood as $\sigma_{s,\cdot}(\mathrm{d}x^u) \cdot \mathrm{d}x^v \,:\, [0,1]^{|v|} \ni r \equiv (r_1, \cdots, r_{|v|}) \mapsto \sigma_{s,r_1}(\mathrm{d}x^u) \cdot \mathrm{d}x^v_r$. Unless mentioned otherwise, this abuse of notation is used throughout.

Proof. (i) : This is an immediate consequence of (5.236). Indeed, using that the absolute continuity of the component paths x^i allows us to rewrite each $\sigma_{s,t}(\mathrm{d}x^w)$ according to (5.231), we obtain

$$\sigma_{s,t}(\sigma_{s,\cdot}(\mathrm{d}x^u)\cdot \mathrm{d}x^v) = \int_{\Delta_{s,t}^{|v|}} \left[\int_{\Delta_{s,r_1}^{|u|}} \dot{x}_q^u \,\mathrm{d}q\right]\cdot \dot{x}_r^v \,\mathrm{d}r \overset{(5.236)}{=} \int_{\Delta_{s,t}^{|u*v|}} \dot{x}_{(q,r)}^{u*v}\,\mathrm{d}(q,r) = \sigma_{s,t}(\mathrm{d}x^{u*v})$$

as claimed (where the outer integration is over $r \equiv (r_1, \ldots, r_{|v|})$).

(ii) : The second statement follows directly from (5.235). Indeed, for any words $u \equiv i_1 \cdots i_{|u|}$ and $v \equiv i_{|u|+1} \cdots i_{|u|+|v|}$ with $|u| =: k$ and $|v| =: \ell$, it holds

$$\begin{aligned}
\sigma_{s,t}(\mathrm{d}x^u)\cdot\sigma_{s,t}(\mathrm{d}x^v) &= \int_{\Delta_{s,t}^k\times\Delta_{s,t}^\ell} \dot{x}^u_{(s_1,\cdots,s_k)}\cdot \dot{x}^v_{(t_1,\cdots,t_\ell)}\,\mathrm{d}s_1\mathrm{d}s_2\cdots\mathrm{d}s_k\mathrm{d}t_1\cdots\mathrm{d}t_\ell \\
&\overset{(5.235)}{=} \sum_{\tau\in S_{k,\ell}} \int_{\Delta_{s,t}^{k+\ell}} \dot{x}^{i_{\tau(1)}\cdots i_{\tau(k)}}_{(r_{\tau(1)},\cdots,r_{\tau(k)})}\cdot \dot{x}^{i_{\tau(k+1)}\cdots i_{\tau(k+\ell)}}_{(r_{\tau(k+1)},\cdots,r_{\tau(k+\ell)})}\mathrm{d}r_{\tau(1)}\cdots\mathrm{d}r_{\tau(k+\ell)} \\
&= \sum_{\tau\in S_{k,\ell}} \int_{\Delta_{s,t}^{k+\ell}} \mathrm{d}x^{i_{\tau(1)}\cdots i_{\tau(k)} * i_{\tau(k+1)}\cdots i_{\tau(k+\ell)}} \\
&= \sigma_{s,t}\big(\textstyle\sum_{\tau\in S_{k,\ell}} \mathrm{d}x^{i_{\tau(1)}\cdots i_{\tau(k)} * i_{\tau(k+1)}\cdots i_{\tau(k+\ell)}}\big) = \sigma_{s,t}(\mathrm{d}x^{u\sqcup\!\sqcup v})
\end{aligned}$$

as desired. (The second of the above equations is due to (5.235) combined with Fubini-Tonelli and the fact that $\mathbb{1}_{\tau(\Delta^m_{s,t})}(s_1, \ldots, s_m) = \prod_{\nu=1}^m \mathbb{1}_{[s,s_{\tau(\nu+1)}]}(s_{\tau(\nu)})$ with $\tau(m + 1) := t$.)

(iii) : The third statement is a direct consequence of Chen's identity. Indeed, since $\sigma_{s,t}(\mathrm{d}x^u) = \langle\mathfrak{sig}(x|_{[s,t]}),\, u\rangle$ for each $u \in [d]^*$, the fact that the path $x|_{[s,t]}$ is a concatenation of $x|_{[s,r]}$ and $x|_{[r,t]}$, any $s < r < t$, implies

$$\sigma_{s,t}(\mathrm{d}x^u) = \langle\mathfrak{sig}(x|_{[s,r]}) * \mathfrak{sig}(x|_{[r,t]}),\, u\rangle = \sum_{u'*u''=u} \langle q_1, u'\rangle\cdot\langle q_2, u''\rangle$$

for $q_1 := \mathfrak{sig}_{s,r}(x)$ and $q_2 := \mathfrak{sig}_{r,t}(x)$, as stated. □

5.9.3 The Co-Signature of Functions on Paths

Building on ideas and results that were first published in [36], we use Proposition 5.9.3 to introduce a signature-based 'dual coordinatisation' of continuous transformations of $\mathcal{BV}$.

We define this coordinatisation for polynomial transforms first, and then widen it to the space of continuous transformations on $\mathcal{BV}$ by continuous extension via Stone-Weierstrass.

Set $\mathcal{R} := \mathbb{R}[u_1, \ldots, u_d]$ for the ring of real polynomials in d variables, and let

$$\mathfrak{P} := \mathcal{R}\times\cdots\times\mathcal{R} \qquad (d\text{-fold direct product})$$

be the set of all polynomial transforms over $\mathbb{R}^d$. Notice that $\mathfrak{P}$ is a free $\mathcal{R}$-module of rank d and a monoid wrt. the composition $\circ\, :\, (p, q) \mapsto p \circ q := (p_1 \circ q, \ldots, p_d \circ q)$.

Remark 5.9.4. Each $g = (g_1, \ldots, g_d) \in \mathfrak{P}$ can be identified with the transformation $g \cong \sum_{i=1}^d g_i \cdot e_i : \mathbb{R}^d \to \mathbb{R}^d$, for $(e_i)_{i\in[d]}$ the standard basis of $\mathbb{R}^d$, and each component g_i of g is of the form

$$g_i \;=\; \textstyle\sum_{\alpha\in[\bar{m}_i]_d^*} c_{\alpha|i} \cdot u^\alpha \;\in\; \mathcal{R}$$

for $[\bar{m}_i]_d^* := (\{0, 1, \ldots, m_i\}, *)$, $m_i \in \mathbb{N}$, and coefficients $(c_{\alpha|i})_\alpha \subset \mathbb{R}$; the number $n_i := \max\{|\alpha| := \alpha_1 + \ldots + \alpha_d \mid \alpha \in [\bar{m}_i]_d^* \text{ s.t. } c_{\alpha|i} \neq 0\}$ is called the *degree* of g_i, and $\deg(g) := \max_{i\in[d]} n_i$ is called the *(total) degree of* g.

Notation 5.9.5. Let $\iota : (\mathcal{R}, +, \cdot) \hookrightarrow (\mathfrak{V}, +, ⧢)$ be the unital ring-monomorphism defined by $\iota(u_i) := i$, and notice that the group $(\mathfrak{V}, +)$ becomes a (right- and left-) $\mathcal{R}$-module wrt.

$$\text{the scalar multiplication} \quad \odot : \mathfrak{V} \times \mathcal{R} \to \mathfrak{V},\; (q, r) \mapsto q \odot r := q ⧢ \iota(r).$$

The first observation of this section states that the signature transform is *polynomially equivariant* in a sense made precise below. To this end, note that for $p \in \mathcal{R}$ a polynomial, $x \equiv (x^i) \in C_d$ absolutely continuous and $\mathfrak{q} \in \mathbb{R}^{[0,1]}$ integrable, it holds a.e. on $\mathbb{I}$ that

$$\mathrm{d}p(x) \;=\; \sum_{\nu=1}^d \partial_\nu p(x) \cdot \dot{x}^\nu \mathrm{d}t \quad \text{and hence} \quad \int \mathfrak{q}\, \mathrm{d}p(x) \;=\; \sum_{\nu=1}^d \int \mathfrak{q} \cdot \partial_\nu p(x)\, \mathrm{d}x^\nu$$

by the chain rule. This motivates the following concept.[44]

Definition 5.9.6 (Co-Signature)**.** For polynomials $\mathfrak{q} \in \mathfrak{V}$ and $p, q \in \mathcal{R}$, define as

$$\fint \mathfrak{q}\, \mathrm{d}p \;:=\; \sum_{\nu=1}^d \left[\mathfrak{q} \odot \partial_\nu p\right] * \nu \quad \text{and} \quad \fint\!\!\fint \mathfrak{q}\, \mathrm{d}p\, \mathrm{d}q := \fint \left[\fint \mathfrak{q}\, \mathrm{d}p\right] \mathrm{d}q \tag{5.237}$$

the (iterated) *polynomial integral of* $\mathfrak{q}$ *along* p (and q). For $g = (g_i)_{i\in[d]} \in \mathfrak{P}$, we set

$$\mathsf{S}(g) \;:=\; \langle \cdot, \epsilon \rangle \cdot \epsilon \;+ \sum_{w\in[d]^*\setminus\{\epsilon\}} \underbrace{\left\langle \,\cdot\,, \fint \cdots \fint\!\!\fint \mathrm{d}g_{i_1}\, \mathrm{d}g_{i_2} \cdots \mathrm{d}g_{i_k} \right\rangle}_{=:\; \mathsf{S}(g)_w \quad (w\,\equiv\, i_1\cdots i_k)} \cdot\, i_1 \cdots i_k \;:\; \overline{\mathfrak{V}} \;\to\; \overline{\mathfrak{V}}$$

and call this operator the *co-signature* of g.

Example 5.9.7 (Linear Equivariance of the Signature)**.** As an illustration, consider the linear polynomial transformation $\boldsymbol{A} \equiv A \cdot x \in \mathfrak{P}$ given by $\boldsymbol{A}(u) = (\sum_{j=1}^d a_{ij} \cdot u_j)_{i\in[d]}$ for $A \equiv (a_{ij}) \in \mathrm{GL}_d(\mathbb{R})$. Then, for $\mathfrak{q} \equiv (\mathfrak{q}_w)_{w\in[d]^*} \in \overline{\mathfrak{V}}$ and $\mathsf{S}(\boldsymbol{A}) \equiv (\mathsf{S}(\boldsymbol{A})_w)_{w\in[d]^*}$ as above,

$$\mathsf{S}(\boldsymbol{A})_w(\mathfrak{q}) = \sum_{\nu_1,\ldots,\nu_k=1}^d a_{i_1\nu_1} \cdots a_{i_k\nu_k} \cdot \mathfrak{q}_{\nu_1\cdots\nu_k} \cdot i_1 * \cdots * i_k \qquad (w \equiv i_1 \cdots i_k). \tag{5.238}$$

[44] In the following, we may think of the symbol $\fint$ as a 'discrete integral'.

Notice that $\mathsf{S}(\boldsymbol{A})$ generalises the transformation matrix A of the linear map $\boldsymbol{A}$. Indeed: for words $j \in [d]^*$ of length one we obtain that $\mathsf{S}(\boldsymbol{A})(j)$ encodes the j^{th}-column of A,

$$\mathsf{S}(\boldsymbol{A})(j) = \sum\nolimits_{w \in [d]^*} \mathsf{S}(\boldsymbol{A})_w(j) \cdot w \overset{(5.238)}{=} \sum_{i=1}^{d} a_{ij} \cdot i \;\cong\; (a_{1j}, \cdots, a_{dj})^{\intercal},$$

whence for any vector $v \equiv \sum_{i=1}^{d} v_i \cdot i \in V_1$ we have

$$\mathsf{S}(\boldsymbol{A})(v) = \sum_{j=1}^{d} v_j \cdot \mathsf{S}(\boldsymbol{A})(j) = \sum_{i=1}^{d} (\textstyle\sum_{j=1}^{d} a_{ij} v_j) \cdot i \;\cong\; A \cdot v.$$

Analogously one can read off that for any $v_1, \ldots, v_k \in V_1$, $k \in \mathbb{N}$,

$$\mathsf{S}(\boldsymbol{A})(v_1 * \cdots * v_k) = (A \cdot v_1) \otimes \cdots \otimes (A \cdot v_k) = A^{\otimes k} \cdot v_1 \otimes \cdots \otimes v_k,$$

whence the affine equivariance of $\mathfrak{sig}$ (Lemma 5.4.7) can be succinctly expressed as

$$\mathfrak{sig}(A \cdot x) = \epsilon + \sum_{k=1}^{\infty} A^{\otimes k} \cdot \mathfrak{sig}_k(x) = \mathsf{S}(\boldsymbol{A}) \cdot \mathfrak{sig}(x) \tag{5.239}$$

for each $x \in \mathcal{BV}$. As we will see next, the above factorisation property (5.239) of the signature generalises from linear to polynomial transformations.

The 'discrete integral' (5.237) of an element $\mathfrak{q} \in \mathfrak{V}$ has the following basic properties.

Lemma 5.9.8. *The operator* $\mathfrak{V} \times \mathcal{R} \ni (\mathfrak{q}, p) \mapsto \displaystyle\nmid\!\!\!\int \mathfrak{q}\, \mathrm{d}p \in \mathfrak{V}$ *is* $\mathbb{R}$*-bilinear, and it holds*

$$\nmid\!\!\!\int \mathrm{d}p \;=\; \iota(p) \tag{5.240}$$

for each $p \in \mathcal{R}$ *with* $p(0) = 0$.

Proof. The bilinearity is clear from Definition 5.9.6. For the remaining assertion, let $p \in \mathcal{R}$ be arbitrary. By the asserted bilinearity together with the $\mathbb{R}$-linearity of $\iota : \mathcal{R} \to V$, it is no loss of generality to assume that p is a monomial, i.e. that $p = u^{\alpha}$ for some $\alpha = (\alpha_\nu) \in (\mathbb{N}_0)^{\times d}$. Note that then,

$$\partial_\nu p = \alpha_\nu \cdot u^{\alpha - e_\nu} \quad \text{and hence} \quad \sum_{\nu=1}^{d} \partial_\nu p \cdot u_\nu = |\alpha| \cdot p \tag{5.241}$$

for the exponent length $|\alpha| = \alpha_1 + \ldots + \alpha_d$. In this case, (ii) follows by induction on $|\alpha|$. Indeed, with the base case $|\alpha| = 1$ being clear (recall that ι is unital), suppose that (ii) holds if $|\alpha| = k$ ($k \in \mathbb{N}$ fixed). Thus if $|\alpha| = k+1$, the identity (5.241) implies

$$\begin{aligned}
\iota(p) &= \frac{1}{|\alpha|} \sum_{\nu=1}^{d} \iota(\partial_\nu p) \sqcup\!\sqcup\, \nu \overset{\text{I.H.}}{=} \frac{1}{|\alpha|} \sum_{\nu,\mu=1}^{d} \left[\iota(\partial_\mu \partial_\nu p) * \mu\right] \sqcup\!\sqcup\, \nu \\
&= \frac{1}{|\alpha|} \sum_{\nu,\mu=1}^{d} \left[(\iota(\partial_\nu \partial_\mu p) \sqcup\!\sqcup\, \nu) * \mu + (\iota(\partial_\mu \partial_\nu p) * \mu) * \nu\right] \\
&= \frac{1}{|\alpha|} \left[(|\alpha| - 1) \sum_{\mu=1}^{d} \iota(\partial_\mu p) * \mu + \sum_{\nu=1}^{d} \iota(\partial_\nu p) * \nu\right] = \sum_{\nu=1}^{d} \iota(\partial_\nu p) * \nu = \nmid\!\!\!\int \mathrm{d}p
\end{aligned}$$

as desired, where the third of the above identities holds by definition of the shuffle product, while the fourth identity uses the decomposition (5.241) and the induction hypothesis. □

An equivalent version of the following observation was first published as [36, Thm 2].

Proposition 5.9.9 (Polynomial Equivariance). *For every $g \in \mathfrak{P}$ with $g(0) = 0$, it holds*

$$\mathfrak{sig}(g(x)) = \mathsf{S}(g) \cdot \mathfrak{sig}(x) \tag{5.242}$$

for each $x \in \mathcal{C}_d$ absolutely continuous with $x_0 = 0$.

Proof. Let $g = (g_i) \in \mathfrak{P}$ and $x = (x^i) \in \mathcal{C}_d$ absolutely continuous with derivative $\dot{x}$ (defined a.e.). We prove the stronger assertion

$$\langle \mathfrak{sig}_{0,t}(g(x)),\, w\rangle = \Big\langle \mathfrak{sig}_{0,t}(x)\,,\varrho_g(w)\Big\rangle \quad \text{with} \quad \varrho_g(w) := \not\!\!\int \cdots \not\!\!\int\!\!\!\not\!\!\int \mathrm{d}g_{i_1}\,\mathrm{d}g_{i_2}\,\cdots\,\mathrm{d}g_{i_k} \tag{5.243}$$

for each $w = i_1 \cdots i_k \in [d]^*_k$, $k \in \mathbb{N}$, and $t \in [0,1]$. We proceed by induction.

For $k = 1$, take any $i \in [d]$ and notice that then, for each $t \in [0,1]$,

$$\langle \mathfrak{sig}_{0,t}(g(x)),\, i\rangle = (g(x_t))_i = g_i(x_t^1, \cdots, x_t^d) = \langle \mathfrak{sig}_{0,t}(x),\, \iota(g_i)\rangle = \langle \mathfrak{sig}_{0,t}(x),\, \varrho_g(i)\rangle,$$

where the penultimate equation is due to Proposition 5.9.3 (ii), i.e. the well-known fact that the map $(\mathfrak{V}, ⧢) \ni u \mapsto \langle \mathfrak{sig}_{0,t}(x),\, u\rangle \in (\mathbb{R}, \cdot)$ is a homomorphism, and the last identity holds by Lemma 5.9.8, eq. (5.241).

Abbreviating $\mathfrak{q}_t := \mathfrak{sig}_{0,t}(x)$, let us next suppose that (5.243) holds for some fixed $k \in \mathbb{N}$. Let $w \in [d]^*_{k+1}$ be arbitrary, whence $w = w' * i$ for some $w' \equiv i_1 \cdots i_k \in [d]^*_k$ and $i \in [d]$. By induction hypothesis, it holds for all $s \in [0,1]$ that

$$\varphi(s) := \sigma_{0,s}(\mathrm{d}g(x)^{w'}) = \sigma_{0,s}(\mathrm{d}x^{\varrho_g(w')}) \tag{5.244}$$

for the differential forms $\mathrm{d}g(x)^{w'} := \mathrm{d}g_{i_1}(x) \wedge \cdots \wedge \mathrm{d}g_{i_k}(x) \in C_{\text{a.e.}}(U_k; \bigwedge^k T^*U_k)$ and $\mathrm{d}x^u \in C_{\text{a.e.}}(U_{|u|}; \bigwedge^{|u|} T^*U_{|u|})$ for $U_\ell := \operatorname{int}(\Delta^\ell_{0,1})$, cf. Notation 5.9.1. Now by definition and Proposition 5.9.3 (i),

$$\langle \mathfrak{sig}_{0,t}(g(x)),\, w\rangle = \sigma_{0,t}(\mathrm{d}g(x)^w) = \sigma_{0,t}(\sigma_{0,\cdot}(\mathrm{d}g(x)^{w'}) \cdot \mathrm{d}g(x)^i) = \sigma_{0,t}(\varphi(\cdot) \cdot \mathrm{d}g_i(x)),$$

where, by the chain rule, the differential $\mathrm{d}g_i(x)$ can be written as

$$\mathrm{d}g_i(x) \;=\; \sum_{\nu=1}^{d} \partial_\nu g_i(x) \cdot \dot{x}^\nu \mathrm{d}s \;=\; \sum_{\nu=1}^{d} \sigma_{0,\cdot}(\mathrm{d}x^{\iota(\partial_\nu g_i)}) \cdot \mathrm{d}x^\nu$$

(the identity $\partial_\nu g_i(x_s) = \sigma_{0,s}(\mathrm{d}x^{\iota(\partial_\nu g_i)})$ is again due to the fact that $\partial_\nu g_i$ is a polynomial and $\mathfrak{V} \ni q \mapsto \sigma_{0,s}(\mathrm{d}x^q) \in \mathbb{R}$ is a shuffle-homomorphism). We thus have

$$\begin{aligned}\langle \mathfrak{sig}_{0,t}(g(x)),\, w\rangle &= \sum_{\nu=1}^{d} \sigma_{0,t}(\varphi(\cdot)\cdot\sigma_{0,\cdot}(\mathrm{d}x^{\iota(\partial_\nu g_i)})\cdot \mathrm{d}x^\nu) = \sum_{\nu=1}^{d}\sigma_{0,t}\Big(\sigma_{0,\cdot}\Big(\mathrm{d}x^{\varrho_g(w')\,\sqcup\!\sqcup\,\iota(\partial_\nu g_i)}\Big)\cdot \mathrm{d}x^\nu\Big)\\ &= \sigma_{0,t}\Big(\mathrm{d}x^{\sum_{\nu=1}^{d}[\varrho_g(w')\,\sqcup\!\sqcup\,\iota(\partial_\nu g_i)]*\nu}\Big) = \Big\langle \mathfrak{sig}_{0,t}(x),\ \textstyle\oint \varrho_g(w')\,\mathrm{d}g_i\Big\rangle\\ &= \Big\langle \mathfrak{sig}_{0,t}(x),\ \varrho_g(w' * i)\Big\rangle\end{aligned}$$

as desired. Note that the second and third of the above identities hold by Proposition 5.9.3 (ii) (via (5.244)) and (i), respectively, and for the penultimate identity recall (5.233). □

For $\varrho_g(\cdot)$ as in (5.243) we have $\mathsf{S}(g) = \sum_{w\in[d]^*}\langle\cdot, \varrho_g(w)\rangle \cdot w$, so that coefficient-wise the identity (5.245) reads

$$\langle \mathfrak{sig}(g(x)), w\rangle \;=\; \langle \mathfrak{sig}(x), \varrho_g(w)\rangle, \qquad \text{for each } w \in [d]^* \tag{5.245}$$

and for each absolutely continuous $x \in \mathcal{BV}$.

5.9.4 Higher-Order Coredinates

We have limited our exposition to the (sufficient) case of second- and third-order signature moments (5.39) for simplicity, but it is of course also possible to aim for the joint diagonalizastion of (contracted) signature moments of order four or higher. As an example, let us derive a corresponding set of signature-based matrix statistics of order four.

Note that for any $B \equiv (b_{ij}) \in \mathbb{R}^{d\times d}$ and each $i, j, k \in [d]$ we by (5.41) have

$$\langle B\cdot\mu\rangle_{kijk} \;=\; \sum_{\alpha,\beta,\gamma,\nu=1}^{d} b_{k\alpha}b_{i\beta}b_{j\gamma}b_{k\nu}\cdot\langle\mu\rangle_{\alpha\beta\gamma\nu} \;=\; \sum_{\beta,\gamma=1}^{d} b_{i\beta}b_{j\gamma}\sum_{\alpha,\nu=1}^{d} b_{k\alpha}b_{k\nu}\cdot\langle\mu\rangle_{\alpha\beta\gamma\nu}.$$

Now for $(\beta,\gamma)\in[d]^2$ fixed define the matrix $\langle\mu\rangle^{\beta\gamma} := \big(\langle\mu\rangle_{\alpha\beta\gamma\nu}\big)_{\alpha\nu}$ and note that

$$\sum_{\alpha,\nu=1}^{d} b_{k\alpha}b_{k\nu}\cdot\langle\mu\rangle_{\alpha\beta\gamma\nu} = \Big[B\cdot\langle\mu\rangle^{\beta\gamma}\cdot B^{\intercal}\Big]_{kk}.$$

Consequently, we obtain for the contracted tensor $C_\mu^{ij} := \sum_{k=1}^{d}\langle\mu\rangle_{kijk}$ that

$$C_\mu^{ij} = \sum_{\beta,\gamma=1}^{d}\sum_{k=1}^{d}\Big[B\cdot\langle\mu\rangle^{\beta\gamma}\cdot B^{\intercal}\Big]_{kk}\cdot b_{i\beta}b_{j\gamma} = \sum_{\beta,\gamma=1}^{d} b_{i\beta}b_{j\gamma}\mathrm{Tr}\left(B^{\intercal}\langle\mu\rangle^{\beta\gamma}B\right) \tag{5.246}$$

where $\mathrm{Tr}(\cdot)$ denotes the trace of a $d\times d$-matrix. Now if A is orthogonal, as may always be assumed in classical (i.e. unperturbed) ICA applications, then

$$\mathrm{Tr}\left(B^{\intercal}\langle\mu\rangle^{\beta\gamma}B\right) = \mathrm{Tr}(\langle\mu\rangle^{\beta\gamma}) =: C^{(1)}_{\beta\gamma}(\mu).$$

Hence for $C^{(1)}(\mu) := (C^{(1)}_{\beta\gamma})$ we have with (5.246) that

$$C^{ij}_{\mu} = \Big[B \cdot C^{(1)}(\mu) \cdot B^{\intercal}\Big]_{ij} \quad \text{for all } i,j \in [d], \quad \text{hence} \quad (C^{ij}_{\mu}) = B \cdot C^{(1)}(\mu) \cdot B^{\intercal}.$$

Denoting $A := B^{-1}$ we found $C^{(1)}(\mu) = A \cdot (C^{ij}_{\mu}) \cdot A^{\intercal}$, hence for $\chi := A \cdot \zeta$, some $\zeta \in \mathcal{D}$,

$$\Big(\mathrm{Tr}\left(\langle\chi\rangle^{\beta\gamma}\right)\Big)_{\beta\gamma} = A \cdot \Big[(\langle\zeta\rangle_{1ij1}) + \cdots + (\langle\zeta\rangle_{dijd})\Big] \cdot A^{\intercal}.$$

Note that if the fourth-order tensor $(\langle\zeta\rangle_{ijkl})$ is of the form (5.70), then the matrix $(C^{ij}_{\zeta}) = \big(\sum_{k=1}^{d}\langle\zeta\rangle_{kijk}\big)$ is diagonal as desired. Likewise, we obtain the equivariant statistics

$$\begin{aligned}
\left(\mathrm{Tr}\left(\langle\chi\rangle^{\gamma\nu}\right)\right) &= A \cdot \Big[(\langle\zeta\rangle_{11ij}) + \cdots + (\langle\zeta\rangle_{ddij})\Big] \cdot A^{\intercal}, \\
\Big(\mathrm{Tr}\left(\langle\chi\rangle^{\beta\nu}\right)\Big) &= A \cdot \Big[(\langle\zeta\rangle_{1i1j}) + \cdots + (\langle\zeta\rangle_{didj})\Big] \cdot A^{\intercal}
\end{aligned}$$

and so on.

Bibliography

[1] B. Afsari. Sensitivity Analysis for the Problem of Matrix Joint Diagonalization. *SIAM J. Matrix Anal. Appl.*, **30**.3:1148–1171, 2008.

[2] Y. Aït-Sahalia, J. Fan, and D. Xiu. High-Frequency Covariance Estimates with Noisy and Asynchronous Financial Data. *J. Amer. Statist. Assoc.*, **105**.492:1504–1517, 2010.

[3] T. Akutsu et al. Application of independent component analysis to the iKAGRA data. *Progress of Theoretical and Experimental Physics*, **2020**.5 (2020): 053F01.

[4] L. B. Almeida. MISEP – Linear and Nonlinear ICA Based on Mutual Information. *J. Mach. Learn. Res.*, **4**:1297–1318, 2003.

[5] L. Ardizzone, J. Kruse, S. Wirkert, D. Rahner, E. W. Pellegrini, R. S. Klessen, L. Maier-Hein, C. Rother, and U. Köthe. Analyzing Inverse Problems With Invertible Neural Networks. *ICLR 2019.*

[6] W. Arendt and K. Urban. *Partielle Differenzialgleichungen.* Springer, 2018.

[7] A.V. Arkhangel'skii and V.V. Fedorchuk. *General Topology I: Basic Concepts and Constructions Dimension Theory.* Encyclopaedia of Mathematical Sciences, Vol. **17**, Springer, 1990.

[8] R. Aziz, C. Verma, and N. Srivastava. A fuzzy based feature selection from independent component subspace for machine learning classification of microarray data. *Genomics Data*, **8**:4–15, 2016.

[9] F. R. Bach and M. I. Jordan. Kernel Independent Component Analysis. *J. Mach. Learn. Res.*, **3**:1–48, 2002.

[10] A. Back and A. Weigend. A first application of independent component analysis to extracting structure from stock returns. *International Journal of Neural Systems*, **8**.4:473–484, 1997.

[11] M. Behr, C. Holmes, and A. Munk. Multiscale Blind Source Separation. *Ann. of Statist.*, **46**.2:711–744, 2018.

[12] M. Belkin, L. Rademacher, and J. Voss. Blind signal separation in the presence of Gaussian noise. *PMLR, Conference on Learning Theory*, pages 270–287, 2013.

[13] A. J. Bell and T. Sejnowski. An information maximisation approach to blind separation and blind deconvolution. *Neural Computation*, **7**.6:1129–1159, 1995.

[14] A. Belouchrani, K. A. Meraim, J. F. Cardoso, and E. Moulines. A blind source separation technique based on second order statistics. *IEEE Trans. Signal Process.*, **45**.2:434–444, 1997.

[15] J. Berner, D. Elbrächter, and P. Grohs. How degenerate is the parametrization of neural networks with the ReLU activation function? *NeurIPS 2019.*

[16] P. Billingsley. *Convergence of Probability Measures.* Second Edition, John Wiley & Sons, 1999.

[17] V. Bogachev. Average Approximations and Moments of Measures. *J. Complexity*, **16**:390–410, 2000.

[18] V. I. Bogachev. *Measure Theory.* Vol. I, II. Springer, 2007.

[19] P. Bonnier and H. Oberhauser. Signature Cumulants, Ordered Partitions, and Independence of Stochastic Processes. *Bernoulli*, **26**.4:2727–2757, 2020.

[20] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence.* Oxford University Press, 2013.

[21] R.C. Bradley. Basic Properties of Strong Mixing Conditions. A Survey and Some Open Questions. *Probability Surveys*, **2**:107–144, 2005.

[22] P. Brakel and Y. Bengio. Learning independent features with adversarial nets for non-linear ICA. *Preprint*, page `arXiv:1710.05050[stat.ML]`, 2017.

[23] J. Bruno and P. Szeptycki. Quantales, Generalised Premetrics and Free Locales. *Appl. Categor. Struct.*, **25**:1045–1058, 2017.

[24] Y. Cai and R.-C. Li. Perturbation Analysis for Matrix Joint Block Diagonalization. *Linear Algebra Appl.*, **581**:163–197, 2019.

[25] Y. Cai and C. Liu. An Algebraic Approach to Nonorthogonal General Joint Block Diagonalization. *SIAM J. Matrix Anal. Appl.*, **38**.1:50–71, 2017.

[26] J. F. Cardoso. High-order Contrasts for Independent Component Analysis. *Neural Computation*, **11**:157–192, 1999.

[27] J. F. Cardoso and A. Souloumiac. Blind beamforming for non Gaussian signals. *IEE Proceedings-F (radar and signal processing)*, **140**.6:362–370, 1993.

[28] G. Chabriel, M. Kleinsteuber, E. Moreau, H. Shen, P. Tichavsky, and A. Yeredor. Joint matrices decompositions and blind source separation: A survey of methods. *IEEE Signal Process. Mag.*, **31**.3:34–43, 2014.

[29] A. Chen and P. J. Bickel. Consistent Independent Component Analysis and Prewhitening. *IEEE Trans. Signal Process.*, **53**.10:3625–3632, 2005.

[30] A. Chen and P-J. Bickel. Efficient Independent Component Analysis. *Ann. Stat.*, **34**.6:2825–2855, 2006.

[31] K.-T. Chen. Integration of paths—a faithful representation of paths by noncommutative formal power series. *Trans. Amer. Math. Soc.*, **89**:395–407, 1958.

[32] X. Chen and Y. Fan. Estimation of Copula-Based Semiparametric Time Series Models. *Journal of Econometrics*, **130**.2:307–335, 2006.

[33] I. Chevyrev and T. Lyons. Characteristic functions of measures on geometric rough paths. *Ann. Probab.*, **44**.6:4049–4082, 2016.

[34] I. Chevyrev and H. Oberhauser. Signature Moments to Characterize Laws of Stochastic Processes. *Preprint*, page `arXiv:1810.1097`, 2018.

[35] S. Choi and A. Cichocki. Blind Separation of Nonstationary Sources in Noisy Mixtures. *Electronics Letters*, **36**.9:848–849, 2000.

[36] L. Colmenarejo and R. Preiß. Signatures of paths transformed by polynomial maps. *Preprint*, page `arXiv:1812.05962v1`, 2018.

[37] P. Comon. Independent Component Analysis, a new concept? *Signal Processing*, **36**.3:287–314, 1994.

[38] P. Comon and C. Jutten, editors. *Handbook of Blind Source Separation: Independent Component Analysis and Applications.* Academic Press, 2010.

[39] J. B. Conway. *Functions of One Complex Variable.* Second Edition. Graduate Texts in Mathematics **11**, Springer, 1978.

[40] K. Cranmer, J. Brehmer, and G. Louppe. The Frontier of Simulation-Based Inference. *Proceedings of the National Academy of Sciences*, **117**.48:30055–30062, 2020.

[41] N. Cvejic, D. Bull, and N. Canagarajah. Improving fusion of surveillance images in sensor networks using independent component analysis. *IEEE Transactions on Consumer Electronics*, **53**.3:1029–1035, 2007.

[42] S. B. Damelin and W. Miller Jr. *The Mathematics of Signal Processing.* Cambridge Texts in Applied Mathematics **48**, Cambridge University Press, 2012.

[43] G. Darmois. Analyse générale des liaisons stochastiques. *Rev. Inst. Intern. Stat.*, **21**:2–8, 1953.

[44] W. F. Darsow, B. Nguyen, and T. Olsen. Copulas and Markov Processes. *Illinois J. Math.*, **36**.4:600–642, 1992.

[45] A. Delorme and S. Makeig. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, **134**.1:9–21, 2004.

[46] M. Denker. The central limit theorem for dynamical systems. *Dynamical Systems and Ergodic Theory, (K. Krzyzewski, ed.)*, pages 33–62, Banach Center Publications, Polish Scientific Publishers, Warsaw, 1989.

[47] J. Diehl, P. Friz, and H. Mai. Pathwise stability of likelihood estimators for diffusions via rough paths. *The Annals of Applied Probability*, 26(4):2169–2192, 2016.

[48] Hua Ding, Yiliang Wang, Zhaojian Yang, and Olivia Pfeiffer. Nonlinear blind source separation and fault feature extraction method for mining machine diagnosis. *Applied Sciences*, 9(9):1852, 2019.

[49] P. Embrechts, A. McNeil, and D. Straumann. Correlation and Dependence in Risk Management: Properties and Pitfalls. *Risk Management: Value at Risk and Beyond*, **1**:176–223, 2002.

[50] T. Emura, T.-H. Long, and L.-H. Sun. R routines for performing estimation and statistical process control under copula-based time series models. *Communications in Statistics - Simulation and Computation*, **46**.4:3067–3087, 2017.

[51] J. Eriksson and V. Koivunen. Identifiability, separability and uniqueness of linear ICA models. *IEEE Signal Processing Letters*, **11**:601–604, 2004.

[52] J. Fan. A Selective Overview of Nonparametric Methods in Financial Econometrics. *Statistical Science*, **20**.4:317–337, 2005.

[53] G. B. Folland. *Real Analysis: Modern Techniques and Their Applications.* John Wiley & Sons, 1999.

[54] P. K. Friz and N. B. Victoir. *Multidimensional Stochastic Processes as Rough Paths: Theory and Applications.* Cambridge Studies in Advanced Mathematics **120**, Cambridge University Press, 2010.

[55] P. Fryzlewicz and S. S. Rao. Mixing properties of ARCH and time-varying ARCH processes. *Bernoulli*, **17**.1:320–346, 2011.

[56] D. J. H. Garling. Radonifying Mappings and Functional Central Limit Theorems. *Z. Wahrscheinlichkeitstheorie verw. Geb.*, **26**:297–308, 1973.

[57] T. Gowers, J. Barrow-Green, and I. Leader (eds.). *The Princeton Companion to Mathematics.* Princeton University Press, 2008.

[58] L. M. Graves. *Theory of functions of real variables.* McGraw-Hill, 1946.

[59] P. Green. *Note on Error Estimation.* Lecture Notes, University of Maryland, available at `http://www.math.umd.edu/~psg/401/errorestimation.pdf`, 2003.

[60] A. Gretton, R. Herbrich, A. Smola, O. Bousquet, and B. Schölkopf. Kernel methods for measuring independence. *J. Mach. Learn. Res.*, **6**:2075–2129, 2005.

[61] H. W Guggenheimer, A. S Edelman, and C. R. Johnson. A simple estimate of the condition number of a linear system. *College Math J.*, **26**:2–5, 1995.

[62] V. Guillemin and A. Pollack. *Differential Topology.* AMS Chelsea Publishing, 1974.

[63] J. Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, pages 49–52, 1902.

[64] P.R. Halmos. *Measure Theory.* Graduate Text in Mathematics **18**, Springer, 1950.

[65] H. Hälvä, S. Le Corff, L. Lehéricy, J. So, Y. Zhu, E. Gassiat, and A. Hyvarinen. Disentangling identifiable features from noisy data with structured nonlinear ICA. *35th Conference on Neural Information Processing Systems (NeurIPS)*, **34**:1624–1633, 2021.

[66] B. Hambly and T. Lyons. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Ann. Math.*, **171**.1:109–167, 2010.

[67] F. R. Hampel. A general qualitative definition of robustness. *Ann. Math. Stat.*, **42**.6:1887–1896, 1971.

[68] S. Harmeling, A. Ziehe, M. Kawanabe, and K.R. Müller. Kernel-based nonlinear blind source separation. *Neural Computation*, **15**.5:1089–1124, 2003.

[69] T. Hastie and R. Tibshirani. Independent component analysis through product density estimation. *Advances in Neural Information Processing Systems*, **15**:649–656, 2003.

[70] Q Peter He and Jin Wang. Statistical Process Monitoring as a Big Data Analytics Tool for Smart Manufacturing. *Journal of Process Control*, 67:35–43, 2018.

[71] N. J. Higham. *Accuracy and Stability of Numerical Algorithms.* SIAM, 2002.

[72] N. J. Higham. *Functions of Matrices : Theory and Computation.* SIAM, 2008.

[73] R. D. Hjelm et al. Learning deep representations by mutual information estimation and maximization. *Preprint*, page `arXiv:1808.06670v5[stat.ML]`, 2018.

[74] R. A. Horn and C. R. Johnson. *Matrix Analysis.* Cambridge University Press, 2012.

[75] P. Huber and E. Ronchetti. *Robust Statistics.* Wiley, 2009.

[76] A. Hyvärinen. Fast ICA for noisy data using Gaussian moments. In *ISCAS'99. Proceedings of the 1999 IEEE International Symposium on Circuits and Systems VLSI.*

[77] A. Hyvärinen. *Independent Component Analysis by Minimization of Mutual Information.* Technical Report (Report A46), Helsiniki University of Technology, Department of Computer Science and Engineering, September 1997.

[78] A. Hyvärinen. New Approximations of Differential Entropy for Independent Component Analysis and Projection Pursuit. *NeurIPS*, 1998.

[79] A. Hyvärinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans. Neural Netw.*, **10**.3:626–634, 1999.

[80] A. Hyvärinen. Independent Component Analysis: Recent Advances. *Phil. Trans. R. Soc. A*, **371**.1984:20110534, 2013.

[81] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis.* John Wiley & Sons, 2001.

[82] A. Hyvärinen and H. Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ICA. *NeurIPS2016*, pages 3765–3773, 2016.

[83] A. Hyvärinen and H. Morioka. Nonlinear ICA of temporally dependent stationary sources. *PMLR*, **54**:460–469, 2017.

[84] A. Hyvärinen and P. Pajunen. Nonlinear Independent Component Analysis: Existence and Uniqueness Results. *Neural Networks*, **12**.3:429–439, 1999.

[85] A. Hyvärinen, H. Sasaki, and R. Turner. Nonlinear ICA Using Auxiliary Variables and Generalized Contrastive Learning. *Proc. Artificial Intelligence and Statistics (AISTATS2019)*, 2019.

[86] Ö. Ican and T. B. Celik. Stock Market Prediction Performance of Neural Networks: A Literature Review. *Int. J. Finance Econ.*, **9**.11:100–108, 2017.

[87] C. Kanzow. *Numerik linearer Gleichungssysteme: Direkte und iterative Verfahren.* Springer-Verlag, 2007.

[88] M. Karow and D. Kressner. On a Perturbation Bound for Invariant Subspaces of Matrices. *SIAM J. Matrix Anal. Appl.*, **35**.2:599–618, 2014.

[89] M. Kawanabe and K.-R. Müller. Estimating Functions for Blind Separation When Sources Have Variance Dependencies. *J. Mach. Learn. Res.*, **6**.4, 2005.

[90] A Kechris. *Classical Descriptive Set Theory.* Graduate Texts in Mathematics **156**, Springer, 1995.

[91] J. L. Kelley. *General Topology.* D. van Nostrand, 1961.

[92] C. Kervazo, N. Gillis, and N. Dobigeon. Provably robust blind source separation of linear-quadratic near-separable mixtures. *SIAM J. Imaging Sci.*, **14**.4:1848–1889, 2021.

[93] I. Khemakhem, D. P. Kingma, R. P. Monti, and A. Hyvärinen. Variational autoencoders and nonlinear ICA: A unifying framework. *Proc. Artificial Intelligence and Statistics (AISTATS2020)*, 2020.

[94] S. A. Khoshnevis and R. Sankar. Applications of Higher Order Statistics in Electroencephalography Signal Processing: A Comprehensive Survey. *IEEE Reviews in Biomedical Engineering*, **13**:169–183, 2019.

[95] P. Kidger and T. Lyons. Signatory: differentiable computations of the signature and logsignature transforms, on both CPU and GPU. *arXiv: 2001. 00706 and https: // github. com/ patrick-kidger/ signatory ; published at ICLR 2021.*

[96] C. Kleiber and J. Stoyanov. Multivariate distributions and the moment problem. *J. Multivariate Anal.*, **113**:7–18, 2013.

[97] A. Kuczmaszewska. On the Strong Law of Large Numbers for ϕ-Mixing and ρ-Mixing Random Variables. *Acta Math. Hungar.*, **132**.1-2:174–189, 2011.

[98] J. M. Lee. *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics **218**, Springer, 2013.

[99] S.-I. Lee and S. Batzoglou. Application of Independent Component Analysis to Microarrays. *Genome Biology*, **4**.11:1–21, 2003.

[100] B. Li, J. Lu, and Z. Yu. Vector-wise joint diagonalization of almost commuting matrices. *arXiv preprint arXiv:2205.15519*, 2022.

[101] H. Li. Asynchronism-based principal component analysis for time series data mining. *Expert Systems with Applications*, **41**.6:2842–2850, 2014.

[102] Z. Li, X. Yan, X. Wang, and Z. Peng. Detection of gear cracks in a complex gearbox of wind turbines using supervised bounded component analysis of vibration signals collected from multi-channel sensors. *Journal of Sound and Vibration*, **371**:406–433, 2016.

[103] W. Linde. Moments of measures on banach spaces. *Math. Ann.*, **258**:277–287, 1982.

[104] T. J. Lyons, M. Caruana, and T. Lévy. *Differential Equations Driven by Rough Paths*. Springer, 2007.

[105] T. J. Lyons and Z. Qian. *System Control and Rough Paths*. Oxford Mathematical Monographs, Oxford University Press, 2002.

[106] S. Makeig and J. Onton. ERP features and EEG dynamics: an ICA perspective. In *Oxford Handbook of Event-Related Potential Components*. Oxford University Press, 2011.

[107] M. B. Marcus and L. A. Shepp. Sample behavior of gaussian processes. *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory*, pages 423–441, 1972.

[108] H. Masuda. Mixing properties of ARMA processes. *Stochastic Process. Appl.*, **19**:297–303, 1988.

[109] H. Masuda. Ergodicity and exponential β-mixing bounds for multidimensional diffusions with jumps. *Stochastic Process. Appl.*, **117**.1:35–56, 2007.

[110] R. A. McCoy. Second countable and separable function spaces. *Am. Math. Mon.*, **85**.6:487–489, 1978.

[111] J. Miettinen, K. Nordhausen, and S. Taskinen. Blind source separation based on joint diagonalization in R: The packages JADE and BSSasymp. *Journal of Statistical Software*, **76**.2, 2017.

[112] E. Moulines, J. F. Cardoso, and E. Gassiat. Maximum likelihood for blind separation and deconvolution of noisy signals using mixture models. *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'97)*, page 3617–3620, 1997.

[113] J. R. Munkres. *Topology.* Second Edition. Prentice Hall, 2000.

[114] R. B. Nelsen. *An Introduction to Copulas.* Springer Series in Statistics. Second Edition. Springer, 2006.

[115] F. Noé, A. Tkatchenko, K.-R. Müller, and C. Clementi. Machine Learning for Molecular Simulation. *Annual Review of Physical Chemistry*, **71**:361–390, 2020.

[116] K. Nordhausen. On robustifying some second order blind source separation methods for nonstationary time series. *Statistical Papers*, **55**:141–156, 2014.

[117] K. Nordhausen and A. Ruiz-Gazen. On the Usage of Joint Diagonalization in Multivariate Statistics. *J. Multivariate Anal.*, **188**:104844, 2022.

[118] D. Nuzillard and A. Bijaoui. Blind source separation and analysis of multispectral astronomical images. *Astron. Astrophys. Suppl. Ser.*, **147**.1:129–138, 2000.

[119] H. Oberhauser and A. Schell. A Robustness Analysis of Blind Source Separation. *in preparation (2022).*

[120] H. Oberhauser and A. Schell. Nonlinear Independent Component Analysis for Continuous-Time Signals. *arXiv preprint* `arXiv:2102.02876`, 2021.

[121] A. Olevskii and A. Ulanovskii. Reconstruction of Signals: Uniqueness and Stable Sampling. In *Sampling: Theory and Applications*, pages 9–49. Springer, 2020.

[122] V. M. Panaretos and Y. Zemel. *An Invitation to Statistics in Wasserstein Space.* Springer Briefs in Probability and Mathematical Statistics. Springer Nature, 2020.

[123] A Papavasiliou and C. Ladroue. Parameter Estimation for Rough Differential Equations. *Ann. Stat.*, **39**.4:2047–2073, 2011.

[124] P. Petersen, M. Raslan, and F. Voigtlaender. Topological Properties of the Set of Functions Generated by Neural Networks of Fixed Size. *Found. Comput. Math*, **21**:375–444, 2021.

[125] N. Pfister, S. Weichwald, P. Bühlmann, and B. Schölkopf. Robustifying Independent Component Analysis by Adjusting for Group-Wise Stationary Noise. *J. Mach. Learn. Res.*, **20**.147:1–50, 2019.

[126] D. T. Pham and P. Garrat. Blind Separation of Mixture of Independent Sources Through a Quasi-Maximum Likelihood Approach. *IEEE Trans. Signal Process.*, **45**.7:1712–1725, 1997.

[127] T. D. Pham and L. T. Tran. Some mixing properties of time series models. *Stochastic Process. Appl.*, **19**:297–303, 1985.

[128] C. E. Rasmussen and C-K. I. Williams. *Gaussian Processes for Machine Learning.* The MIT Press, 2006.

[129] M. Reed and B. Simon. *Methods of Modern Mathematical Physics: Vol. I, Functional Analysis.* Academic Press, 1972.

[130] O. Reiersøl. Identifiability of a linear relation between variables which are subject to error. *Econometrica: Journal of the Econometric Society*, pages 375–389, 1950.

[131] R. Remmert. *Theory of Complex Functions.* (4th corrected printing.) Springer, 1998.

[132] C. Reutenauer. *Free Lie Algebras.* London Mathematical Society Monographs, New Series. Oxford Science Publications **7**, The Clarendon Press, 1993.

[133] rmcerafl (https://dsp.stackexchange.com/users/59697/rmcerafl) and Gillespie (https://dsp.stackexchange.com/users/55647/gillespie). Blind source separation for asynchronously observed mixture channels. dsp.stackexchange. `https://dsp.stackexchange.com/questions/78796` (version: 2022-09-16).

[134] L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes, and Martingales.* Volume 1, Cambridge University Press, 2000.

[135] M. Rosenblatt. *Markov Processes, Structure and Asymptotic Behavior.* Springer-Verlag, New York, 1971.

[136] W. Rudin. *Principles of Mathematical Analysis.* (3rd ed.) McGraw-Hill, 1976.

[137] W. Rudin. *Real and Complex Analysis.* (3rd ed.) McGraw-Hill, 1987.

[138] A. Samarov and A. Tsybakov. Nonparametric Independent Component Analysis. *Bernoulli*, **10**.4:565–582, 2004.

[139] R.J. Samworth and M. Yuan. Independent Component Analysis via Nonparametric Maximum Likelihood Estimation. *Ann. Stat.*, **40**.6:2973–3002, 2012.

[140] A. Schell. *SigNICA.* GitHub repository, `https://github.com/alexander-schell/SigNICA.git`; Python Code for Section 4.8, 2021.

[141] E. Schlemm and R. Stelzer. Multivariate CARMA processes, continuous-time state space models and complete regularity of the innovations of the sampled processes. *Bernoulli*, **18**.1:46–63, 2012.

[142] L. Schwartz. Applications radonifiantes. *Séminaire de l'École Polytechnique, Paris*, 1969-1970.

[143] L. Schwartz. Probabilités cylindriques et applications radonifiantes. *J. Fac. Sci. Univ. Tokyo, Sect. IA*, **18**:139–286, 1971.

[144] D. Shi, Y. Cai, and S. Xu. Some Perturbation Results for a Normalized Non-Orthogonal Joint Diagonalization Problem. *Linear Algebra Appl.*, **484**:457–476, 2015.

[145] V. P. Skitovich. On a Property of a Normal Distribution. *Doklady Akad. Nauk. SSSR*, **89**:217–219 (in Russian), 1953.

[146] Y. Takahashi. On the relation between Radonifying mappings and kernels of probability measures on Banach spaces. *Hokkaido Mathematical Journal*, 14(1):97–106, 1985.

[147] A. Takahata, E. Nadalin, R. Ferrari, L. Duarte, R. Suyama, R. Lopes, J. Romano, and M. Tygel. Unsupervised processing of geophysical signals: A review of some key aspects of blind deconvolution and blind source separation. *IEEE Signal Processing Magazine*, **29**.4:27–35, 2012.

[148] Y. Tan, J. Wang, and J. M. Zurada. Nonlinear blind source separation using a radial basis function network. *IEEE Trans. Neural Netw.*, **12**.1:124–134, 2001.

[149] T. et al. Teshima. Coupling-based invertible neural networks are universal diffeomorphism approximators. *arXiv: 2006. 11469, accepted at NeurIPS 2020*, 2020.

[150] F. J. Theis and Y. Inouye. On the use of joint diagonalization in blind signal processing. *IEEE International Symposium on Circuits and Systems*, 2006.

[151] K. Usevich, J. Li, and P. Comon. Approximate matrix and tensor diagonalization by unitary transformations: Convergence of jacobi-type algorithms. *SIAM J. Optim.*, **30**.4:2998–3028, 2020.

[152] A. W. van der Vaart. *Asymptotic Statistics.* Cambridge Series in Statistical and Probabilistic Mathematics, 1998.

[153] A. W. van der Vaart. *Time Series.* Version 242013, Lecture Notes, Universiteit Leiden, 2013. (Available at `http://www.math.leidenuniv.nl/~avdvaart/timeseries/dictaat.pdf`.).

[154] W. Walter. *Analysis 2.* (3rd ed.) Springer, 1992.

[155] J. A. Wellner. *Empirical Processes: Theory and Applications.* Special Topics Course Spring 2005, Delft Technical University, June 2005. (Available at `https://www.stat.washington.edu/people/jaw/RESEARCH/TALKS/Delft/emp-proc-delft-big.pdf`).

[156] C. Xiaohong, L. P. Hansen, and M. Carrasco. Nonlinearity and temporal dependence. *Econometrics*, **155**.2:155–169, 2010.

[157] J. Yeh. *Martingales and Stochastic Analysis.* Series on Multivariate Analysis **1**, World Scientific Publishing, 1995.

[158] G. Zhou, S. Xie, Z. Yang, and J. Zhang. Nonorthogonal Approximate Joint Diagonalization With Well-Conditioned Diagonalizers. *IEEE Trans. Neural Netw. Learn. Syst.*, **20**.11:1810–1819, 2009.

[159] Q. Zhou and J. Yu. Asymptotic theory for linear diffusions under alternative sampling schemes. *Econonomics Letters*, **128**:1–5, 2015.