Canonical Spaces and Hybrid Operator Algebras in Fractional Calculus

From Impossibility to Boundary Trace Theory

Ariel Daley

March 2026

Contents

Preface

This book grew out of the research program documented in the papers AD01 through AD12. Those twelve papers, written over a concentrated period, address a single organizing question: on which spaces do fractional differential and integral operators acquire a structurally transparent algebraic form? The papers develop the answer progressively, from a negative result on finite-dimensional polynomial spaces to a fully developed theory of canonical bases, hybrid operator algebras, weighted Banach completions, transform models, semigroup generation, and ordered boundary trace calculus.

The papers were written for a research audience. This book is not. Its purpose is to make the ideas, results, and internal logic of the entire program accessible to a mathematically mature undergraduate reader—someone in the third year of a degree in the mathematical sciences, with solid preparation in calculus, real analysis, point-set topology, linear algebra, and the basics of group theory and ring theory, but who may have little or no prior background in complex analysis, functional analysis, spectral theory, semigroup theory, or discrete fractional calculus.

What this book is

This is a textbook, not a survey article and not a research monograph. Its exposition is theorem–proof–example–remark driven, in the style of a graduate-level text, but pitched at an advanced undergraduate level. Every concept is motivated before it is defined. Every theorem is placed in context before it is proved. Prerequisite ideas—from complex powers and closed operators to C_0-semigroups—are introduced exactly when they become necessary, not front-loaded in a massive preliminary chapter.

The book is organized as a single coherent narrative, not as a chapter-by-chapter paraphrase of the twelve papers. The order of exposition has been rearranged wherever pedagogy demands it, proofs have been rewritten for clarity, and special cases and low-dimensional examples have been added throughout. When a paper's results are more general than is needed for a first presentation, the book presents a clean illuminating case first and widens the scope afterward.

The arc of the book

The story has six stages, and the chapter structure reflects them.

The first stage (Chapters 1–3) sets the scene. Chapter 1 explains the central question and previews the entire program through two familiar pictures from ordinary calculus: the shift picture on normalized monomials and the spectral picture on exponentials. Chapter 2 provides the minimum preparation in continuous fractional calculus—Riemann–Liouville integrals, Caputo derivatives, and the Gamma, Beta, and Mittag–Leffler functions. Chapter 3 proves the first main result: the impossibility of an internal fractional model on finite-dimensional polynomial spaces.

The second stage (Chapters 4–7) builds the core algebraic theory. Chapter 4 constructs the one-variable canonical shift model. Chapter 5 extends it to several variables. Chapter 6 develops the whole-space spectral model. Chapter 7 unifies the two on mixed domains, producing the hybrid shift-spectral algebra.

The third stage (Chapters 8–9) passes from algebra to analysis. Chapter 8 introduces weighted Banach completions, proves that the algebraic relations persist, constructs the fiberwise holomorphic transform model, and shows that the generating eigenvectors become genuine Banach-space elements. Chapter 9 enlarges the completed space by adjoining ordered boundary-trace layers and classifies the maximal graded invariant sector on which the extended Caputo tuple remains commuting.

The fourth stage (Chapters 10–13) develops the discrete counterpart. Chapter 10 introduces the discrete building blocks—rising factorials, nabla fractional sums, Caputo nabla differences, and lattice characters. Chapters 11, 12, and 13 mirror the continuous sequence exactly: the discrete hybrid algebra, its Banach completion and Z-transform model, and its boundary-augmented maximal commuting sectors.

The fifth stage (Chapter 14) unifies the continuous and discrete theories in a single abstract coefficient-space framework. It proves the optimal-weight theorem, constructs the common transform model, and develops the semigroup generation theory for diagonal multipliers and mixed shift-spectral generators.

The sixth and final stage (Chapters 15–16) develops the culminating algebraic structure. Chapter 15 introduces the ordered boundary trace calculus, identifies the ordering-defect space as the range of the commutator ideal, and constructs the universal commuting quotient. Chapter 16 looks back at the whole program, distills the main themes, and suggests further reading.

Recurring themes

Several ideas recur throughout the book, and the reader should watch for them:

- the failure of finite-dimensional ordinary polynomial spaces and the necessity of infinite-dimensional, differently graded canonical spaces;
- the search for canonical bases that reveal the true form of the operators;
- the geometric distinction between one-sided domains (which produce shifts) and whole-space domains (which produce spectral diagonalization);
- the hybridization of the shift and spectral pictures on mixed domains;
- the passage from algebraic model-building to genuine Banach-space operator theory through weighted completion;
- the exact structural parallelism between the continuous and discrete theories;
- the fact that noncommutativity in the multi-variable theory is entirely boundary-generated;
- the two complementary resolutions of noncommutativity: restriction to a maximal commuting sector and quotient by the commutator ideal.

Prerequisites

The reader is assumed to be comfortable with: multivariable calculus, including the Gamma and Beta integrals; real analysis at the level of Rudin's *Principles* or an equivalent course, including uniform convergence and basic metric-space topology; linear algebra through Jordan normal form and dual spaces; and the definitions of groups, rings, and ideals at an introductory level.

No prior knowledge of fractional calculus is assumed. No prior knowledge of functional analysis, operator theory, complex analysis beyond the level of power series, or semigroup theory is assumed; these topics are developed from scratch, to the extent needed, as they arise in the narrative.

How to read this book

The book is designed to be read linearly, but a first reading need not cover every chapter with the same depth. Chapter 16 contains specific suggestions for a first pass through the core chapters and a second pass through the full theory. As a general principle: the reader who understands the one-variable shift model (Chapter 4), the hybrid algebra (Chapter 7), the Banach completion (Chapter 8), and the boundary-augmented theory (Chapter 9) has grasped the conceptual spine of the program. Everything else—the multi-variable extensions, the discrete mirror, the abstract unification, and the commutator ideal theory—adds breadth and depth to that spine.

The two-variable case $r = 2$ is the smallest setting in which every nontrivial phenomenon of the theory is visible. The reader is encouraged to work through the $r = 2$ examples by hand whenever they appear.

Notation

A summary of the principal notation is given in Section 1.4 of Chapter 1. The notation has been kept as consistent as possible across all sixteen chapters; where a notational choice from the original research papers has been adjusted for textbook clarity, this is noted explicitly in the text.

Ariel Daley
Seoul, March 2026

Chapter 1

Problem Setting and the Global Map

This book is about a single organizing question: on which spaces do fractional differential and integral operators acquire a transparent algebraic form? The question may sound technical, but its roots are firmly in the most classical part of calculus. In this opening chapter we explain the question, show why it is natural, and describe the path that the rest of the book will follow.

1.1 Two familiar pictures from ordinary calculus

A basic lesson of linear algebra is that the choice of basis determines how transparent a linear map appears. The same operator can look complicated in one basis and almost trivial in another. Ordinary calculus already provides two striking illustrations of this principle, and they will guide the entire book.

The shift picture

Consider the sequence of normalized monomials

$$e_n(x) = \frac{x^n}{n!}, \qquad n = 0, 1, 2, \ldots$$

These are the ordinary monomials $1, x, x^2, \ldots$ divided by the factorials that would otherwise clutter the formulas. In this basis, differentiation takes a remarkably simple form:

$$\frac{d}{dx}\, e_n(x) = e_{n-1}(x) \qquad (n \geq 1), \qquad \frac{d}{dx}\, e_0(x) = 0.$$

Differentiation simply moves each basis vector one step down the chain, and the constant function $e_0 = 1$ is annihilated. Integration $\int_0^x$ moves each basis vector one step up:

$$\int_0^x e_n(t)\, dt = e_{n+1}(x).$$

In the language of operator theory, integration is a *forward shift* and differentiation is a *backward shift* on this chain, with e_0 playing the role of a *vacuum vector*—the bottom of the ladder, destroyed by the backward shift.

Remark 1.1.1. The word "shift" is used deliberately. In a unilateral shift algebra, one has an operator S that raises an index and an operator S^* that lowers it, subject to the relations $S^*S = I$ and $SS^* = I - P_0$, where P_0 projects onto the vacuum. The normalized monomials make differentiation and integration realize exactly this pattern. Much of this book is devoted to finding the analogous pattern for operators of fractional order.

The spectral picture

Now consider the exponential functions

$$\varphi_\lambda(x) = e^{\lambda x}, \qquad \lambda \in \mathbb{C}.$$

Each exponential is an eigenvector of d/dx:

$$\frac{d}{dx} e^{\lambda x} = \lambda \, e^{\lambda x}.$$

In the family $\{e^{\lambda x}\}_\lambda$, differentiation does not move vectors along a chain; instead it multiplies each vector by its own eigenvalue. The operator is *diagonalized*: knowing the eigenvalue λ is the same as knowing the effect of the operator.

This is the *spectral picture*. The word "spectral" comes from the analogy with spectral theory in linear algebra and functional analysis, where an operator is understood by describing its eigenvalues (its spectrum) and the corresponding eigenspaces.

Remark 1.1.2. The distinction between the shift picture and the spectral picture is not merely a matter of taste. As we shall see, these two pictures are tied to different geometric settings. The shift picture is the natural one on a half-line $(0, \infty)$ or more generally on a domain with a boundary, because the boundary creates a vacuum. The spectral picture is the natural one on the whole real line $\mathbb{R}$ or on $\mathbb{R}^d$, where there is no distinguished boundary and the relevant symmetry is translation invariance. One of the main themes of this book is that *both* pictures have exact fractional-order analogues, and that these analogues can be combined into a single unified framework.

What makes these pictures useful

Why do we care about finding a "nice" basis for an operator? Because a transparent representation makes computations simple, reveals structural properties, and often provides the right language for solving equations. For instance:

- In the shift picture, solving the equation $y' = f$ reduces to raising the index of each coefficient by one, which is a purely algebraic operation.
- In the spectral picture, solving $y' = \lambda y$ is immediate, and more complicated constant-coefficient equations reduce to polynomial algebra in the spectral variable.

The question that drives this book is whether the same kind of structural transparency can be achieved for *fractional-order* operators—operators that generalize differentiation and integration to non-integer orders.

1.2 The central question of the book

Fractional calculus extends the classical notions of differentiation and integration to arbitrary real (or even complex) orders. For instance, one can define operators J^α and D^α that act as "integration of order α" and "differentiation of order α," where α need not be an integer. We will define these operators precisely in Chapter 2, but for now it suffices to know that they exist, that they are well established in the mathematical literature, and that they are genuinely useful in applications ranging from diffusion equations to viscoelasticity.

The classical formulas for these operators, however, involve integrals with singular kernels and special-function coefficients that can obscure their structural nature. The experience of ordinary calculus suggests that the right remedy is not to compute harder but to find the right spaces.

> *The central question of this book is:* On which spaces, and with respect to which bases, do fractional differential and integral operators acquire a structurally transparent algebraic form—either as exact shifts or as exact diagonal multipliers?

This is *not* the same as asking for closed-form solutions to specific fractional differential equations, although such solutions will emerge naturally along the way. It is a question about the *architecture* of fractional calculus: what are the correct state spaces, and what is the correct algebraic skeleton of the theory?

The answer, as developed through the rest of this book, has several layers:

(i) The ordinary polynomial space P_n of polynomials of degree at most n is *not* the right space; this is the negative conclusion of Chapter 3.

(ii) On the half-line $(0, \infty)$, there exists a canonical graded space on which fractional integration and differentiation become exact forward and backward shifts. This is the *shift model* (Chapters 4–5).

(iii) On the whole space $\mathbb{R}^d$, exponential functions diagonalize the relevant operators, producing an exact *spectral model* (Chapter 6).

(iv) On mixed domains such as $(0, \infty)^r \times \mathbb{R}^s$, the two models combine into a *hybrid shift-spectral algebra* (Chapter 7).

(v) These algebraic models extend to genuine Banach-space operator theory through *weighted completions* and *transform models* (Chapters 8, 12, and 14).

(vi) The boundary-generated obstructions to commutativity, which appear when several one-sided coordinates interact, are captured by an *ordered boundary trace calculus*, *commutator ideals*, and *commuting quotient reductions* (Chapters 9, 13, and 15).

In short, the book builds, step by step, the right spaces and the right algebraic structures for fractional calculus, moving from the simplest setting to a fully general theory.

1.3 The flow from AD01 to AD12

This book is based on a sequence of twelve research papers, referred to throughout as AD01 through AD12. Each paper addresses one stage of the program described above. It is useful to see the whole arc before entering the details.

For the reader who wishes to consult the original papers alongside this book, we list them here with their full titles and persistent identifiers.

AD01 *The Nonexistence of Internal Fractional Models on P_n.*

https://doi.org/10.5281/zenodo.19020897

AD02 *Fractional Shift Algebra on a Canonical α-Graded Space.*

https://doi.org/10.5281/zenodo.19021257

AD03 *Partial Fractional Integrals and Caputo Derivatives as a Commuting Shift Algebra on a Canonical Multi-Graded Space.*

https://doi.org/10.5281/zenodo.19030049

AD04 *A Spectral Algebra for Multidimensional Weyl Fractional Operators on Exponential Characters.*

https://doi.org/10.5281/zenodo.19031699

AD05 *Hybrid Shift-Spectral Algebra on Mixed Regions.*

https://doi.org/10.5281/zenodo.19112602

AD06 *Banach Completions of Hybrid Shift-Spectral Algebra.*

https://doi.org/10.5281/zenodo.19132720

AD07 *Maximal Commuting Sectors in Boundary-Augmented Hybrid Fractional Calculus.*

https://doi.org/10.5281/zenodo.19133120

AD08 *Discrete Hybrid Shift-Spectral Algebra.*

https://doi.org/10.5281/zenodo.19133409

AD09 *Banach Completions and Z-Transform Model for Discrete Hybrid Shift-Spectral Algebra.*

https://doi.org/10.5281/zenodo.19151690

AD10 *Maximal Commuting Sectors in Boundary-Augmented Discrete Hybrid Fractional Calculus.*

https://doi.org/10.5281/zenodo.19174686

AD11 *Transform Models, Semigroup Generation, and Optimal Weights for Hybrid Fractional Operator Algebras.*

https://doi.org/10.5281/zenodo.19184833

AD12 *Ordered Boundary Trace Calculus and Commutator Ideals for Partial Caputo Tuples.*

https://doi.org/10.5281/zenodo.19185137

All twelve papers are open-access and can be retrieved from Zenodo using the DOIs above. The labels AD01–AD12 will be used consistently throughout this book whenever the original papers are cited.

We now describe the role each paper plays in the overall program.

Stage 1: Clearing the ground

The first paper (AD01) proves that the most obvious candidate—the finite-dimensional polynomial space P_n—cannot support an internal fractional model. Two independent obstructions are identified. First, the classical Riemann–Liouville and Caputo operators simply do not preserve P_n. Second, the differentiation operator on P_n is a nilpotent Jordan block, and a nilpotent matrix has no nontrivial roots. This double negative result is not a dead end; it is a precise signpost telling us that the correct space must be infinite-dimensional and differently graded.

Stage 2: Building the shift model

The second paper (AD02) constructs the one-variable canonical model. The key insight is to replace the ordinary grading $0, 1, 2, 3, \dots$ by the α-grading $0, \alpha, 2\alpha, 3\alpha, \dots$ and to work with the basis

$$e_n(x) = \frac{x^{n\alpha}}{\Gamma(n\alpha + 1)}, \qquad n = 0, 1, 2, \dots$$

On this basis, the Riemann–Liouville integral of order α is an exact forward shift and the Caputo derivative of order α is an exact backward shift. The constant function $e_0 = 1$ is the vacuum vector, annihilated by the backward shift, just as in the integer-order case. Moreover, this canonical basis is essentially unique.

The third paper (AD03) extends the shift model to several variables, replacing the chain $\mathbb{N}_0$ by the multi-index lattice $\mathbb{N}_0^r$. The coordinatewise fractional operators become commuting shifts in different lattice directions, and the one-variable vacuum vector becomes a family of boundary hyperplanes.

Stage 3: The spectral model

The fourth paper (AD04) addresses what happens on the whole space $\mathbb{R}^d$, where there is no boundary. Here the shift picture is not the natural one. Instead, the exponential characters $e_\lambda(x) = e^{\langle \lambda, x \rangle}$ diagonalize the Weyl-type fractional operators: each exponential is an eigenvector with eigenvalue determined by the spectral parameter λ. The operator algebra becomes a diagonal multiplier algebra, and there is no vacuum defect.

Stage 4: Hybridization

The fifth paper (AD05) unifies the shift and spectral models on mixed domains of the form $(0, \infty)^r \times \mathbb{R}^s$. The canonical basis is a tensor product of normalized fractional monomials (one-sided directions) and exponential characters (whole-space directions). The resulting *hybrid shift-spectral algebra* has a shift block and a spectral block that commute with each other. All defect phenomena remain localized in the shift block.

Stage 5: Completion and transform theory

The algebraic models built so far live on spaces of finite linear combinations of basis vectors. For analysis—for convergence, for operator norms, for spectral theory, for semigroups—one needs completions. The sixth paper (AD06) introduces weighted Banach coefficient spaces and shows that the hybrid operators extend to bounded or closed operators on these completions. It also constructs a fiberwise holomorphic transform model on a polydisk.

The ninth paper (AD09) carries out the same program in the discrete setting, using the Z-transform in place of the holomorphic model. The eleventh paper (AD11) then recognizes that the continuous and discrete completion theories are two realizations of a single abstract coefficient-space framework. In this unified setting, it proves sharp optimal-weight theorems, constructs a common transform model, and develops a C_0-semigroup generation theory for diagonal multipliers and mixed generators.

Stage 6: Boundary structure and commutator theory

On the algebraic core, the coordinatewise operators in a multi-variable setting commute. But once one enlarges the space to include boundary-trace information, commutativity can fail. The seventh paper (AD07) introduces ordered boundary-trace words to record the order in which different coordinates reach their boundary, computes explicit commutator formulas, and classifies the maximal graded invariant sector on which the extended tuple remains commuting. The tenth paper (AD10) establishes the same results in the discrete theory.

The twelfth paper (AD12) goes further. It isolates the abstract mechanism by defining an ordered boundary trace calculus, identifies the ordering-defect space as the range of the commutator ideal, and constructs a universal commuting quotient. In this quotient, ordered trace words that share the same underlying support are identified, the induced operator tuple becomes commuting, and the quotient enjoys a universal property. The earlier maximal commuting sectors (obtained by restriction) and the commuting quotient (obtained by factoring out the commutator ideal) provide two complementary ways of recovering commutativity from a noncommutative ambient space.

The global picture

The reader will notice a deliberate parallelism: the continuous theory (AD05–AD07) and the discrete theory (AD08–AD10) follow the same three-step pattern of algebraic model, Banach completion, and boundary structure. This parallelism is not a coincidence; it reflects a common underlying architecture that AD11 makes explicit.

1.4 Reading conventions and notation

We collect here the notational conventions that will be used throughout the book. The reader may wish to return to this section as needed rather than memorize it on a first reading.

Numbers and index sets

We write $\mathbb{N} = \{1, 2, 3, \dots\}$ for the positive integers and $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ for the nonnegative integers. The symbols $\mathbb{R}$, $\mathbb{C}$, $\mathbb{Z}$ have their standard meanings. The open right half-plane is

$$\mathbb{C}_+ = \{z \in \mathbb{C} : \operatorname{Re} z > 0\}.$$

Multi-index notation

For $r \geq 1$, a multi-index is a vector $\mathbf{k} = (k_1, \dots, k_r) \in \mathbb{N}_0^r$. We write $|\mathbf{k}| = k_1 + \cdots + k_r$, $\mathbf{k}! = k_1! \cdots k_r!$, and $\mathbf{k} \leq \mathbf{m}$ to mean $k_i \leq m_i$ for every i. The coordinatewise partial order $\leq$ on $\mathbb{N}_0^r$ will play a central role in describing boundary layers.

Fractional orders

Throughout most of the book, α denotes a fixed fractional order with $0 < \alpha < 1$. In the multi-variable setting, the order tuple is $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r)$ with each $\alpha_i \in (0, 1)$.

Canonical basis vectors

The one-variable canonical basis vectors are

$$e_n(x) \;=\; \frac{x^{n\alpha}}{\Gamma(n\alpha + 1)}, \qquad n \in \mathbb{N}_0.$$

In several variables, the canonical basis on $(0, \infty)^r$ is

$$e_{\mathbf{k}}(x) \;=\; \prod_{i=1}^{r} \frac{x_i^{k_i\alpha_i}}{\Gamma(k_i\alpha_i + 1)}, \qquad \mathbf{k} \in \mathbb{N}_0^r.$$

When both one-sided and whole-space variables are present, the hybrid basis element on $(0, \infty)^r \times \mathbb{R}^s$ is

$$e_{\mathbf{k},\lambda}(x, y) \;=\; e_{\mathbf{k}}(x)\, e^{\langle \lambda, y \rangle}, \qquad \mathbf{k} \in \mathbb{N}_0^r,\ \lambda \in \Lambda \subset (\mathbb{C}_+)^s.$$

Operators

The principal operators in the book are:

- $J_i = {}_0I_{x_i}^{\alpha_i}$: the Riemann–Liouville fractional integral of order α_i in the i-th coordinate (forward shift).
- $C_i = {}_0^{\mathrm{C}}D_{x_i}^{\alpha_i}$: the Caputo fractional derivative of order α_i in the i-th coordinate (backward shift).
- $W^{\boldsymbol{\beta}}$: the Weyl fractional operator of multi-order $\boldsymbol{\beta}$ (spectral multiplier).
- M_σ: the diagonal spectral multiplier by a symbol σ.
- Π_0, $\Pi_{<\mathbf{m}}$: projections onto the vacuum component and onto boundary-layer sectors.

Vacuum and boundary language

In the one-variable shift model, the *vacuum* is the grade-zero subspace spanned by $e_0 = 1$. The backward shift (Caputo derivative) annihilates the vacuum: $C_\alpha e_0 = 0$. More generally, a *boundary layer* or *defect sector* is the subspace on which the full semigroup property of the backward shift fails. In several variables, the boundary layers are indexed by coordinate hyperplanes where one or more of the grade indices k_i is zero.

Ordered trace words and commutator language

In the boundary-augmented theories of Chapters 9, 13, and 15, the order in which different coordinates reach their boundary is recorded by *ordered words* over the alphabet $\{1, \dots, r\}$. For instance, the words ij and ji both have *support* $\{i, j\}$ but record different boundary orderings. The *commutator ideal* is generated by the differences arising from distinct orderings of the same support, and the *commuting quotient* is obtained by identifying all orderings that share the same support.

Weights and completions

A *weight* is a positive function $\omega\colon \mathbb{N}_0^r \times \Lambda \to (0, \infty)$ used to define a norm on coefficient sequences. A weight is *shift-admissible* if it makes both the forward and backward shifts bounded. A weight is *geometric* if $\omega(\mathbf{k}, \lambda) = \eta(\lambda)\, \rho_1^{k_1} \cdots \rho_r^{k_r}$ for some positive constants ρ_i and a fiber function η.

1.5 Why the word "model" is appropriate

Throughout this book we refer to the algebraic and analytic realizations of fractional operators as *models*. This word is chosen with care, and it is worth pausing to explain what it means.

In operator theory, a *model* for a class of operators is a concrete, canonical realization on a specific Hilbert or Banach space in which the operators take a particularly simple and revealing form. The most celebrated example is the *Sz.-Nagy–Foiaş model theory* for contractions on Hilbert spaces, in which every contraction is realized as the compression of a multiplication operator on a Hardy space. Model theory transforms abstract operator-theoretic problems into concrete function-theoretic problems.

The constructions in this book serve the same purpose for fractional operators. When we say that the canonical graded space $\mathcal{G}_\alpha^{\mathrm{alg}}$ is a *shift model* for the Caputo–Riemann–Liouville pair, we mean that these operators are realized *exactly* as a forward and backward shift on that space. This is not a symbolic analogy or an approximation; it is an identity. The model reveals the true algebraic skeleton of the operators, just as a Jordan normal form reveals the true structure of a nilpotent matrix.

Similarly, the spectral model on exponential characters realizes Weyl operators as exact diagonal multipliers. The hybrid model combines both realizations on a mixed domain. The Banach completions and transform models are analytic models in which the algebraic relations persist in the presence of topologies, norms, and convergence.

Remark 1.5.1. One should resist the temptation to read "model" as "approximate picture" or "heuristic analogy." In this book, a model is always an exact algebraic or analytic realization. When we write $J_\alpha e_n = e_{n+1}$, this is not a metaphor; it is an identity of functions on $(0, \infty)$.

Looking ahead

The next chapter introduces the minimum preparation from continuous fractional calculus that is needed to state and prove the results of the following chapters. We will define the Riemann–Liouville integral and the Caputo derivative, recall the essential properties of the Gamma and Mittag-Leffler functions, and perform the first concrete calculations that preview the canonical basis. The reader who already has some background in fractional calculus may read Chapter 2 quickly, pausing only at the calculations in its final section, which connect directly to the shift model of Chapter 4.

Chapter 2

Preparation in Continuous Fractional Calculus

This chapter introduces the basic operators and special functions of fractional calculus that are needed in the rest of the book. The treatment is self-contained: we assume the reader knows single-variable real analysis (including the theory of the Riemann integral on $(0, \infty)$) and the definition of the Gamma function, but nothing more. The goal is not to survey fractional calculus in general—the literature is vast—but to develop exactly the tools that will appear in the canonical models of the following chapters.

Readers who already have experience with fractional calculus may wish to skim this chapter quickly, pausing at the concrete calculations in the final section, which lead directly into the shift model of Chapter 4.

2.1 From iterated integrals to fractional integrals

The most natural route to fractional-order integration begins with a classical observation about repeated integration.

The Cauchy formula for iterated integrals

Let f be a continuous function on $(0, \infty)$. Define the integration operator based at 0 by

$$(If)(x) = \int_0^x f(t)\,dt.$$

Applying this operator twice gives

$$(I^2 f)(x) = \int_0^x \int_0^{t_1} f(t_2)\,dt_2\,dt_1.$$

By exchanging the order of integration (Fubini's theorem), this double integral simplifies to a single integral:

$$(I^2 f)(x) = \int_0^x (x-t)\,f(t)\,dt.$$

More generally, the n-fold iterated integral can be written as a single convolution-type integral:

Proposition 2.1.1 (Cauchy's formula for repeated integration). *For every $n \in \mathbb{N}$ and every continuous f on $(0, \infty)$,*

$$(I^n f)(x) = \frac{1}{(n-1)!} \int_0^x (x-t)^{n-1} f(t)\,dt, \qquad x > 0.$$

Proof. The case $n = 1$ is the definition of I. For $n = 2$, we have just verified the formula. Suppose the formula holds for some $n \geq 1$. Then

$$(I^{n+1} f)(x) = \int_0^x (I^n f)(s)\, ds = \int_0^x \frac{1}{(n-1)!} \int_0^s (s-t)^{n-1} f(t)\, dt\, ds.$$

Exchanging the order of integration gives

$$(I^{n+1} f)(x) = \frac{1}{(n-1)!} \int_0^x f(t) \int_t^x (s-t)^{n-1}\, ds\, dt = \frac{1}{n!} \int_0^x (x-t)^n\, f(t)\, dt,$$

which is the formula for $n + 1$. ☐

Replacing n by a real number

Cauchy's formula expresses I^n using only the exponent $n - 1$ and the factorial $(n-1)!$. The key observation is that neither of these requires n to be an integer. If we replace $(n-1)!$ by $\Gamma(n) = (n-1)!$ (valid for positive integers) and then allow n to be any positive real number μ, we obtain a natural candidate for "integration of order μ."

Definition 2.1.2 (Riemann–Liouville fractional integral)**.** Let $\mu > 0$. The *(left-sided) Riemann–Liouville fractional integral of order* μ*, based at* 0, is

$$({}_0I_x^\mu f)(x) := \frac{1}{\Gamma(\mu)} \int_0^x (x-t)^{\mu-1} f(t)\, dt, \qquad x > 0,$$

whenever the integral on the right exists.

When $\mu = n \in \mathbb{N}$, the definition reduces to Cauchy's formula, so the Riemann–Liouville integral genuinely extends n-fold integration to arbitrary positive real orders.

Remark 2.1.3. A central structural property of the Riemann–Liouville integral is the *semigroup law*: for $\mu, \nu > 0$ and suitably integrable f,

$${}_0I_x^\mu \big({}_0I_x^\nu f\big) = {}_0I_x^{\mu+\nu} f.$$

This can be verified by writing both sides as double integrals, exchanging the order of integration, and evaluating the inner integral using the Beta function identity that we shall recall in Section 2.3. The semigroup law says that fractional integration of orders μ and ν composes to give fractional integration of order $\mu + \nu$, just as $I^m \circ I^n = I^{m+n}$ in the integer case.

Remark 2.1.4. Throughout this book we write $J_\alpha := {}_0I_x^\alpha$ when we wish to emphasize the operator-algebraic role of the fractional integral as a *forward shift*. The letter J is chosen for consistency with the notation of the research papers on which the book is based.

2.2 The Caputo derivative and the role of initial data

Having extended integration to non-integer orders, we turn to differentiation. There are several ways to define a "derivative of order α" for non-integer $\alpha > 0$. The two most important are the

Riemann–Liouville derivative and the Caputo derivative. Both will appear in this book, but the Caputo derivative plays the more central role in the canonical shift models.

The Riemann–Liouville derivative

Fix $\alpha > 0$, and let $m := \lceil \alpha \rceil$ (the smallest integer greater than or equal to α). A first idea for a fractional derivative of order α is to integrate to order $m - \alpha$ (which brings us to an integer order m) and then differentiate m times in the ordinary sense:

$$ {}^{\mathrm{RL}}_{0}D^{\alpha}_{x} f \;:=\; \frac{d^m}{dx^m}\big({}_0I^{m-\alpha}_x f\big). $$

This is the *Riemann–Liouville fractional derivative*. It is perfectly well defined on smooth functions, and it satisfies elegant composition laws. However, it has one feature that is inconvenient for initial-value problems: it does not annihilate constants when $0 < \alpha < 1$. Indeed, since

$$ {}_0I^{1-\alpha}_x(1) \;=\; \frac{x^{1-\alpha}}{\Gamma(2-\alpha)}, $$

differentiating once gives

$$ {}^{\mathrm{RL}}_{0}D^{\alpha}_{x}(1) \;=\; \frac{x^{-\alpha}}{\Gamma(1-\alpha)}, $$

which is nonzero and, moreover, not even a polynomial. This observation already hints at the difficulties that will be analyzed carefully in Chapter 3.

The Caputo derivative

The Caputo derivative reverses the order of operations: differentiate first, then integrate.

Definition 2.2.1 (Caputo fractional derivative)**.** Let $\alpha > 0$ with $\alpha \notin \mathbb{N}$, and set $m := \lceil \alpha \rceil$. The *(left-sided) Caputo fractional derivative of order* α*, based at* 0, is

$$ ({}^{\mathrm{C}}_{0}D^{\alpha}_{x} f)(x) \;:=\; \frac{1}{\Gamma(m-\alpha)} \int_0^x (x-t)^{m-\alpha-1} f^{(m)}(t)\, dt, \qquad x > 0, $$

whenever f is m times differentiable and the integral exists. If $\alpha = m \in \mathbb{N}$, we set ${}^{\mathrm{C}}_{0}D^{m}_{x} f := f^{(m)}$.

Equivalently, the Caputo derivative is

$$ {}^{\mathrm{C}}_{0}D^{\alpha}_{x} f \;=\; {}_0I^{m-\alpha}_x\big(f^{(m)}\big). $$

The crucial difference from the Riemann–Liouville definition is that ordinary differentiation acts *first*, so that constant and low-degree polynomial terms in f are killed before the fractional integral is applied.

Proposition 2.2.2. *If* $0 < \alpha < 1$*, then*

$$ {}^{\mathrm{C}}_{0}D^{\alpha}_{x}(1) \;=\; 0. $$

More generally, if $\alpha > 0$*,* $m = \lceil \alpha \rceil$*, and* $0 \le r \le m-1$*, then*

$$ {}_0^{\mathrm{C}}D_x^{\alpha}(x^r) \;=\; 0. $$

Proof. When $0 \le r \le m-1$, the m-th ordinary derivative of x^r is identically zero. Therefore ${}_0^{\mathrm{C}}D_x^{\alpha}(x^r) = {}_0I_x^{m-\alpha}(0) = 0$. □

This property makes the Caputo derivative especially natural for initial-value problems, because initial data of the form $f(0), f'(0), \ldots, f^{(m-1)}(0)$ are invisible to the operator.

Remark 2.2.3. The relationship between the two derivatives is

$$ {}_0^{\mathrm{C}}D_x^{\alpha} f \;=\; {}_0^{\mathrm{RL}}D_x^{\alpha} f \;-\; \sum_{k=0}^{m-1} \frac{f^{(k)}(0)}{\Gamma(k+1-\alpha)}\, x^{k-\alpha}. $$

Thus the Caputo derivative differs from the Riemann–Liouville derivative by a correction term that depends only on the Taylor data of f at the origin. This correction is exactly the "initial-data artifact" that the Caputo definition removes.

The fundamental composition identities

The following two identities describe how the fractional integral and the Caputo derivative interact. They are the precursors of the shift relations that will be established in Chapter 4.

Proposition 2.2.4. *Let* $0 < \alpha < 1$*. For every sufficiently smooth* f *on* $(0, \infty)$*:*

(i) ${}_0^{\mathrm{C}}D_x^{\alpha}\big({}_0I_x^{\alpha} f\big) \;=\; f.$

(ii) ${}_0I_x^{\alpha}\big({}_0^{\mathrm{C}}D_x^{\alpha} f\big)(x) \;=\; f(x) - f(0).$

The proof uses the semigroup law for fractional integrals and the definition of the Caputo derivative; we omit the details, which can be found in any standard reference such as Podlubny [9] or Diethelm [3].

Remark 2.2.5. Identity (i) says that the Caputo derivative is a left inverse of the fractional integral: $C_\alpha J_\alpha = I$ (using the operator notation of Chapter 1). Identity (ii) says that the fractional integral is only a *partial* right inverse: $J_\alpha C_\alpha f = f - f(0)$. The defect $f(0)$ is precisely the component of f in the kernel of C_α—the "vacuum" in the shift language that we will develop. These identities will become $C_\alpha J_\alpha = I$ and $J_\alpha C_\alpha = I - \Pi_0$ on the canonical graded space.

Remark 2.2.6. We write $C_\alpha := {}_0^{\mathrm{C}}D_x^{\alpha}$ when we wish to emphasize the backward-shift role. The letter C stands for both "Caputo" and "contraction" (lowering of the grade).

2.3 Gamma, Beta, and Mittag-Leffler functions

Three special functions appear repeatedly in this book. We collect here only the properties that will be used later.

The Gamma function

The Gamma function is defined for $\operatorname{Re} z > 0$ by

$$ \Gamma(z) \;=\; \int_0^\infty t^{z-1} e^{-t}\, dt. $$

Its most important properties for us are:

(a) *Recurrence:* $\Gamma(z+1) = z\,\Gamma(z)$.

(b) *Normalization:* $\Gamma(1) = 1$, so that $\Gamma(n+1) = n!$ for $n \in \mathbb{N}_0$.

(c) *Positivity:* $\Gamma(x) > 0$ for all real $x > 0$.

(d) *Log-convexity:* $\log\Gamma$ is convex on $(0,\infty)$.

(e) *Meromorphic extension:* Γ extends to a meromorphic function on $\mathbb{C}$ with simple poles at $0, -1, -2, \ldots$ and no zeros.

The Gamma function is the structural reason why the factorials $n!$ that appear in Cauchy's formula can be replaced by $\Gamma(\mu)$ for non-integer μ.

Remark 2.3.1. Throughout this book, the Gamma function appears almost exclusively in the denominator of normalized monomials $x^\beta/\Gamma(\beta+1)$. This normalization is not a cosmetic choice; it is *structurally indispensable*. Without it, the action of the fractional integral on a power function x^β would carry a ratio of Gamma values as a coefficient, and the clean shift relation $J_\alpha e_n = e_{n+1}$ would be replaced by a formula with a nontrivial multiplicative factor at every step. The Gamma normalization absorbs these factors once and for all.

The Beta function

The Beta function is defined for $\operatorname{Re} a > 0$, $\operatorname{Re} b > 0$ by

$$B(a,b) \;=\; \int_0^1 u^{a-1}(1-u)^{b-1}\,du.$$

Its connection to the Gamma function is given by the identity

$$B(a,b) \;=\; \frac{\Gamma(a)\,\Gamma(b)}{\Gamma(a+b)}. \tag{2.1}$$

This identity is the essential computational tool in every proof that involves fractional operators acting on power functions: it converts a convolution integral into a ratio of Gamma values.

Example 2.3.2. Let $\mu > 0$ and $\beta > -1$. The integral

$$\int_0^1 (1-u)^{\mu-1}\,u^\beta\,du \;=\; B(\beta+1,\mu) \;=\; \frac{\Gamma(\beta+1)\,\Gamma(\mu)}{\Gamma(\beta+\mu+1)}$$

will appear every time we compute the action of a Riemann–Liouville integral on a power function. The reader is encouraged to verify this identity at least once by direct computation for a specific pair of values, say $\mu = 1/2$ and $\beta = 1$.

The Mittag-Leffler function

The third special function that we need is the one-parameter Mittag-Leffler function:

$$E_\alpha(z) \;:=\; \sum_{n=0}^{\infty} \frac{z^n}{\Gamma(\alpha n+1)}, \qquad \alpha > 0,\; z \in \mathbb{C}. \tag{2.2}$$

The series converges for all $z \in \mathbb{C}$ (the function is entire), and it reduces to the ordinary exponential when $\alpha = 1$:

$$E_1(z) \;=\; \sum_{n=0}^{\infty} \frac{z^n}{n!} \;=\; e^z.$$

Why does the Mittag-Leffler function matter for this book? Consider the canonical basis vectors $e_n(x) = x^{n\alpha}/\Gamma(n\alpha+1)$ that will be introduced in Chapter 4. The Mittag-Leffler function is precisely the generating series of these vectors:

$$E_\alpha(\lambda\, x^\alpha) \;=\; \sum_{n=0}^{\infty} \lambda^n\, e_n(x).$$

Thus the Mittag-Leffler function is to fractional calculus what the exponential function $e^{\lambda x} = \sum \lambda^n x^n/n!$ is to ordinary calculus. Just as the exponential is an eigenfunction of d/dx with eigenvalue λ, the Mittag-Leffler function is, in a suitable sense, an eigenfunction of the Caputo derivative C_α with eigenvalue λ. We will make this precise in Chapter 4 (and in greater analytic detail in Chapters 8 and 14).

Remark 2.3.3. The Mittag-Leffler series has infinite radius of convergence, but the canonical graded space $\mathcal{G}_\alpha^{\mathrm{alg}}$ that we will construct consists only of *finite* linear combinations of the e_n. The Mittag-Leffler function itself is therefore *not* an element of the algebraic space; it lives in a completion. This tension between algebraic models (finite sums) and analytic models (convergent series) will be resolved in the completion theory of Chapter 8.

Remark 2.3.4. There is also a two-parameter Mittag-Leffler function $E_{\alpha,\beta}(z) = \sum_{n=0}^{\infty} z^n/\Gamma(\alpha n+\beta)$, which reduces to $E_\alpha(z)$ when $\beta = 1$. We will encounter it occasionally, but the one-parameter version suffices for most of the book.

2.4 Boundary versus translation

Before performing our first calculations, it is important to understand a conceptual distinction that shapes the entire book: the difference between *one-sided* and *whole-space* fractional operators.

One-sided operators and the role of the boundary

The Riemann–Liouville integral and the Caputo derivative, as defined above, are *one-sided* operators: they integrate from a fixed base point 0 to the variable x. The base point is a boundary in the following sense: the integration kernel $(x-t)^{\mu-1}$ is supported on $0 \le t \le x$, and the lower limit 0 acts as a wall that the integration path cannot cross.

This one-sidedness has profound algebraic consequences. The Caputo derivative annihilates the constant function $e_0 = 1$, creating a *vacuum*: a vector that is destroyed by the backward operator. The existence of a vacuum is directly tied to the presence of the boundary.

Whole-space operators and translation invariance

On the whole real line $\mathbb{R}$ (or on $\mathbb{R}^d$), there is no distinguished boundary point. The natural fractional operators in this setting are the *Weyl* operators, which integrate over infinite tails:

$$(W_+^\mu f)(x) \;:=\; \frac{1}{\Gamma(\mu)} \int_0^\infty s^{\mu-1}\, f(x-s)\, ds.$$

Notice that the integration variable s runs from 0 to ∞, not from 0 to x. The effect is that the operator commutes with translations: if $\tau_h f(x) = f(x-h)$, then $W_+^{\mu}(\tau_h f) = \tau_h(W_+^{\mu} f)$.

Translation invariance means that the Weyl operator has no preferred base point and hence no boundary. As we shall see in Chapter 6, this leads to a *spectral* picture rather than a shift picture: the exponential characters $e^{\lambda x}$ are eigenvectors, and the operator algebra is diagonal.

The geometric principle

The conceptual lesson is:

> *The geometry of the domain determines the algebra of the operators.*
>
> - *A domain with a boundary produces a shift algebra with a vacuum vector.*
> - *A domain without a boundary produces a spectral (diagonal) algebra with no vacuum.*

When the domain has *both* types of directions—for instance, $(0,\infty)^r \times \mathbb{R}^s$—the natural framework is a *hybrid* algebra that combines shifts in the one-sided directions with diagonal multipliers in the whole-space directions. This is the subject of Chapters 7 through 9.

2.5 First concrete calculations

We now carry out the explicit calculations that will serve as the foundation for the shift model of Chapter 4. The central task is to determine how the fractional integral and the Caputo derivative act on power functions with Gamma-normalized coefficients.

The key lemma: fractional integral of a power function

Lemma 2.5.1. *Let $\mu > 0$ and $\beta > -1$. Then*

$$ {}_0I_x^{\mu}\left(\frac{x^{\beta}}{\Gamma(\beta+1)}\right) = \frac{x^{\beta+\mu}}{\Gamma(\beta+\mu+1)}. $$

Proof. By Definition 2.1.2,

$$ {}_0I_x^{\mu}\left(\frac{x^{\beta}}{\Gamma(\beta+1)}\right) = \frac{1}{\Gamma(\mu)\,\Gamma(\beta+1)} \int_0^x (x-t)^{\mu-1}\, t^{\beta}\, dt. $$

Substitute $t = xu$, so that $dt = x\,du$ and the limits become $0 \le u \le 1$:

$$ \frac{x^{\beta+\mu}}{\Gamma(\mu)\,\Gamma(\beta+1)} \int_0^1 (1-u)^{\mu-1}\, u^{\beta}\, du = \frac{x^{\beta+\mu}}{\Gamma(\mu)\,\Gamma(\beta+1)}\, B(\beta+1,\mu). $$

The Beta–Gamma identity (2.1) gives

$$ B(\beta+1,\mu) = \frac{\Gamma(\beta+1)\,\Gamma(\mu)}{\Gamma(\beta+\mu+1)}, $$

and substituting yields

$$ {}_0I_x^{\mu}\left(\frac{x^{\beta}}{\Gamma(\beta+1)}\right) = \frac{x^{\beta+\mu}}{\Gamma(\beta+\mu+1)}. \qquad \square $$

Notice the remarkable simplicity of the result: the Gamma normalization in the numerator and the Beta–Gamma identity in the denominator cancel each other perfectly, producing the same functional form with the exponent shifted by μ.

The Caputo derivative of a power function

Lemma 2.5.2. *Let* $\mu > 0$ *with* $m := \lceil \mu \rceil$, *and let* $\beta > m - 1$. *Then*

$$ {}_0^{\mathrm{C}}D_x^{\mu}\left(\frac{x^{\beta}}{\Gamma(\beta+1)}\right) = \frac{x^{\beta-\mu}}{\Gamma(\beta+1-\mu)}. $$

In addition, if $0 < \mu < 1$ *and* $0 \le r \le 0$ *(that is,* $r = 0$*), then*

$$ {}_0^{\mathrm{C}}D_x^{\mu}(1) = 0. $$

Proof. Assume first that $\mu \notin \mathbb{N}$. Since $\beta > m-1$, the m-th ordinary derivative of $x^{\beta}/\Gamma(\beta+1)$ is

$$ \frac{d^m}{dt^m}\left(\frac{t^{\beta}}{\Gamma(\beta+1)}\right) = \frac{t^{\beta-m}}{\Gamma(\beta+1-m)}. $$

(Here we use the recurrence $\Gamma(\beta+1) = \beta\,\Gamma(\beta)$ repeatedly.) Therefore

$$ {}_0^{\mathrm{C}}D_x^{\mu}\left(\frac{x^{\beta}}{\Gamma(\beta+1)}\right) = {}_0I_x^{m-\mu}\left(\frac{x^{\beta-m}}{\Gamma(\beta+1-m)}\right). $$

Applying Lemma 2.5.1 with the substitutions $\mu \mapsto m - \mu$ and $\beta \mapsto \beta - m$ gives

$$ \frac{x^{(\beta-m)+(m-\mu)}}{\Gamma((\beta-m)+(m-\mu)+1)} = \frac{x^{\beta-\mu}}{\Gamma(\beta+1-\mu)}. $$

When $\mu = m \in \mathbb{N}$, the Caputo derivative is the ordinary m-th derivative, and the formula follows directly from the product $\beta(\beta-1)\cdots(\beta-m+1) = \Gamma(\beta+1)/\Gamma(\beta+1-m)$.

Finally, if $0 < \mu < 1$, then $m = 1$ and the Caputo derivative involves only the first ordinary derivative. Since $d(1)/dx = 0$, we have ${}_0^{\mathrm{C}}D_x^{\mu}(1) = 0$. □

Previewing the canonical basis

Lemmas 2.5.1 and 2.5.2 immediately suggest the basis that will be the protagonist of Chapter 4. Fix $0 < \alpha < 1$, and define

$$ e_n(x) := \frac{x^{n\alpha}}{\Gamma(n\alpha+1)}, \qquad n = 0, 1, 2, \ldots $$

Proposition 2.5.3 (Preview of the shift relations)**.** *For all* $n \ge 0$,

$$ {}_0I_x^{\alpha}\, e_n(x) = e_{n+1}(x). $$

For $n \ge 1$,

$$ {}_0^{\mathrm{C}}D_x^{\alpha}\, e_n(x) = e_{n-1}(x), $$

and

$$ {}_0^{\mathrm{C}}D_x^{\alpha}\, e_0(x) = 0. $$

Proof. The first identity is Lemma 2.5.1 with $\mu = \alpha$ and $\beta = n\alpha$:

$$ {}_0I_x^{\alpha}\left(\frac{x^{n\alpha}}{\Gamma(n\alpha+1)}\right) = \frac{x^{(n+1)\alpha}}{\Gamma((n+1)\alpha+1)} = e_{n+1}(x). $$

For $n \geq 1$, the exponent $\beta = n\alpha > 0 = m-1$ (since $m = 1$ for $0 < \alpha < 1$), so Lemma 2.5.2 applies:

$$ {}_0^{\mathrm{C}}D_x^{\alpha}\left(\frac{x^{n\alpha}}{\Gamma(n\alpha+1)}\right) = \frac{x^{(n-1)\alpha}}{\Gamma((n-1)\alpha+1)} = e_{n-1}(x). $$

Finally, $e_0(x) = 1$ and ${}_0^{\mathrm{C}}D_x^{\alpha}(1) = 0$ by Lemma 2.5.2. □

Example 2.5.4 (The case $\alpha = 1/2$)**.** Take $\alpha = 1/2$. The first few basis vectors are

$$ e_0(x) = 1, \qquad e_1(x) = \frac{x^{1/2}}{\Gamma(3/2)} = \frac{2\sqrt{x}}{\sqrt{\pi}}, \qquad e_2(x) = \frac{x}{\Gamma(2)} = x, \qquad e_3(x) = \frac{x^{3/2}}{\Gamma(5/2)} = \frac{4x^{3/2}}{3\sqrt{\pi}}. $$

The half-derivative of e_1 is $e_0 = 1$, the half-derivative of $e_2 = x$ is $e_1 = 2\sqrt{x}/\sqrt{\pi}$, and the half-derivative of $e_0 = 1$ is zero. The half-integral moves each vector one step up: $J_{1/2}(1) = e_1$, $J_{1/2}(e_1) = e_2 = x$, and so on. The “ladder” picture is just as clean as in the integer case.

Example 2.5.5 (Fractional integral and derivative of a monomial)**.** Let $\alpha = 1/3$. We compute

$$ {}_0I_x^{1/3}\left(\frac{x^{2/3}}{\Gamma(5/3)}\right) = \frac{x^1}{\Gamma(2)} = x, $$

and

$$ {}_0^{\mathrm{C}}D_x^{1/3}\left(\frac{x^{2/3}}{\Gamma(5/3)}\right) = \frac{x^{1/3}}{\Gamma(4/3)}. $$

In the notation $e_n(x) = x^{n/3}/\Gamma(n/3+1)$, the first identity reads $J_{1/3}\, e_2 = e_3$ and the second reads $C_{1/3}\, e_2 = e_1$. The pattern is unmistakable.

What the calculations reveal

Proposition 2.5.3 and the examples above demonstrate that the Gamma-normalized power functions e_n are the “right” basis for fractional operators of order α: the fractional integral raises the index by one, the Caputo derivative lowers it by one, and the constant function is the vacuum. This is exactly the shift-algebra structure that was described informally in Chapter 1.

The next chapter will show that this optimistic first impression is *not* the end of the story. First, we must understand why the most obvious candidate space—the ordinary polynomial space P_n—fails to support any such model. That failure, analyzed in Chapter 3, is what makes the canonical basis of Chapter 4 both necessary and remarkable.

Chapter 3

Failure on Finite-Dimensional Polynomial Spaces

The calculations of the previous chapter ended on an optimistic note: the Gamma-normalized power functions $e_n(x) = x^{n\alpha}/\Gamma(n\alpha + 1)$ produce a clean shift picture for fractional operators. Before developing that picture fully, however, we must understand why the most obvious candidate space fails. The finite-dimensional polynomial space P_n—the space in which ordinary differentiation is most comfortably at home—cannot support any fractional-calculus model in the sense we need.

The purpose of this chapter is to prove two independent obstruction theorems and to extract from them a clear directive for the rest of the book. The material follows the research paper AD01.

3.1 The space P_n and the idea of an internal model

Fix an integer $n \geq 1$, and consider the space of complex polynomials of degree at most n:

$$P_n \;:=\; \{p(x) \in \mathbb{C}[x] : \deg p \leq n\}.$$

This is a vector space of dimension $n + 1$, with basis $\{1, x, x^2, \ldots, x^n\}$.

Ordinary differentiation preserves P_n: the derivative of a polynomial of degree at most n is a polynomial of degree at most $n - 1$. Therefore the restriction

$$D_n \;:=\; \left.\frac{d}{dx}\right|_{P_n}$$

is a well-defined linear endomorphism of P_n.

It is natural to ask whether fractional differentiation can be organized in the same way. The most optimistic form of this question is:

> *Does there exist a family of operators* $(T_\alpha)_{\alpha \geq 0} \subset \mathrm{End}(P_n)$ *such that* $T_0 = I$, $T_1 = D_n$, *and* $T_{\alpha+\beta} = T_\alpha T_\beta$ *for all* $\alpha, \beta \geq 0$?

Such a family would be an *internal fractional model* for D_n on P_n: the word "internal" indicates that every operator in the family acts on the same finite-dimensional space. If such a model existed, it would provide a purely algebraic, finite-dimensional framework for fractional calculus, with no need for infinite-dimensional function spaces or singular integral kernels.

Definition 3.1.1. An *internal fractional model* for D_n on P_n is a family $(T_\alpha)_{\alpha\ge 0} \subset \operatorname{End}(P_n)$ satisfying:

(i) $T_0 = I$ (the identity on P_n),

(ii) $T_1 = D_n$,

(iii) $T_{\alpha+\beta} = T_\alpha T_\beta$ for all $\alpha, \beta \ge 0$.

No continuity in α is assumed.

Remark 3.1.2. The term "internal model" is used in a purely descriptive sense throughout this chapter: it means a semigroup of endomorphisms acting on a fixed finite-dimensional state space. No connection with the control-theoretic internal model principle is intended.

As we shall prove in this chapter, no such model exists. Two independent obstructions prevent it.

3.2 Why classical fractional operators do not stay inside P_n

The first obstruction is direct: the classical Riemann–Liouville and Caputo fractional derivatives simply do not map P_n into itself.

Recall from Chapter 2 that the Riemann–Liouville and Caputo derivatives of order α act on a monomial x^r by

$$ {}_0^{\mathrm{RL}}D_x^\alpha(x^r) = \frac{\Gamma(r+1)}{\Gamma(r+1-\alpha)}\, x^{r-\alpha}, \tag{3.1} $$

$$ {}_0^{\mathrm{C}}D_x^\alpha(x^r) = \begin{cases} 0, & 0 \le r \le \lceil\alpha\rceil - 1, \\ \dfrac{\Gamma(r+1)}{\Gamma(r+1-\alpha)}\, x^{r-\alpha}, & r \ge \lceil\alpha\rceil. \end{cases} \tag{3.2} $$

The decisive feature is the exponent $r - \alpha$ in the output.

The Riemann–Liouville obstruction

Proposition 3.2.1. *For every $n \ge 1$ and every noninteger $\alpha > 0$,*

$$ {}_0^{\mathrm{RL}}D_x^\alpha(P_n) \not\subset P_n. $$

Proof. Apply (3.1) to the constant polynomial $1 = x^0$:

$$ {}_0^{\mathrm{RL}}D_x^\alpha(1) = \frac{1}{\Gamma(1-\alpha)}\, x^{-\alpha}. $$

Since α is not an integer, $-\alpha$ is not a nonnegative integer, and $x^{-\alpha}$ is not an ordinary polynomial. Therefore ${}_0^{\mathrm{RL}}D_x^\alpha(1) \notin P_n$. □

This is a blunt obstruction: the Riemann–Liouville derivative sends even the simplest polynomial—the constant function—out of P_n.

The Caputo obstruction

The Caputo derivative is more polite: it annihilates constants (and, more generally, polynomials of degree less than $\lceil \alpha \rceil$). But this courtesy is not enough.

Proposition 3.2.2. *Let* $n \geq 1$.

(i) For every noninteger α *with* $0 < \alpha < n$,

$$ {}_0^{\mathrm{C}}D_x^{\alpha}(P_n) \not\subset P_n. $$

(ii) For every $\alpha > n$,

$$ {}_0^{\mathrm{C}}D_x^{\alpha}\big|_{P_n} = 0. $$

Proof. For (i), let $m = \lceil \alpha \rceil$. Since $\alpha < n$, we have $m \leq n$, so the monomial x^n belongs to P_n and satisfies $n \geq m$. By (3.2),

$$ {}_0^{\mathrm{C}}D_x^{\alpha}(x^n) = \frac{\Gamma(n+1)}{\Gamma(n+1-\alpha)}\, x^{n-\alpha}. $$

Since $\alpha \notin \mathbb{N}$, the exponent $n - \alpha$ is not a nonnegative integer, and $x^{n-\alpha} \notin P_n$.

For (ii), if $\alpha > n$, then $m = \lceil \alpha \rceil > n$. Every $p \in P_n$ satisfies $p^{(m)} = 0$, so ${}_0^{\mathrm{C}}D_x^{\alpha} p = {}_0I_x^{m-\alpha}(0) = 0$. ☐

Remark 3.2.3. Proposition 3.2.2 presents a dilemma with no escape. For orders $0 < \alpha < n$ with $\alpha \notin \mathbb{N}$, the Caputo derivative leaves P_n (because it produces non-polynomial powers). For orders $\alpha > n$, it collapses P_n entirely to zero. Neither behavior is compatible with an internal model.

Example 3.2.4. Consider $P_2 = \operatorname{span}\{1, x, x^2\}$. The Caputo derivative of order $\alpha = 1/2$ acts as follows:

$$ \begin{aligned} {}_0^{\mathrm{C}}D_x^{1/2}(1) &= 0, \\ {}_0^{\mathrm{C}}D_x^{1/2}(x) &= \frac{1}{\Gamma(3/2)}\, x^{1/2} = \frac{2}{\sqrt{\pi}}\, x^{1/2}, \\ {}_0^{\mathrm{C}}D_x^{1/2}(x^2) &= \frac{2}{\Gamma(5/2)}\, x^{3/2} = \frac{8}{3\sqrt{\pi}}\, x^{3/2}. \end{aligned} $$

The outputs $x^{1/2}$ and $x^{3/2}$ are not polynomials. The Caputo derivative immediately escapes from P_2.

3.3 Jordan form and nilpotent operators

The second obstruction is purely algebraic and does not depend on any specific formula for fractional derivatives. It requires only the Jordan structure of the differentiation operator on P_n. We pause to recall the relevant facts from linear algebra.

The Jordan canonical form

Let V be a finite-dimensional complex vector space and let $A \in \operatorname{End}(V)$. The Jordan canonical form theorem (which the reader has met in a standard linear algebra course) states that A is similar to a direct sum of Jordan blocks:

$$A \sim J_{s_1}(\lambda_1) \oplus J_{s_2}(\lambda_2) \oplus \cdots \oplus J_{s_m}(\lambda_m),$$

where each $J_s(\lambda)$ is the $s \times s$ upper-triangular matrix with λ on the diagonal and 1 on the superdiagonal. A *nilpotent* operator is one for which all eigenvalues are zero; its Jordan form is a direct sum of blocks $J_s(0)$, which we abbreviate to J_s.

Differentiation on P_n is a single nilpotent block

Proposition 3.3.1. *With respect to the normalized basis $e_k(x) = x^k/k!$ $(k = 0, 1, \dots, n)$, the differentiation operator $D_n = (d/dx)|_{P_n}$ is represented by the $(n+1) \times (n+1)$ nilpotent Jordan block:*

$$[D_n] = J_{n+1} = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 0 & 1 \\ 0 & \cdots & \cdots & 0 & 0 \end{pmatrix}.$$

In particular, $D_n^{n+1} = 0$ and $D_n^n \neq 0$.

Proof. Since $(d/dx)(x^k/k!) = x^{k-1}/(k-1)!$ for $k \geq 1$ and $(d/dx)(1) = 0$, the operator acts by $D_n e_k = e_{k-1}$ for $k \geq 1$ and $D_n e_0 = 0$. The matrix of D_n with respect to the ordered basis $(e_0, e_1, \dots, e_n)$ therefore has the entry 1 in position $(k-1, k)$ for $k = 1, \dots, n$ and 0 elsewhere. This is exactly J_{n+1}.

The nilpotency index is $n+1$ because $D_n^n(x^n/n!) = 1 \neq 0$, while $D_n^{n+1} = 0$ since D_n^{n+1} acts on a space of dimension $n+1$ and every basis vector has been lowered to zero after at most $n+1$ steps. □

Remark 3.3.2. The relation $D_n e_k = e_{k-1}$ (with $D_n e_0 = 0$) is exactly the backward-shift relation that we saw in Chapter 2 for ordinary differentiation. The point of this chapter is that the analogous fractional relation cannot be realized on the same space P_n.

3.4 The nonexistence theorem

We now prove the main structural result: the differentiation operator D_n admits no nontrivial roots, and consequently no internal fractional model can exist on P_n.

Powers of a Jordan block

The key algebraic fact is that raising a single nilpotent Jordan block to a power $q \geq 2$ always produces *multiple* Jordan blocks.

Lemma 3.4.1. *Let $s \geq 2$ and $q \geq 2$. The Jordan canonical form of J_s^q is*

$$J_s^q \sim \bigoplus_{r=1}^{\min\{q,s\}} J_{\ell_r}, \qquad \ell_r := \left\lfloor \frac{s-r}{q} \right\rfloor + 1 \quad (1 \leq r \leq \min\{q, s\}).$$

In particular, J_s^q has at least $\min\{q, s\} \geq 2$ Jordan blocks.

Proof. Let $u_1, \ldots, u_s$ be the standard basis of $\mathbb{C}^s$, with the convention $J_s u_j = u_{j-1}$ (setting $u_0 := 0$). Then $J_s^q u_j = u_{j-q}$ (with $u_k = 0$ for $k \leq 0$).

For each $r = 1, \ldots, \min\{q, s\}$, define the subspace

$$V_r \;:=\; \operatorname{span}\{u_r,\, u_{r+q},\, u_{r+2q},\, \ldots\}.$$

The subspaces $V_1, \ldots, V_{\min\{q,s\}}$ are J_s^q-invariant, pairwise linearly independent, and their direct sum is $\mathbb{C}^s$. On the ordered basis $u_r, u_{r+q}, u_{r+2q}, \ldots$, the restriction $J_s^q|_{V_r}$ acts as a one-step backward shift (sending each vector to its predecessor and annihilating the first), so it is a nilpotent Jordan block of size $\ell_r = \lfloor (s-r)/q \rfloor + 1$.

Since $s \geq 2$, the number of blocks is $\min\{q, s\} \geq 2$. ☐

Example 3.4.2. Let $s = 4$ and $q = 2$. The standard basis is u_1, u_2, u_3, u_4, and J_4^2 acts by $J_4^2 u_j = u_{j-2}$. The invariant subspaces are

$$V_1 = \operatorname{span}\{u_1, u_3\}, \qquad V_2 = \operatorname{span}\{u_2, u_4\}.$$

On V_1, the operator sends $u_3 \mapsto u_1 \mapsto 0$, giving a Jordan block J_2. On V_2, it sends $u_4 \mapsto u_2 \mapsto 0$, giving another J_2. Thus $J_4^2 \sim J_2 \oplus J_2$. The single block J_4 has been replaced by two blocks.

Nonexistence of roots

Theorem 3.4.3. *For every $n \geq 1$ and every integer $q \geq 2$, there is no $B \in \operatorname{End}(P_n)$ such that $B^q = D_n$. Equivalently, D_n has no nontrivial q-th root.*

Proof. Suppose $B^q = D_n$ for some $B \in \operatorname{End}(P_n)$. Every eigenvalue λ of B satisfies $\lambda^q \in \sigma(D_n) = \{0\}$, so $\lambda = 0$. Hence B is nilpotent and similar to a direct sum of nilpotent Jordan blocks:

$$B \;\sim\; J_{s_1} \oplus \cdots \oplus J_{s_m}, \qquad s_1 + \cdots + s_m = n+1.$$

Consequently,

$$D_n \;=\; B^q \;\sim\; J_{s_1}^q \oplus \cdots \oplus J_{s_m}^q.$$

Since $D_n \neq 0$ (for instance, $D_n(x) = 1 \neq 0$), at least one s_i satisfies $s_i \geq 2$. By Lemma 3.4.1, $J_{s_i}^q$ has at least two Jordan blocks. Each remaining $J_{s_j}^q$ contributes at least one block (possibly the 1×1 zero block if $s_j = 1$). Therefore B^q has at least $m + 1 \geq 2$ Jordan blocks in total.

But Proposition 3.3.1 shows that D_n is a *single* Jordan block J_{n+1}. A matrix with at least two Jordan blocks cannot be similar to one with exactly one block. This contradiction proves that no such B exists. ☐

Remark 3.4.4. The obstruction is a matter of *Jordan type*, not merely of nilpotency. Some nilpotent matrices do admit nontrivial roots. For instance, $J_2 \oplus J_2$ is the square of a single J_4 (as shown in Example 3.4.2). What fails here is that D_n is a *single* block; the "fragmentation" effect of Lemma 3.4.1 makes it impossible for any power of another nilpotent matrix to produce exactly one block.

The main nonexistence theorem

Theorem 3.4.5 (Nonexistence of internal fractional models on P_n)**.** *For every $n \geq 1$, the operator D_n admits no internal fractional model on P_n.*

Proof. Suppose $(T_\alpha)_{\alpha\geq 0} \subset \mathrm{End}(P_n)$ is an internal fractional model. For any integer $q \geq 2$, the semigroup law gives

$$(T_{1/q})^q = T_{q\cdot(1/q)} = T_1 = D_n.$$

Thus $T_{1/q}$ is a q-th root of D_n, contradicting Theorem 3.4.3. □

Combining with the analytic obstructions of Section 3.2, we obtain a definitive conclusion:

Corollary 3.4.6. *The differentiation operator D_n on P_n cannot serve as an internal model for classical fractional calculus, either in the Riemann–Liouville or in the Caputo sense. Specifically:*

(i) The classical fractional derivatives do not act as endomorphisms of P_n at noninteger orders.

(ii) Even abstractly, without requiring agreement with any particular analytic formula, no semigroup of endomorphisms on P_n with $T_1 = D_n$ can exist.

3.5 Why the negative result is productive

It would be easy to read Theorem 3.4.5 as a dead end. In fact, it is the opposite: a precise diagnosis that points directly to the remedy.

The theorem tells us *exactly what goes wrong*:

(a) *The grading is wrong.* The space P_n is graded by integer exponents $0, 1, 2, \dots, n$. But the fractional operators produce non-integer exponents $r - \alpha$, which do not belong to this grading. The correct space must be graded by the *α-lattice* $0, \alpha, 2\alpha, 3\alpha, \dots$

(b) *The space is too small.* The space P_n is finite-dimensional, but the α-lattice is infinite. The correct space must be infinite-dimensional: it must accommodate all grades $n\alpha$ for $n \in \mathbb{N}_0$.

(c) *The Jordan type is wrong.* The operator D_n is a single nilpotent Jordan block, which admits no roots. But the shift operator on an infinite-dimensional graded chain *does* admit natural "fractional steps"—moving along the chain by α instead of by 1.

In short, the failure is not in the idea of finding a space where fractional operators act as shifts. The failure is in the choice of P_n as the space. The next chapter will show that the space

$$\mathcal{G}_\alpha^{\mathrm{alg}} = \bigoplus_{n=0}^{\infty} \mathbb{C}\,\frac{x^{n\alpha}}{\Gamma(n\alpha+1)}$$

corrects all three defects simultaneously.

Remark 3.5.1. The passage from a finite-dimensional polynomial space to an infinite-dimensional graded space is a significant conceptual step. The reader should note that it is not merely a technical enlargement. The infinite-dimensional space has a *qualitatively different* operator theory: it supports genuine unilateral shifts, which have no finite-dimensional counterpart. This is exactly the feature that makes the fractional shift model possible.

3.6 Matrix roots and why they are not enough

We close this chapter with a brief discussion of a related approach that a reader with a background in matrix analysis might consider: defining "fractional differentiation" through matrix functions.

Matrix functions and fractional powers

If A is a matrix with no eigenvalue on $(-\infty, 0]$, one can define A^α for any $\alpha > 0$ through the Cauchy integral formula or through the principal logarithm:

$$A^\alpha := e^{\alpha \log A}.$$

More generally, if f is a function defined on the spectrum of A (together with enough derivatives at each eigenvalue to account for Jordan block sizes), then $f(A)$ is well defined; see Higham's monograph [6] for a thorough treatment.

The obstruction at the eigenvalue zero

For the differentiation operator D_n on P_n, the only eigenvalue is 0, and the Jordan block has size $n+1$. The function $f(z) = z^\alpha$ (with $0 < \alpha < 1$) has a singularity at $z = 0$: the function is defined on $\mathbb{C} \setminus (-\infty, 0]$, but it does not extend holomorphically (or even continuously) to any neighborhood of the origin. In the standard theory of primary matrix functions, applying f to D_n would require the values $f(0), f'(0), \ldots, f^{(n)}(0)$—but already $f'(0)$ does not exist.

Proposition 3.6.1. *For every $n \geq 1$ and every $0 < \alpha < 1$, the function $z \mapsto z^\alpha$ cannot be applied to D_n as a primary matrix function.*

Proof. The primary matrix function $f(A)$ at an eigenvalue λ_0 of A with Jordan block of size s requires that f be $(s-1)$ times differentiable at λ_0. Here $\lambda_0 = 0$ and $s = n+1$. But $f(z) = z^\alpha$ is not even once differentiable at $z = 0$ when $0 < \alpha < 1$, since $f'(z) = \alpha z^{\alpha-1} \to \infty$ as $z \to 0$. □

Why matrix-root theory does not help

One might try to circumvent this by using a different definition of matrix roots. For instance, if A is invertible, one can define $A^{1/q}$ via the binomial series or via contour integrals that avoid the origin. But D_n is *not* invertible—it is nilpotent. The nilpotency is inherent: differentiation maps constants to zero, and no redefinition can change this.

Theorem 3.4.3 shows that the failure goes deeper than any particular definition of matrix roots. The result is that no endomorphism of P_n whatsoever—regardless of how it is constructed—can serve as a q-th root of D_n. The obstruction is a property of the Jordan type of D_n itself, not a limitation of any one technique for constructing matrix functions.

Remark 3.6.2. The situation is dramatically different in infinite dimensions. On the canonical graded space $\mathcal{G}_\alpha^{\mathrm{alg}}$, the backward shift C_α is *not* nilpotent: it annihilates only the vacuum vector e_0, and it acts nontrivially on all higher grades. The space is infinite-dimensional, so the shift has infinite rank, and the Jordan-type obstruction disappears entirely. This is the fundamental reason why the correct state space for fractional calculus must be infinite-dimensional.

Looking ahead

This chapter has established that the polynomial space P_n cannot host any internal fractional model. The next chapter begins the positive program: we construct the canonical α-graded space on which the fractional integral and derivative become exact forward and backward shifts, realizing the pattern that the calculations at the end of Chapter 2 already suggested.

Chapter 4

The Canonical One-Variable Shift Model

The previous chapter showed that finite-dimensional polynomial spaces cannot support an internal fractional model. The grading is wrong, the space is too small, and the Jordan type of the differentiation operator forbids roots. This chapter constructs the space that overcomes all three obstructions.

The key idea is simple: replace the integer grading $0, 1, 2, \ldots$ by the α-grading $0, \alpha, 2\alpha, \ldots,$ and normalize each basis vector by the Gamma function. On the resulting infinite-dimensional graded space, the Riemann–Liouville integral of order α becomes an exact forward shift and the Caputo derivative of order α becomes an exact backward shift. The constant function $e_0 = 1$ plays the role of the vacuum vector: it is annihilated by the backward shift.

The construction is not merely one example among many. We shall prove that the canonical basis is essentially forced by the shift relations: any graded monomial chain on which the fractional operators act as shifts must be a scalar multiple of this one.

The material of this chapter is based on the research paper AD02. Throughout, we fix a fractional order $0 < \alpha < 1$ and work on the half-line $(0, \infty)$.

4.1 The canonical basis

We begin by identifying the correct basis. Recall from Chapter 2 that the Gamma-normalized power function $x^\beta/\Gamma(\beta + 1)$ has the remarkable property of absorbing the combinatorial factors that arise when fractional operators act on monomials. This motivates the following definition.

Definition 4.1.1. For $n \in \mathbb{N}_0$, the n-th *canonical basis vector* is

$$e_n(x) := \frac{x^{n\alpha}}{\Gamma(n\alpha + 1)}, \qquad x > 0.$$

The *canonical α-graded monomial space* is the algebraic direct sum

$$\mathcal{G}_\alpha^{\mathrm{alg}} := \bigoplus_{n=0}^{\infty} \mathbb{C}\, e_n.$$

We call $e_0 = 1$ the *vacuum vector*.

Several features of this definition deserve comment.

First, the exponents $0, \alpha, 2\alpha, 3\alpha, \ldots$ form an arithmetic progression with common difference

α. This is the α-grading that replaces the integer grading of ordinary polynomials. Each homogeneous component $\mathbb{C}\, e_n$ is one-dimensional and has "grade" n.

Second, the space $\mathcal{G}_\alpha^{\mathrm{alg}}$ is infinite-dimensional. Every element is, by definition, a *finite* linear combination of the e_n, but there is no upper bound on which e_n may appear. This is the algebraic direct sum (as opposed to a completed space of infinite series); no topology is imposed.

Third, when $\alpha = 1$, the basis vectors reduce to the familiar normalized monomials $e_n(x) = x^n/n!$, and $\mathcal{G}_1^{\mathrm{alg}}$ is the algebraic span of these monomials—the same space on which ordinary differentiation and integration form a classical shift pair.

Remark 4.1.2. One might ask: why must the exponents begin at 0? Could we start with a different initial exponent $\rho > 0$ and consider the shifted chain

$$e_n^{(\rho)}(x) \;:=\; \frac{x^{\rho+n\alpha}}{\Gamma(\rho+n\alpha+1)}, \qquad n \in \mathbb{N}_0\,?$$

The answer is that the forward shift J_α preserves every such chain (it raises the exponent by α regardless of ρ), but the backward shift C_α preserves the chain only when $\rho = 0$. Indeed, if $\rho > 0$, then

$$C_\alpha\, e_0^{(\rho)} \;=\; \frac{x^{\rho-\alpha}}{\Gamma(\rho+1-\alpha)},$$

whose exponent $\rho - \alpha$ is not of the form $\rho + n\alpha$ for any $n \in \mathbb{N}_0$ (since that would require $n = -1$). Therefore the Caputo derivative escapes the chain. The condition $\rho = 0$ is necessary and sufficient for invariance under both operators. This observation, which will be proved as part of the shift theorem below, already shows that the starting exponent is not a free parameter but is forced by the algebra.

4.2 The shift theorem

We now prove the central result of this chapter: on the canonical space, the fractional integral and Caputo derivative realize an exact unilateral shift pair.

Theorem 4.2.1 (Shift theorem). *On $\mathcal{G}_\alpha^{\mathrm{alg}}$, the order-$\alpha$ Riemann–Liouville integral $J_\alpha = {}_0I_x^\alpha$ and the order-α Caputo derivative $C_\alpha = {}_0^{\mathrm{C}}D_x^\alpha$ act as follows:*

(i) Forward shift: *$J_\alpha\, e_n = e_{n+1}$ for every $n \geq 0$.*

(ii) Backward shift with vacuum annihilation: *$C_\alpha\, e_0 = 0$, and $C_\alpha\, e_n = e_{n-1}$ for $n \geq 1$.*

In particular, both J_α and C_α define linear endomorphisms of $\mathcal{G}_\alpha^{\mathrm{alg}}$.

Proof. (i) By Lemma 2.5.1 (Chapter 2) with $\mu = \alpha$ and $\beta = n\alpha$,

$$J_\alpha\, e_n \;=\; {}_0I_x^\alpha\left(\frac{x^{n\alpha}}{\Gamma(n\alpha+1)}\right) \;=\; \frac{x^{(n+1)\alpha}}{\Gamma((n+1)\alpha+1)} \;=\; e_{n+1}.$$

(ii) For $n = 0$: $e_0 = 1$, so $C_\alpha\, e_0 = {}_0^{\mathrm{C}}D_x^\alpha(1) = 0$ (Proposition 2.2.2).

For $n \geq 1$: the exponent $\beta = n\alpha$ satisfies $\beta > 0 = m - 1$ (since $m = \lceil \alpha \rceil = 1$ for $0 < \alpha < 1$). Lemma 2.5.2 gives

$$C_\alpha\, e_n \;=\; {}_0^{\mathrm{C}}D_x^\alpha\left(\frac{x^{n\alpha}}{\Gamma(n\alpha+1)}\right) \;=\; \frac{x^{(n-1)\alpha}}{\Gamma((n-1)\alpha+1)} \;=\; e_{n-1}.$$

Since $\mathcal{G}_\alpha^{\mathrm{alg}}$ consists of finite linear combinations of the e_n, the formulas extend by linearity to all elements, defining endomorphisms of $\mathcal{G}_\alpha^{\mathrm{alg}}$. □

Example 4.2.2. For $\alpha = 1/2$, the first few transitions are:

$$J_{1/2} : \underbrace{1}_{e_0} \xrightarrow{J_{1/2}} \underbrace{\frac{2\sqrt{x}}{\sqrt{\pi}}}_{e_1} \xrightarrow{J_{1/2}} \underbrace{x}_{e_2} \xrightarrow{J_{1/2}} \underbrace{\frac{4x^{3/2}}{3\sqrt{\pi}}}_{e_3} \xrightarrow{J_{1/2}} \cdots$$

and the Caputo half-derivative reverses each arrow, with $C_{1/2}$ annihilating $e_0 = 1$. The "ladder" is exactly as clean as for ordinary differentiation on $x^n/n!$.

The shift theorem gives $\mathcal{G}_\alpha^{\mathrm{alg}}$ the structure of an abstract unilateral shift algebra. To make this connection precise, we record the following.

Corollary 4.2.3. *Let c_{00} denote the vector space of finitely supported complex sequences, with standard basis $\{u_n\}_{n\geq 0}$. Define the abstract forward and backward shifts by $S_+ u_n = u_{n+1}$ and $S_- u_0 = 0$, $S_- u_n = u_{n-1}$ for $n \geq 1$. Then the map*

$$U : \mathcal{G}_\alpha^{\mathrm{alg}} \longrightarrow c_{00}, \qquad U(e_n) = u_n,$$

is a vector-space isomorphism satisfying

$$U\, J_\alpha\, U^{-1} = S_+, \qquad U\, C_\alpha\, U^{-1} = S_-.$$

Proof. Since $\{e_n\}$ and $\{u_n\}$ are bases of their respective spaces, U is a vector-space isomorphism. The intertwining identities follow directly from the shift theorem: $UJ_\alpha U^{-1} u_n = UJ_\alpha e_n = Ue_{n+1} = u_{n+1} = S_+ u_n$, and similarly for C_α. □

This corollary says that (J_α, C_α) on $\mathcal{G}_\alpha^{\mathrm{alg}}$ is *unitarily equivalent* (in the algebraic sense) to the standard unilateral shift pair on c_{00}. The shift model has been realized.

4.3 Vacuum and defect projection

The shift theorem immediately yields the fundamental algebraic relations that govern the interaction between J_α and C_α.

Definition 4.3.1. For $n \in \mathbb{N}_0$, let

$$\Pi_n : \mathcal{G}_\alpha^{\mathrm{alg}} \longrightarrow \mathcal{G}_\alpha^{\mathrm{alg}}$$

denote the projection onto the n-th homogeneous component:

$$\Pi_n\left(\sum_{k=0}^{N} a_k\, e_k\right) := a_n\, e_n.$$

In particular, Π_0 is the *vacuum projection*: it extracts the grade-zero component of a vector.

Theorem 4.3.2 (Shift-algebra relations)**.** *On* $\mathcal{G}_\alpha^{\mathrm{alg}}$*,*

$$\begin{aligned} C_\alpha J_\alpha &= I, && (4.1)\\ J_\alpha C_\alpha &= I - \Pi_0, && (4.2)\\ [C_\alpha, J_\alpha] := C_\alpha J_\alpha - J_\alpha C_\alpha &= \Pi_0. && (4.3)\end{aligned}$$

Proof. It suffices to verify each identity on the basis $\{e_n\}_{n\geq 0}$.

For (4.1): for every $n \geq 0$,

$$C_\alpha J_\alpha e_n = C_\alpha e_{n+1} = e_n,$$

since $n+1 \geq 1$. Hence $C_\alpha J_\alpha = I$.

For (4.2): if $n \geq 1$, then

$$J_\alpha C_\alpha e_n = J_\alpha e_{n-1} = e_n.$$

If $n = 0$, then

$$J_\alpha C_\alpha e_0 = J_\alpha 0 = 0.$$

Thus $J_\alpha C_\alpha$ acts as the identity on grades $n \geq 1$ and kills the vacuum component. This is exactly $I - \Pi_0$.

The commutator identity (4.3) is immediate: $[C_\alpha, J_\alpha] = I - (I - \Pi_0) = \Pi_0$. □

Remark 4.3.3. The asymmetry between (4.1) and (4.2) is the algebraic signature of the boundary. The Caputo derivative is a *left* inverse of the fractional integral, but not a right inverse: $J_\alpha C_\alpha$ fails to be the identity precisely on the vacuum. In the language of Chapter 2, the "defect" $f(0)$ in the composition identity $J_\alpha C_\alpha f = f - f(0)$ is now visible as the projection Π_0.

On a whole-space domain, where there is no boundary, there is no vacuum, and the corresponding operators are both-sided inverses. This is the fundamental difference between the shift picture (one-sided, with boundary) and the spectral picture (whole-space, without boundary) that will be developed in Chapter 6.

Remark 4.3.4. The relations $C_\alpha J_\alpha = I$, $J_\alpha C_\alpha = I - \Pi_0$, $[C_\alpha, J_\alpha] = \Pi_0$ are exactly the algebraic relations of a *unilateral shift pair* (sometimes called a creation–annihilation pair). They are the one-sided, discrete analogue of the canonical commutation relation $[a, a^\dagger] = 1$ familiar from quantum mechanics. The vacuum projection plays the role of the ground-state projector.

4.4 Higher powers and the semigroup picture

The shift theorem tells us how the *basic* operators J_α and C_α act. What about their higher powers? For instance, is the m-th power J_α^m the same as the Riemann–Liouville integral of order $m\alpha$?

Definition 4.4.1. For $m \in \mathbb{N}$, the *tail subspace* of $\mathcal{G}_\alpha^{\mathrm{alg}}$ is

$$\mathcal{G}_\alpha^{(\geq m)} := \bigoplus_{n=m}^{\infty} \mathbb{C}\, e_n.$$

The complementary *low-grade sector* is

$$\mathcal{G}_\alpha^{(<m)} := \bigoplus_{n=0}^{m-1} \mathbb{C}\, e_n,$$

which is a finite-dimensional subspace of dimension m.

Lemma 4.4.2. *For every $m \in \mathbb{N}$ and every $n \in \mathbb{N}_0$,*

$$J_\alpha^m\, e_n = e_{n+m},$$

and

$$C_\alpha^m\, e_n = \begin{cases} e_{n-m}, & n \geq m, \\ 0, & 0 \leq n < m. \end{cases}$$

Proof. Both identities follow by induction on m, using $J_\alpha\, e_k = e_{k+1}$ and $C_\alpha\, e_k = e_{k-1}$ (with $C_\alpha\, e_0 = 0$) at each step. □

Theorem 4.4.3 (Semigroup law for the forward shift)**.** *For every $m \in \mathbb{N}$,*

$$J_\alpha^m = {}_0I_x^{m\alpha}$$

on all of $\mathcal{G}_\alpha^{\mathrm{alg}}$.

Proof. By Lemma 4.4.2, $J_\alpha^m\, e_n = e_{n+m}$. On the other hand, Lemma 2.5.1 with $\mu = m\alpha$ and $\beta = n\alpha$ gives

$${}_0I_x^{m\alpha}\, e_n = \frac{x^{(n+m)\alpha}}{\Gamma((n+m)\alpha + 1)} = e_{n+m}.$$

Since both operators agree on the basis, they agree on $\mathcal{G}_\alpha^{\mathrm{alg}}$. □

Theorem 4.4.4 (Partial semigroup law for the backward shift)**.** *For every $m \in \mathbb{N}$,*

$$C_\alpha^m = {}_0^{\mathrm{C}}D_x^{m\alpha}$$

on the tail subspace $\mathcal{G}_\alpha^{(\geq m)}$.

Proof. Let $n \geq m$. By Lemma 4.4.2, $C_\alpha^m\, e_n = e_{n-m}$. We compare with the Caputo derivative of order $\mu = m\alpha$. Set $p = \lceil m\alpha \rceil$. Since $n \geq m$, the exponent $\beta = n\alpha$ satisfies $\beta = n\alpha \geq m\alpha > p - 1$. Lemma 2.5.2 therefore gives

$${}_0^{\mathrm{C}}D_x^{m\alpha}\, e_n = \frac{x^{(n-m)\alpha}}{\Gamma((n-m)\alpha + 1)} = e_{n-m}.$$

The two operators agree on every e_n with $n \geq m$, hence on all of $\mathcal{G}_\alpha^{(\geq m)}$. □

Remark 4.4.5. Theorems 4.4.3 and 4.4.4 make the semigroup picture completely explicit. The Riemann–Liouville integral has a *global* semigroup law: $J_\alpha^m = {}_0I_x^{m\alpha}$ on the entire space. The Caputo derivative has only a *partial* semigroup law: $C_\alpha^m = {}_0^{\mathrm{C}}D_x^{m\alpha}$ on the tail $\mathcal{G}_\alpha^{(\geq m)}$, but *not* on the low-grade sector $\mathcal{G}_\alpha^{(<m)}$.

The failure is localized in a finite-dimensional defect space of dimension m. This is entirely analogous to the situation in Chapter 3, where the finite-dimensional polynomial space was too rigid. Here the low-grade sector plays the role of a "remnant" of that finite-dimensional rigidity, but it no longer obstructs the theory because the full space is infinite-dimensional.

The following corollary records the higher-order composition identities.

Corollary 4.4.6. *For every $m \in \mathbb{N}$,*

$$C_\alpha^m J_\alpha^m = I,$$

and

$$J_\alpha^m C_\alpha^m = I - \Pi_{<m}, \qquad \text{where} \quad \Pi_{<m} := \sum_{k=0}^{m-1} \Pi_k.$$

Proof. By Lemma 4.4.2, $C_\alpha^m J_\alpha^m e_n = C_\alpha^m e_{n+m} = e_n$ for all $n \geq 0$, so $C_\alpha^m J_\alpha^m = I$.

For the second identity,

$$J_\alpha^m C_\alpha^m e_n = \begin{cases} J_\alpha^m e_{n-m} = e_n, & n \geq m, \\ J_\alpha^m 0 = 0, & 0 \leq n < m. \end{cases}$$

This is the action of $I - \Pi_{<m}$. □

The projection $\Pi_{<m}$ onto the low-grade sector is the higher-order analogue of the vacuum projection Π_0. It measures the "defect" of the semigroup relation at order m.

4.5 Uniqueness of the canonical basis

The construction of the canonical space might appear to be one clever choice among many. In fact, it is essentially forced. This section proves a rigidity theorem: among graded monomial chains, the canonical basis is the unique one (up to a single global scalar) on which the fractional operators act as shifts.

Definition 4.5.1. A *graded monomial chain* is a graded vector space

$$\mathcal{M} = \bigoplus_{n=0}^{\infty} \mathbb{C} f_n,$$

where each f_n is a nonzero monomial: $f_n(x) = c_n x^{\beta_n}$ with $c_n \in \mathbb{C} \setminus \{0\}$ and $\beta_n \geq 0$.

Lemma 4.5.2. *Let $\mathcal{M} = \bigoplus_{n \geq 0} \mathbb{C} f_n$ be a graded monomial chain with $f_n(x) = c_n x^{\beta_n}$. Suppose that $J_\alpha f_n = f_{n+1}$ for every $n \geq 0$. Then:*

(i) $\beta_{n+1} = \beta_n + \alpha$ *for all* $n \geq 0$, *so* $\beta_n = \beta_0 + n\alpha$.

(ii) $c_{n+1} = c_n \cdot \Gamma(\beta_n + 1)/\Gamma(\beta_n + \alpha + 1)$.

Proof. By Lemma 2.5.1,

$$J_\alpha f_n = c_n \, {}_0I_x^\alpha(x^{\beta_n}) = c_n \frac{\Gamma(\beta_n + 1)}{\Gamma(\beta_n + \alpha + 1)} x^{\beta_n + \alpha}.$$

Since $J_\alpha f_n = f_{n+1} = c_{n+1}\, x^{\beta_{n+1}}$ and two nonzero monomials on $(0,\infty)$ agree if and only if both their exponents and coefficients coincide, the result follows. □

Lemma 4.5.3. *Let $\mathcal{M}$ be as above, and suppose that $C_\alpha f_0 = 0$. Then $\beta_0 = 0$.*

Proof. Suppose $\beta_0 > 0$. Since $0 < \alpha < 1$ and $\beta_0 > 0$, Lemma 2.5.2 gives

$$C_\alpha f_0 \;=\; c_0 \frac{x^{\beta_0-\alpha}}{\Gamma(\beta_0+1-\alpha)} \;\neq\; 0$$

on $(0,\infty)$, contradicting $C_\alpha f_0 = 0$. Hence $\beta_0 = 0$. □

Theorem 4.5.4 (Uniqueness of the canonical chain)**.** *Let $\mathcal{M} = \bigoplus_{n\geq 0}\mathbb{C} f_n$ be a graded monomial chain with $f_n(x) = c_n\, x^{\beta_n}$ ($c_n \neq 0, \beta_n \geq 0$). Suppose that*

$$J_\alpha f_n = f_{n+1} \qquad (n \geq 0), \qquad C_\alpha f_0 = 0, \qquad C_\alpha f_n = f_{n-1} \qquad (n \geq 1).$$

Then there exists a unique nonzero scalar $c \in \mathbb{C}$ such that

$$f_n(x) \;=\; c\,\frac{x^{n\alpha}}{\Gamma(n\alpha+1)} \qquad (n \geq 0).$$

In particular, up to a global scalar normalization, $\mathcal{M} = \mathcal{G}_\alpha^{\mathrm{alg}}$.

Proof. By Lemma 4.5.2, $\beta_n = \beta_0 + n\alpha$. By Lemma 4.5.3, $\beta_0 = 0$, so $\beta_n = n\alpha$.

The coefficient recurrence from Lemma 4.5.2 becomes

$$c_{n+1} \;=\; c_n\,\frac{\Gamma(n\alpha+1)}{\Gamma((n+1)\alpha+1)}.$$

Iterating from $n = 0$:

$$c_n \;=\; c_0 \prod_{k=0}^{n-1} \frac{\Gamma(k\alpha+1)}{\Gamma((k+1)\alpha+1)} \;=\; \frac{c_0}{\Gamma(n\alpha+1)},$$

where the last equality follows from the telescoping of the Gamma ratios. Writing $c = c_0$, we obtain $f_n(x) = c\, x^{n\alpha}/\Gamma(n\alpha+1)$. □

Corollary 4.5.5. *Among graded monomial chains with one-dimensional homogeneous components, the canonical chain $\mathcal{G}_\alpha^{\mathrm{alg}}$ is, up to multiplication of the entire basis by a single nonzero scalar, the unique chain on which J_α acts as a forward shift and C_α acts as a backward shift with vacuum annihilation.*

Remark 4.5.6. The uniqueness theorem is important for two reasons. First, it shows that the canonical basis is not an ad hoc construction but the *only* graded monomial basis with the desired shift property. Second, it guarantees that the multi-variable extensions of Chapters 5–7 (which are built by taking products of one-variable canonical chains) inherit the same rigidity. The entire edifice rests on a single, forced choice.

4.6 Generating series and the first appearance of Mittag-Leffler

The canonical basis $\{e_n\}_{n\geq 0}$ is intimately connected with the Mittag-Leffler function.

The generating series

Recall from Chapter 2 that the one-parameter Mittag-Leffler function is

$$E_\alpha(z) \;=\; \sum_{n=0}^{\infty} \frac{z^n}{\Gamma(\alpha n+1)}.$$

If we set $z = \lambda\, x^\alpha$ for a parameter λ, we obtain

$$E_\alpha(\lambda\, x^\alpha) \;=\; \sum_{n=0}^{\infty} \lambda^n\, e_n(x). \tag{4.4}$$

Thus the canonical basis vectors are exactly the coefficient functions of the Mittag-Leffler generating series, and λ^n plays the role of the n-th Fourier coefficient.

Why the connection matters

In ordinary calculus, the exponential function $e^{\lambda x} = \sum_{n=0}^{\infty} \lambda^n\, x^n/n!$ is the generating function of the normalized monomials $x^n/n!$, and it is an eigenfunction of d/dx with eigenvalue λ. The Mittag-Leffler function plays the same role for the Caputo derivative:

Proposition 4.6.1 (Formal eigenvalue relation)**.** *If termwise application of C_α to the series $\sum_{n=0}^{\infty} \lambda^n e_n$ is justified, then*

$$C_\alpha\, E_\alpha(\lambda\, x^\alpha) \;=\; \lambda\, E_\alpha(\lambda\, x^\alpha).$$

Proof. Applying C_α term by term:

$$C_\alpha\left(\sum_{n=0}^{\infty} \lambda^n e_n\right) \;=\; \sum_{n=1}^{\infty} \lambda^n\, e_{n-1} \;=\; \lambda \sum_{n=0}^{\infty} \lambda^n\, e_n \;=\; \lambda\, E_\alpha(\lambda\, x^\alpha).$$

The first step uses $C_\alpha e_0 = 0$ and $C_\alpha e_n = e_{n-1}$ for $n \geq 1$. □

Remark 4.6.2. There is a subtlety here. The space $\mathcal{G}_\alpha^{\mathrm{alg}}$ consists of *finite* linear combinations of the e_n, so the infinite series $E_\alpha(\lambda x^\alpha) = \sum \lambda^n e_n$ is *not* an element of $\mathcal{G}_\alpha^{\mathrm{alg}}$. The formal eigenvalue relation therefore lives outside the algebraic space. To make it rigorous, one must pass to a *completion*—a Banach space in which the series converges.

This is the subject of Chapter 8, where weighted Banach completions of $\mathcal{G}_\alpha^{\mathrm{alg}}$ are constructed. In those completions, under appropriate weight conditions, the Mittag-Leffler series converges to an actual element of the space, the termwise application of C_α is justified, and Proposition 4.6.1 becomes a genuine identity in Banach-space operator theory.

For now, the proposition should be understood as a strong hint that the algebraic shift model is the "skeleton" of a richer analytic theory.

The analogy with Fourier analysis

The identity (4.4) also suggests a transform-theoretic perspective. Just as an ordinary power series $\sum a_n z^n$ encodes a function on a disk, the formal series $\sum a_n e_n$ encodes an element of $\mathcal{G}_\alpha^{\mathrm{alg}}$ through its "coefficients" (a_n). In the completed theory, the map $f \mapsto (a_n)$ from functions to coefficient sequences becomes a concrete isomorphism between a function space and a weighted sequence space. Under this transform:

- the forward shift J_α corresponds to multiplication of the coefficient sequence by a shift operator (or, in the transform-model picture, to multiplication by a spectral variable);
- the backward shift C_α corresponds to a backward quotient operation on the transform side.

This transform model will be developed in detail in Chapter 8 for the continuous theory and in Chapter 12 for the discrete theory. The unified framework is the subject of Chapter 14.

Looking ahead

This chapter has established the one-variable canonical shift model: the space $\mathcal{G}_\alpha^{\mathrm{alg}}$ with its basis $\{e_n\}$ on which J_α is the forward shift, C_α is the backward shift, and e_0 is the vacuum. The model is unique (up to a global scalar), the semigroup relations are explicit, and the Mittag-Leffler function appears as the natural generating series.

The next chapter extends this construction from one variable to several variables. The graded chain $\mathbb{N}_0$ becomes a multi-index lattice $\mathbb{N}_0^r$, the single vacuum vector becomes a family of boundary hyperplanes, and the one-variable commutativity (which was trivial, since there was only one operator in each direction) becomes a genuine theorem: the coordinatewise shifts commute. This multi-variable extension is the content of AD03 and the foundation on which the hybrid and completion theories are built.

Chapter 5

Multi-Variable One-Sided Fractional Calculus

The previous chapter established a canonical shift model for one-variable fractional calculus: a single graded chain, a single forward shift, a single backward shift, and a single vacuum vector. The present chapter extends this picture from one variable to several variables.

The passage is guided by a concrete geometric question. On $(0, \infty)^d$, each coordinate carries its own fractional integral and Caputo derivative. Is there a canonical space on which all of these coordinatewise operators act as shifts—and, crucially, on which they *commute*?

The answer is yes. The graded chain $\mathbb{N}_0$ is replaced by the multi-index lattice $\mathbb{N}_0^d$; the single vacuum vector becomes a family of coordinate-vacuum hyperplanes; and the one-variable uniqueness theorem extends to a multi-variable rigidity result. The material follows the research paper AD03.

5.1 From a chain to a lattice

In one variable, the canonical basis is indexed by $n \in \mathbb{N}_0$, and the forward and backward shifts move the index up or down by 1. Geometrically, the index set $\mathbb{N}_0$ is a half-line—a one-dimensional lattice with a boundary at $n = 0$.

In d variables, the natural index set is the d-dimensional nonnegative integer lattice

$$\mathbb{N}_0^d = \big\{\mathbf{k} = (k_1, \dots, k_d) : k_i \in \mathbb{N}_0\big\}.$$

Each coordinate direction in $\mathbb{N}_0^d$ carries its own shift pair. The forward shift in the j-th direction moves $\mathbf{k}$ to $\mathbf{k} + \mathbf{e}_j$ (where $\mathbf{e}_j$ is the j-th standard basis vector), and the backward shift moves $\mathbf{k}$ to $\mathbf{k} - \mathbf{e}_j$—but only when $k_j \geq 1$. When $k_j = 0$, the backward shift in the j-th direction annihilates the vector: the lattice has a boundary wall at $k_j = 0$.

This boundary structure is richer than in one variable. In one dimension, the boundary is a single point ($n = 0$). In d dimensions, the boundary of $\mathbb{N}_0^d$ consists of d coordinate hyperplanes $\{k_j = 0\}$, their pairwise intersections $\{k_i = 0,\ k_j = 0\}$, and so on, down to the single corner point $\mathbf{0} = (0, \dots, 0)$. This hierarchy of boundary strata will be important when we study defects and, later, boundary-trace words.

Remark 5.1.1. The lattice $\mathbb{N}_0^d$ carries a natural partial order: $\mathbf{k} \leq \mathbf{m}$ if and only if $k_j \leq m_j$ for every j. This partial order governs the boundary-layer decomposition. The region $\{\mathbf{k} : \mathbf{k} \geq \mathbf{m}\}$ is the "interior" relative to the threshold $\mathbf{m}$, and the complement $\{\mathbf{k} : \mathbf{k} \ngeq \mathbf{m}\}$ is the "boundary layer" where at least one coordinate index is too small.

5.2 Partial Riemann–Liouville integrals and Caputo derivatives

Fix a dimension $d \geq 1$ and an order tuple $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_d) \in (0,1)^d$. We work on the open positive orthant $(0,\infty)^d$.

The coordinatewise fractional operators are defined by applying the one-variable definitions of Chapter 2 to a single coordinate while treating the remaining coordinates as parameters.

Definition 5.2.1. For $j \in \{1, \dots, d\}$ and $\mu > 0$:

(i) The *partial Riemann–Liouville integral of order μ in the j-th coordinate* is

$$({}_0I^{\mu}_{x_j} f)(x) \;:=\; \frac{1}{\Gamma(\mu)} \int_0^{x_j} (x_j - t)^{\mu-1}\, f(x_1, \dots, x_{j-1}, t, x_{j+1}, \dots, x_d)\, dt.$$

(ii) The *partial Caputo derivative of order μ in the j-th coordinate* (with $\mu \notin \mathbb{N}$, $p = \lceil \mu \rceil$) is

$$({}^{\mathrm{C}}_{0}D^{\mu}_{x_j} f)(x) \;:=\; \frac{1}{\Gamma(p-\mu)} \int_0^{x_j} (x_j - t)^{p-\mu-1}\, \partial_j^p f(x_1, \dots, x_{j-1}, t, x_{j+1}, \dots, x_d)\, dt.$$

Throughout this chapter, we abbreviate

$$J_j \;:=\; {}_0I^{\alpha_j}_{x_j}, \qquad C_j \;:=\; {}^{\mathrm{C}}_{0}D^{\alpha_j}_{x_j}, \qquad 1 \leq j \leq d.$$

The key computational fact is that partial fractional operators act on separable products by affecting only the relevant coordinate.

Lemma 5.2.2. *Let $j \in \{1, \dots, d\}$, $\mu > 0$, and let $\phi = \phi(\widehat{x_j})$ be a function independent of x_j.*

(i) If $\beta > -1$, then

$${}_0I^{\mu}_{x_j}(\phi\, x_j^{\beta}) \;=\; \phi\, \frac{\Gamma(\beta+1)}{\Gamma(\beta+\mu+1)}\, x_j^{\beta+\mu}.$$

(ii) If $p = \lceil \mu \rceil$ and $\beta > p - 1$, then

$${}^{\mathrm{C}}_{0}D^{\mu}_{x_j}(\phi\, x_j^{\beta}) \;=\; \phi\, \frac{\Gamma(\beta+1)}{\Gamma(\beta+1-\mu)}\, x_j^{\beta-\mu}.$$

(iii) If $0 < \mu < 1$, then ${}^{\mathrm{C}}_{0}D^{\mu}_{x_j}\phi = 0$.

Proof. Since ϕ is independent of x_j, it factors out of the integral in each case. The remaining one-variable integrals are exactly the computations from Lemmas 2.5.1 and 2.5.2 of Chapter 2. For (iii), the Caputo derivative involves $\partial_j \phi = 0$. □

5.3 The canonical multivariable basis

The multivariable canonical basis is constructed by taking products of one-variable canonical basis vectors.

Definition 5.3.1. For $\mathbf{k} = (k_1, \dots, k_d) \in \mathbb{N}_0^d$, define

$$e_{\mathbf{k}}(x) := \prod_{j=1}^{d} \frac{x_j^{k_j \alpha_j}}{\Gamma(k_j \alpha_j + 1)}, \qquad x = (x_1, \dots, x_d) \in (0, \infty)^d.$$

The *canonical multi-graded monomial space* is the algebraic direct sum

$$\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}} := \bigoplus_{\mathbf{k} \in \mathbb{N}_0^d} \mathbb{C}\, e_{\mathbf{k}}.$$

The vector $e_{\mathbf{0}} = 1$ is the *vacuum vector*.

Each basis vector $e_{\mathbf{k}}$ is a product of d one-variable canonical basis functions, one for each coordinate. The index $\mathbf{k}$ specifies the "grade" in each direction: the j-th factor has exponent $k_j \alpha_j$ and Gamma normalization $\Gamma(k_j \alpha_j + 1)$.

Remark 5.3.2. The space $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$ is the algebraic tensor product of d copies of the one-variable canonical spaces:

$$\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}} \cong \mathcal{G}_{\alpha_1}^{\mathrm{alg}} \otimes_{\mathrm{alg}} \cdots \otimes_{\mathrm{alg}} \mathcal{G}_{\alpha_d}^{\mathrm{alg}},$$

where each factor carries its own fractional order α_j. The tensor-product structure is the reason why the coordinatewise operators will commute: operators acting on different tensor factors automatically commute.

Example 5.3.3. For $d = 2$ with $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$, the first few basis vectors are:

$$e_{(0,0)} = 1, \qquad e_{(1,0)} = \frac{x_1^{\alpha_1}}{\Gamma(\alpha_1 + 1)}, \qquad e_{(0,1)} = \frac{x_2^{\alpha_2}}{\Gamma(\alpha_2 + 1)},$$

$$e_{(1,1)} = \frac{x_1^{\alpha_1}}{\Gamma(\alpha_1 + 1)} \cdot \frac{x_2^{\alpha_2}}{\Gamma(\alpha_2 + 1)}, \qquad e_{(2,0)} = \frac{x_1^{2\alpha_1}}{\Gamma(2\alpha_1 + 1)}, \qquad e_{(0,2)} = \frac{x_2^{2\alpha_2}}{\Gamma(2\alpha_2 + 1)}.$$

The lattice $\mathbb{N}_0^2$ is a quarter-plane. The boundary consists of two walls: the "x_1-wall" $\{k_1 = 0\}$ (the left edge) and the "x_2-wall" $\{k_2 = 0\}$ (the bottom edge). Their intersection $\{(0, 0)\}$ is the corner, occupied by the vacuum vector.

5.4 Coordinate shifts and the commuting tuple

We now prove the main theorem of this chapter: the partial fractional operators act as commuting coordinate shifts on the canonical multi-graded space.

Theorem 5.4.1 (Multi-variable shift theorem). *For every* $\mathbf{k} \in \mathbb{N}_0^d$ *and every* $j \in \{1, \dots, d\}$*:*

(i) $J_j\, e_{\mathbf{k}} = e_{\mathbf{k}+\mathbf{e}_j}$.

(ii) $C_j\, e_{\mathbf{k}} = 0$ *if* $k_j = 0$*, and* $C_j\, e_{\mathbf{k}} = e_{\mathbf{k}-\mathbf{e}_j}$ *if* $k_j \geq 1$.

In particular, each J_j *and each* C_j *is a linear endomorphism of* $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$.

Proof. Fix $\mathbf{k}$ and j. Write $e_{\mathbf{k}}(x) = \phi(\widehat{x}_j)\, x_j^{k_j\alpha_j}/\Gamma(k_j\alpha_j + 1)$, where $\phi(\widehat{x}_j) = \prod_{i\neq j} x_i^{k_i\alpha_i}/\Gamma(k_i\alpha_i + 1)$ is independent of x_j.

For (i), Lemma 5.2.2(i) with $\mu = \alpha_j$ and $\beta = k_j\alpha_j$ gives

$$J_j\, e_{\mathbf{k}} \;=\; \phi(\widehat{x}_j)\, \frac{x_j^{(k_j+1)\alpha_j}}{\Gamma((k_j + 1)\alpha_j + 1)} \;=\; e_{\mathbf{k}+\mathbf{e}_j}.$$

For (ii), if $k_j = 0$, then $e_{\mathbf{k}}$ is independent of x_j and Lemma 5.2.2(iii) gives $C_j\, e_{\mathbf{k}} = 0$. If $k_j \geq 1$, then $k_j\alpha_j > 0 = p - 1$ (since $p = 1$ for $0 < \alpha_j < 1$), so Lemma 5.2.2(ii) gives $C_j\, e_{\mathbf{k}} = \phi(\widehat{x}_j)\, x_j^{(k_j-1)\alpha_j}/\Gamma((k_j - 1)\alpha_j + 1) = e_{\mathbf{k}-\mathbf{e}_j}$. □

Theorem 5.4.2 (Commutativity). *For all $i, j \in \{1, \ldots, d\}$ with $i \neq j$,*

$$J_iJ_j = J_jJ_i, \qquad C_iC_j = C_jC_i, \qquad C_iJ_j = J_jC_i.$$

Hence the tuples $(J_1, \ldots, J_d)$ and $(C_1, \ldots, C_d)$ are commuting tuples on $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$.

Proof. It suffices to verify each identity on basis vectors. For the integrals:

$$J_iJ_j\, e_{\mathbf{k}} \;=\; J_i\, e_{\mathbf{k}+\mathbf{e}_j} \;=\; e_{\mathbf{k}+\mathbf{e}_j+\mathbf{e}_i} \;=\; J_j\, e_{\mathbf{k}+\mathbf{e}_i} \;=\; J_jJ_i\, e_{\mathbf{k}}.$$

For the Caputo derivatives, $C_iC_j\, e_{\mathbf{k}}$ equals $e_{\mathbf{k}-\mathbf{e}_i-\mathbf{e}_j}$ when both $k_i \geq 1$ and $k_j \geq 1$, and zero otherwise. This expression is symmetric in i and j, so $C_iC_j = C_jC_i$.

For the mixed identity, if $i \neq j$:

$$C_iJ_j\, e_{\mathbf{k}} \;=\; C_i\, e_{\mathbf{k}+\mathbf{e}_j} \;=\; \begin{cases} e_{\mathbf{k}+\mathbf{e}_j-\mathbf{e}_i}, & k_i \geq 1, \\ 0, & k_i = 0, \end{cases}$$

and the same expression results from $J_jC_i\, e_{\mathbf{k}}$, since J_j affects coordinate j while C_i affects coordinate $i \neq j$. □

The shift-algebra relations also extend coordinatewise.

Definition 5.4.3. For $j \in \{1, \ldots, d\}$, the *j-th coordinate-vacuum projection* is the linear map

$$\Pi_j : \mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}} \longrightarrow \mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$$

defined by

$$\Pi_j\, e_{\mathbf{k}} \;:=\; \begin{cases} e_{\mathbf{k}}, & k_j = 0, \\ 0, & k_j \geq 1. \end{cases}$$

Its range is the *j-th coordinate-vacuum hyperplane*

$$\mathcal{V}_j \;:=\; \operatorname{span}\{e_{\mathbf{k}} : k_j = 0\}.$$

Theorem 5.4.4 (Coordinatewise shift-algebra relations). *For every $j \in \{1, \ldots, d\}$,*

$$C_jJ_j = I, \qquad J_jC_j = I - \Pi_j, \qquad [C_j, J_j] = \Pi_j.$$

Proof. The proof is identical to the one-variable case (Theorem 4.3.2 of Chapter 4), applied coordinatewise. For every $\mathbf{k}$: $C_j J_j\, e_{\mathbf{k}} = C_j\, e_{\mathbf{k}+\mathbf{e}_j} = e_{\mathbf{k}}$, so $C_j J_j = I$. And $J_j C_j\, e_{\mathbf{k}} = e_{\mathbf{k}}$ when $k_j \geq 1$ but $= 0$ when $k_j = 0$, which is $(I - \Pi_j)\, e_{\mathbf{k}}$. □

Remark 5.4.5. The commutator $[C_j, J_j] = \Pi_j$ is nonzero and has infinite rank (for $d \geq 2$, the hyperplane $\{k_j = 0\}$ contains infinitely many lattice points). By contrast, the *cross-commutators* $[C_i, J_j]$ with $i \neq j$ vanish: $C_i J_j = J_j C_i$ by Theorem 5.4.2. Thus the only nonzero commutators are the "diagonal" ones, and each is the projection onto the corresponding coordinate boundary.

5.5 Defect hyperplanes and boundary layers

The multi-variable structure introduces a richer defect geometry than the one-variable theory.

Definition 5.5.1. For $\mathbf{m} = (m_1, \ldots, m_d) \in \mathbb{N}_0^d$, define:

(i) The *tail subspace*: $\mathcal{G}_{\boldsymbol{\alpha}}^{(\geq \mathbf{m})} := \bigoplus_{\mathbf{k} \geq \mathbf{m}} \mathbb{C}\, e_{\mathbf{k}}$.

(ii) The *boundary-layer sector*: $\mathcal{G}_{\boldsymbol{\alpha}}^{(<\mathbf{m})} := \bigoplus_{\mathbf{k} \not\geq \mathbf{m}} \mathbb{C}\, e_{\mathbf{k}}$.

(iii) The multi-shift operators: $J^{\mathbf{m}} := J_1^{m_1} \cdots J_d^{m_d}$ and $C^{\mathbf{m}} := C_1^{m_1} \cdots C_d^{m_d}$.

The operators $J^{\mathbf{m}}$ and $C^{\mathbf{m}}$ are well defined because the coordinate operators commute.

Proposition 5.5.2. *For every $\mathbf{m} \in \mathbb{N}_0^d$ and every $\mathbf{k} \in \mathbb{N}_0^d$,*

$$J^{\mathbf{m}}\, e_{\mathbf{k}} = e_{\mathbf{k}+\mathbf{m}},$$

and

$$C^{\mathbf{m}}\, e_{\mathbf{k}} = \begin{cases} e_{\mathbf{k}-\mathbf{m}}, & \mathbf{k} \geq \mathbf{m}, \\ 0, & \mathbf{k} \not\geq \mathbf{m}. \end{cases}$$

Moreover, $J^{\mathbf{m}} = {}_0I_x^{\mathbf{m}\odot\boldsymbol{\alpha}}$ on all of $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$, and $C^{\mathbf{m}} = {}_0^{\mathrm{C}}D_x^{\mathbf{m}\odot\boldsymbol{\alpha}}$ on the tail $\mathcal{G}_{\boldsymbol{\alpha}}^{(\geq \mathbf{m})}$. Here $\mathbf{m} \odot \boldsymbol{\alpha} = (m_1\alpha_1, \ldots, m_d\alpha_d)$.

Proof. The shift formulas follow by repeated application of Theorem 5.4.1. The identification with the classical operators uses Lemma 5.2.2 applied coordinatewise, exactly as in the proofs of Theorems 4.4.3 and 4.4.4 of Chapter 4. □

Corollary 5.5.3. *For every $\mathbf{m} \in \mathbb{N}_0^d$,*

$$C^{\mathbf{m}} J^{\mathbf{m}} = I, \qquad J^{\mathbf{m}} C^{\mathbf{m}} = I - \Pi_{<\mathbf{m}},$$

where $\Pi_{<\mathbf{m}}$ is the projection onto the boundary-layer sector $\mathcal{G}_{\boldsymbol{\alpha}}^{(<\mathbf{m})}$.

Remark 5.5.4. When $d = 1$, the boundary-layer projection $\Pi_{<m}$ has finite rank m: the defect occupies the first m grades $\{0, 1, \ldots, m-1\}$. For $d \geq 2$, the situation is fundamentally different. The set $\{\mathbf{k} \in \mathbb{N}_0^d : \mathbf{k} \not\geq \mathbf{m}\}$ is the complement of a shifted octant and in general contains infinitely many lattice points. Consequently, $\Pi_{<\mathbf{m}}$ has *infinite* rank.

The one-dimensional vacuum has become a family of boundary walls. The defect is no longer concentrated in a finite-dimensional low-grade sector but is spread along the coordinate hyperplanes. This boundary structure will become even more prominent in Chapters 9 and 15, where ordered trace words record the sequential interaction of a vector with different boundary walls.

Example 5.5.5. Let $d = 2$ and $\mathbf{m} = (1, 1)$. The tail subspace $\mathcal{G}_{\boldsymbol{\alpha}}^{(\geq(1,1))}$ consists of the span of $\{e_{\mathbf{k}} : k_1 \geq 1,\ k_2 \geq 1\}$—the "interior" of the quarter-plane. The boundary-layer sector $\mathcal{G}_{\boldsymbol{\alpha}}^{(<(1,1))}$ is spanned by the basis vectors on the two edges:

$$\{e_{(0,k_2)} : k_2 \geq 0\} \ \cup\ \{e_{(k_1,0)} : k_1 \geq 1\}.$$

This is an infinite-dimensional L-shaped region along the boundary of $\mathbb{N}_0^2$. Compare this with the one-variable case $\mathbf{m} = (1)$, where the defect is just the single vector e_0.

5.6 Relation with broader multidimensional fractional calculus

The results of this chapter should be placed in proper context.

What the theorem says and what it does not say

Theorem 5.4.2 establishes that the coordinatewise fractional operators commute *on the canonical multi-graded space* $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$. It does *not* claim that mixed partial fractional derivatives commute on arbitrary function spaces. In fact, on general function classes, mixed partial Riemann–Liouville or Caputo derivatives can fail to commute; see the examples discussed in the broader literature on multidimensional fractional calculus, for instance in the work of Kostić.

The canonical space $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$ is therefore not a universal commuting domain; it is an *optimal invariant sector* on which commutativity holds exactly. Later chapters will examine what happens when this sector is enlarged. In Chapter 9, we will see that boundary-augmented extensions of the canonical space can break commutativity, and that the breaking is precisely controlled by ordered boundary-trace words.

Uniqueness

The uniqueness theorem for the multi-graded lattice extends the one-variable result of Chapter 4.

Theorem 5.6.1 (Uniqueness of the canonical multi-graded lattice). *Let* $\mathcal{M} = \bigoplus_{\mathbf{k}\in\mathbb{N}_0^d} \mathbb{C}\, f_{\mathbf{k}}$ *be a graded monomial lattice with* $f_{\mathbf{k}}(x) = c_{\mathbf{k}} \prod_{j=1}^{d} x_j^{\beta_j(\mathbf{k})}$ *(*$c_{\mathbf{k}} \neq 0$, $\beta_j(\mathbf{k}) \geq 0$*). Suppose that for every* $\mathbf{k}$ *and every* j,

$$J_j\, f_{\mathbf{k}} = f_{\mathbf{k}+\mathbf{e}_j}, \qquad C_j\, f_{\mathbf{k}} = \begin{cases} 0, & k_j = 0, \\ f_{\mathbf{k}-\mathbf{e}_j}, & k_j \geq 1. \end{cases}$$

Then there exists a unique nonzero $c \in \mathbb{C}$ *such that* $f_{\mathbf{k}} = c\, e_{\mathbf{k}}$ *for all* $\mathbf{k}$.

Proof sketch. The forward-shift assumption forces each exponent function $\beta_j(\mathbf{k})$ to increase by α_j when k_j increases by one, and to remain unchanged when any other coordinate changes. Hence $\beta_j(\mathbf{k}) = \beta_j(\mathbf{0}) + k_j\alpha_j$, and the vacuum annihilation assumption forces $\beta_j(\mathbf{0}) = 0$ for every j. Thus $\beta_j(\mathbf{k}) = k_j\alpha_j$.

For the coefficients, the Gamma ratios from the forward-shift recurrence yield $c_{\mathbf{k}+\mathbf{e}_j} = c_{\mathbf{k}} \cdot \Gamma(k_j\alpha_j + 1)/\Gamma((k_j + 1)\alpha_j + 1)$. Defining $d_{\mathbf{k}} = c_{\mathbf{k}} \prod_j \Gamma(k_j\alpha_j + 1)$, one shows that $d_{\mathbf{k}+\mathbf{e}_j} = d_{\mathbf{k}}$ for every j. Since $\mathbb{N}_0^d$ is connected by additions of standard basis vectors, $d_{\mathbf{k}}$ is constant: $d_{\mathbf{k}} = c$ for all $\mathbf{k}$. Therefore $c_{\mathbf{k}} = c/\prod_j \Gamma(k_j\alpha_j + 1)$, giving $f_{\mathbf{k}} = c\, e_{\mathbf{k}}$. □

The product Mittag-Leffler generating function

As in one variable, the canonical basis is the coefficient lattice of a natural generating function. For $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_d) \in \mathbb{C}^d$,

$$\prod_{j=1}^{d} E_{\alpha_j}(\lambda_j\, x_j^{\alpha_j}) = \sum_{\mathbf{k}\in\mathbb{N}_0^d} \boldsymbol{\lambda}^{\mathbf{k}}\, e_{\mathbf{k}}(x), \qquad \boldsymbol{\lambda}^{\mathbf{k}} := \prod_{j=1}^{d} \lambda_j^{k_j}.$$

Formally applying C_j termwise yields

$$C_j \prod_{i=1}^{d} E_{\alpha_i}(\lambda_i\, x_i^{\alpha_i}) = \lambda_j \prod_{i=1}^{d} E_{\alpha_i}(\lambda_i\, x_i^{\alpha_i}),$$

so the product Mittag-Leffler function is a joint eigenfunction of the commuting backward-shift tuple $(C_1, \ldots, C_d)$ with eigenvalue tuple $(\lambda_1, \ldots, \lambda_d)$. As in one variable, this identity is formal at the algebraic level and becomes rigorous in the Banach completions of Chapter 8.

Looking ahead

This chapter has built the multi-variable one-sided shift model. The canonical lattice $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{alg}}$ carries commuting forward and backward shifts in each coordinate direction, with coordinate-vacuum projections on the boundary walls.

So far, every coordinate has been a “one-sided” coordinate: the domain is $(0, \infty)^d$, and each direction has a boundary at the origin. The next chapter asks what happens when some coordinates live on the whole real line $\mathbb{R}$ instead of $(0, \infty)$. On $\mathbb{R}$ there is no boundary, so the shift picture is no longer natural; instead, the spectral picture—diagonalization by exponential characters—takes over. Chapter 6 develops this whole-space spectral model, and Chapter 7 will then combine the two pictures on mixed domains $(0, \infty)^r \times \mathbb{R}^s$.

Chapter 6

Whole-Space Spectral Models and Elementary Complex Analysis

The previous two chapters developed the shift picture for fractional operators on one-sided domains. On $(0, \infty)^r$, the canonical basis consists of Gamma-normalized monomials, the operators act as coordinate shifts, and the boundary at the origin creates a vacuum.

This chapter asks a different question: what is the natural picture on the *whole* space $\mathbb{R}^d$, where there is no boundary?

The answer is fundamentally different. On $\mathbb{R}^d$ the relevant fractional operators are the Weyl operators, which are translation-invariant. Translation invariance means that the natural basis is not monomial but *exponential*: the exponential characters $e_\Lambda(x) = e^{\langle \Lambda, x \rangle}$ diagonalize the operators. The resulting algebra is a *spectral* (diagonal multiplier) algebra rather than a shift algebra, and it carries no vacuum, no boundary defect, and no projection anomaly.

The material follows the research paper AD04. We will need a small amount of complex analysis—specifically, principal branches of complex powers—which we develop along the way.

6.1 Why whole-space geometry is different

In Chapter 2 we noted that one-sided operators (Riemann–Liouville, Caputo) integrate from a fixed base point 0, while whole-space operators (Weyl type) integrate over an infinite tail without reference to a base point. This distinction has deep algebraic consequences.

On a one-sided domain, the base point is a boundary. The boundary creates a vacuum vector (the constant function $e_0 = 1$, annihilated by the Caputo derivative) and produces the asymmetry $J_\alpha C_\alpha = I - \Pi_0 \neq I$ that we studied in Chapters 4 and 5.

On $\mathbb{R}^d$ there is no distinguished base point. Fractional operators are defined by convolution with a kernel over the whole line, and the resulting operators commute with translations. In linear algebra terms: on $\mathbb{R}^d$, the natural symmetry is the translation group, and the natural basis consists of the simultaneous eigenvectors of that group—the exponential characters.

This is exactly the spectral picture that was previewed in Section 1.1 of Chapter 1. The present chapter makes it precise.

Remark 6.1.1. The contrast can be stated as a geometric principle:

Boundary $\Longrightarrow$ *shift algebra (ladder structure, vacuum, defect).*
No boundary $\Longrightarrow$ *spectral algebra (diagonal structure, no vacuum, no defect).*

This principle will be unified in Chapter 7, where mixed domains $(0, \infty)^r \times \mathbb{R}^s$ carry both structures simultaneously.

6.2 Exponential characters and the minimum complex analysis needed

Exponential characters

Let $d \geq 1$. For $\lambda = (\lambda_1, \dots, \lambda_d) \in \mathbb{C}^d$ and $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, define the *exponential character*

$$e_\lambda(x) \ := \ e^{\langle \lambda, x \rangle}, \qquad \langle \lambda, x \rangle := \lambda_1 x_1 + \cdots + \lambda_d x_d.$$

These are the simultaneous eigenfunctions of the partial derivative operators: $\partial_j \, e_\lambda = \lambda_j \, e_\lambda$ for each j. In fact, a stronger uniqueness holds.

Proposition 6.2.1. *If $f \in C^1(\mathbb{R}^d)$ satisfies $\partial_j f = \lambda_j f$ for every $j = 1, \dots, d$, then $f(x) = c\, e^{\langle \lambda, x \rangle}$ for some constant $c \in \mathbb{C}$.*

Proof. Define $g(x) = e^{-\langle \lambda, x \rangle} f(x)$. Then $\partial_j g = -\lambda_j g + \lambda_j g = 0$ for every j, so all first partial derivatives of g vanish and g is constant on $\mathbb{R}^d$. □

This result justifies calling exponential characters the "canonical spectral basis": they are the only smooth joint eigenfunctions of the derivative tuple $(\partial_1, \dots, \partial_d)$.

The spectral index set

We restrict attention to spectral parameters in the open right half-plane:

$$\mathbb{C}_+ \ := \ \{z \in \mathbb{C} : \operatorname{Re} z > 0\}.$$

For a subset $\Lambda \subset (\mathbb{C}_+)^d$, the *algebraic spectral module* is

$$\mathcal{E}_\Lambda^{\mathrm{alg}} \ := \ \bigoplus_{\lambda \in \Lambda} \mathbb{C}\, e_\lambda.$$

Every element of $\mathcal{E}_\Lambda^{\mathrm{alg}}$ is a finite linear combination of exponential characters, and the family $\{e_\lambda\}_{\lambda \in \Lambda}$ is linearly independent (distinct exponentials are linearly independent, which can be seen by restricting to a line and applying the classical result for one-variable exponentials).

Remark 6.2.2. The restriction $\Lambda \subset (\mathbb{C}_+)^d$ ensures that the Weyl-type integrals converge: the kernel integration over $s \in [0, \infty)$ involves $e^{-\lambda_j s}$, which is absolutely integrable when $\operatorname{Re} \lambda_j > 0$. It also ensures that the principal branch of z^γ (introduced below) is well defined.

Principal branches and complex powers

We need one piece of complex analysis: the definition of z^γ for $z \in \mathbb{C}_+$ and $\gamma \geq 0$.

Definition 6.2.3. For $z \in \mathbb{C}_+$ and $\gamma \geq 0$, define

$$z^\gamma \ := \ e^{\gamma \operatorname{Log} z},$$

where $\operatorname{Log} z = \ln|z| + i \operatorname{Arg} z$ is the principal logarithm, with $\operatorname{Arg} z \in (-\pi, \pi)$.

Since $z \in \mathbb{C}_+$ implies $\operatorname{Arg} z \in (-\pi/2, \pi/2)$, the principal logarithm is holomorphic on $\mathbb{C}_+$ and the power z^γ is well defined and nonzero. The key algebraic property we need is:

Proposition 6.2.4. *For $z \in \mathbb{C}_+$ and $\gamma_1, \gamma_2 \geq 0$,*

$$z^{\gamma_1} \cdot z^{\gamma_2} \;=\; z^{\gamma_1+\gamma_2}.$$

Proof. Both sides equal $e^{(\gamma_1+\gamma_2)\operatorname{Log} z}$, since $e^a e^b = e^{a+b}$ for complex numbers and $\operatorname{Log} z$ is the same in both factors. □

This is the "law of exponents" for the principal branch. It holds without ambiguity on $\mathbb{C}_+$ because the principal logarithm is single-valued there.

Remark 6.2.5. We emphasize that the complex analysis used in this chapter is minimal. We need only the principal branch of the logarithm on the right half-plane, the resulting definition of complex powers, and the law of exponents (Proposition 6.2.4). No contour integration, no residue calculus, and no deep function-theoretic machinery is required. A reader who has seen the complex exponential and the principal logarithm in an introductory complex analysis or advanced calculus course has all the prerequisites.

6.3 Weyl fractional operators

We now define the Weyl-type operators that act on $\mathcal{E}_\Lambda^{\mathrm{alg}}$.

Coordinate Weyl integrals

Definition 6.3.1. Let $j \in \{1, \ldots, d\}$ and let a_j be a locally integrable function on $[0, \infty)$ whose Laplace transform

$$\widehat{a_j}(z) \;:=\; \int_0^\infty a_j(s)\, e^{-zs}\, ds$$

exists for every $z \in \mathbb{C}_+$. The *coordinate Weyl integral with kernel* a_j acts on $\mathcal{E}_\Lambda^{\mathrm{alg}}$ by

$$(I_j^{a_j} u)(x) \;:=\; \int_0^\infty a_j(s)\, u(x - s\,\mathbf{e}_j)\, ds,$$

where $\mathbf{e}_j$ is the j-th standard basis vector in $\mathbb{R}^d$.

The defining feature of this operator is that it integrates over $s \in [0, \infty)$ against $u(x - s\,\mathbf{e}_j)$, *not* from 0 to x_j. There is no reference to a base point.

Lemma 6.3.2. *For every $\lambda \in \Lambda$,*

$$I_j^{a_j}\, e_\lambda \;=\; \widehat{a_j}(\lambda_j)\, e_\lambda.$$

Proof.

$$(I_j^{a_j}\, e_\lambda)(x) \;=\; \int_0^\infty a_j(s)\, e^{\langle \lambda, x - s\mathbf{e}_j\rangle}\, ds \;=\; e^{\langle \lambda, x\rangle} \int_0^\infty a_j(s)\, e^{-\lambda_j s}\, ds \;=\; \widehat{a_j}(\lambda_j)\, e_\lambda(x).$$

The integral converges absolutely because $\operatorname{Re} \lambda_j > 0$. □

The key observation is immediate: the Weyl integral acts on each exponential character by *multiplication by a scalar*—the Laplace transform of the kernel evaluated at the spectral parameter. The operator does not shift the index; it multiplies the eigenvalue.

Coordinate Weyl derivatives

To incorporate derivatives, we compose the Weyl integral with ordinary partial differentiation.

Definition 6.3.3. Let a_j be an admissible kernel as above, and let $m_j \in \mathbb{N}_0$. The *coordinate generalized Weyl derivative* is

$$D_j^{a_j,m_j} := \partial_j^{m_j} \circ I_j^{a_j}.$$

For a d-tuple of kernels $a = (a_1, \ldots, a_d)$ and a multi-index $m = (m_1, \ldots, m_d) \in \mathbb{N}_0^d$, define

$$I_{W,a} := I_1^{a_1} \cdots I_d^{a_d}, \qquad D_{W,a,m} := D_1^{a_1,m_1} \cdots D_d^{a_d,m_d}.$$

Lemma 6.3.4. *For every $\lambda \in \Lambda$,*

$$D_j^{a_j,m_j} e_\lambda = \lambda_j^{m_j} \widehat{a_j}(\lambda_j) e_\lambda.$$

Proof. By Lemma 6.3.2, $I_j^{a_j} e_\lambda = \widehat{a_j}(\lambda_j) e_\lambda$. Applying $\partial_j^{m_j}$ to a scalar multiple of e_λ gives

$$\partial_j^{m_j} \big(\widehat{a_j}(\lambda_j) e_\lambda\big) = \widehat{a_j}(\lambda_j) \lambda_j^{m_j} e_\lambda. \qquad \square$$

6.4 The diagonal action theorem and spectral algebra

The lemmas of the previous section combine to give the main diagonalization result.

Theorem 6.4.1 (Diagonal action of generalized Weyl operators)**.** *Let $a = (a_1, \ldots, a_d)$ be an admissible kernel tuple, and let $m \in \mathbb{N}_0^d$. For every $\lambda \in \Lambda$,*

$$I_{W,a} e_\lambda = \widehat{a}(\lambda) e_\lambda, \qquad D_{W,a,m} e_\lambda = \sigma_{a,m}(\lambda) e_\lambda,$$

where the symbols are

$$\widehat{a}(\lambda) := \prod_{j=1}^{d} \widehat{a_j}(\lambda_j), \qquad \sigma_{a,m}(\lambda) := \prod_{j=1}^{d} \lambda_j^{m_j} \widehat{a_j}(\lambda_j).$$

Proof. By Lemma 6.3.2, each $I_j^{a_j}$ multiplies e_λ by $\widehat{a_j}(\lambda_j)$. Composing all d coordinates gives $I_{W,a} e_\lambda = \prod_j \widehat{a_j}(\lambda_j) e_\lambda = \widehat{a}(\lambda) e_\lambda$. The proof for $D_{W,a,m}$ is identical, using Lemma 6.3.4. $\square$

Corollary 6.4.2. *All operators of the form $I_{W,a}$ and $D_{W,a,m}$ pairwise commute on $\mathcal{E}_\Lambda^{\mathrm{alg}}$.*

Proof. Each such operator acts on each basis vector e_λ by multiplication by a scalar. Operators that are simultaneously diagonal in a fixed basis commute. $\square$

The diagonal form naturally leads to a multiplier algebra.

Definition 6.4.3. For any function $\sigma : \Lambda \to \mathbb{C}$, the *spectral multiplier* M_σ is the linear map on $\mathcal{E}_\Lambda^{\mathrm{alg}}$ defined by

$$M_\sigma\left(\sum_{\lambda \in F} c_\lambda e_\lambda\right) := \sum_{\lambda \in F} c_\lambda \sigma(\lambda) e_\lambda,$$

where $F \subset \Lambda$ is finite.

Theorem 6.4.4 (The spectral multiplier algebra)**.** *The set* $\mathscr{M}(\Lambda) = \{M_\sigma : \sigma\colon \Lambda \to \mathbb{C}\}$ *is a commutative subalgebra of* $\mathrm{End}(\mathcal{E}_\Lambda^{\mathrm{alg}})$*, and the map* $\sigma \mapsto M_\sigma$ *is an algebra isomorphism from the pointwise function algebra on* Λ *onto* $\mathscr{M}(\Lambda)$*. Moreover, every generalized Weyl operator on* $\mathcal{E}_\Lambda^{\mathrm{alg}}$ *belongs to* $\mathscr{M}(\Lambda)$*.*

Proof. For any $\sigma, \tau : \Lambda \to \mathbb{C}$, $M_\sigma M_\tau\, e_\lambda = \sigma(\lambda)\tau(\lambda)\, e_\lambda = M_{\sigma\tau}\, e_\lambda$. Hence $M_\sigma M_\tau = M_{\sigma\tau}$, proving that $\mathscr{M}(\Lambda)$ is a commutative algebra and the map preserves multiplication. The map is injective (if $M_\sigma = 0$, then $\sigma(\lambda) = 0$ for all λ) and surjective by definition. The final statement is Theorem 6.4.1 restated as $I_{W,a} = M_{\widehat{a}}$ and $D_{W,a,m} = M_{\sigma_{a,m}}$. □

Remark 6.4.5. Compare this with the shift algebra of Chapters 4–5. There, the operators moved basis vectors along the lattice—they were off-diagonal. Here, every operator is diagonal: it multiplies each basis vector by a scalar determined by the spectral parameter. The algebra has changed from a shift algebra to a multiplier algebra. This is the algebraic reflection of the geometric change from a domain with boundary to a domain without boundary.

6.5 No vacuum, no boundary defect

In the one-sided theory, the vacuum vector $e_0 = 1$ is annihilated by the Caputo derivative, producing the projection Π_0 and the asymmetry $J_\alpha C_\alpha = I - \Pi_0$. What is the analogue in the spectral picture?

There is none. On $\mathcal{E}_\Lambda^{\mathrm{alg}}$, the Weyl operators satisfy a clean law of exponents with no correction term.

Theorem 6.5.1 (Law of exponents for generalized Weyl operators)**.** *Let* a, b *be admissible kernel tuples, and let* $m, n \in \mathbb{N}_0^d$*. Then*

$$I_{W,a}\, I_{W,b} \;=\; I_{W,a *_0 b}, \qquad D_{W,a,m}\, D_{W,b,n} \;=\; D_{W,a *_0 b, m+n},$$

where $a *_0 b = (a_1 *_0 b_1, \dots, a_d *_0 b_d)$ *denotes coordinatewise one-sided convolution.*

Proof. By Theorem 6.4.4, $I_{W,a} I_{W,b} = M_{\widehat{a}} M_{\widehat{b}} = M_{\widehat{a}\,\widehat{b}}$. The standard Laplace-transform identity gives $\widehat{a_j *_0 b_j}(z) = \widehat{a}_j(z)\,\widehat{b}_j(z)$, so $\widehat{a} \cdot \widehat{b} = \widehat{a *_0 b}$, and therefore $I_{W,a} I_{W,b} = M_{\widehat{a *_0 b}} = I_{W,a *_0 b}$. The proof for $D_{W,a,m} D_{W,b,n}$ is analogous, using $\sigma_{a,m} \cdot \sigma_{b,n} = \sigma_{a *_0 b, m+n}$. □

For the standard Weyl derivative of multi-order $\boldsymbol{\beta} = (\beta_1, \dots, \beta_d) \in [0, \infty)^d$, the diagonalization takes an especially transparent form.

Definition 6.5.2. For $j \in \{1, \dots, d\}$ and $\beta_j \geq 0$, define the *standard Weyl derivative* $W_j^{\beta_j}$ on $\mathcal{E}_\Lambda^{\mathrm{alg}}$ as the generalized Weyl derivative whose kernel is $g_\nu(s) = s^{\nu - 1}/\Gamma(\nu)$ with $\nu = m_j - \beta_j$ and $m_j = \lceil \beta_j \rceil$ (with the convention $W_j^0 = I$ and $W_j^k = \partial_j^k$ for $k \in \mathbb{N}$). The multidimensional standard Weyl derivative is $W^{\boldsymbol{\beta}} = W_1^{\beta_1} \cdots W_d^{\beta_d}$.

The Laplace transform of g_ν is $z^{-\nu}$ (a standard calculation using the Gamma integral), so the symbol of $W_j^{\beta_j}$ on e_λ is $\lambda_j^{m_j} \cdot \lambda_j^{-(m_j - \beta_j)} = \lambda_j^{\beta_j}$.

Theorem 6.5.3 (Spectral form of the standard Weyl derivative)**.** *For every $\boldsymbol{\beta} \in [0, \infty)^d$ and every $\lambda \in \Lambda$,*

$$W^{\boldsymbol{\beta}} e_\lambda \;=\; \lambda^{\boldsymbol{\beta}} e_\lambda, \qquad \lambda^{\boldsymbol{\beta}} := \prod_{j=1}^{d} \lambda_j^{\beta_j}.$$

Consequently, $W^{\boldsymbol{\alpha}} W^{\boldsymbol{\beta}} = W^{\boldsymbol{\alpha}+\boldsymbol{\beta}}$ on $\mathcal{E}_\Lambda^{\mathrm{alg}}$ for all $\boldsymbol{\alpha}, \boldsymbol{\beta} \in [0, \infty)^d$.

Proof. The first identity follows from the symbol computation above. For the law of exponents, on each e_λ: $W^{\boldsymbol{\alpha}} W^{\boldsymbol{\beta}} e_\lambda = \lambda^{\boldsymbol{\alpha}} \lambda^{\boldsymbol{\beta}} e_\lambda = \lambda^{\boldsymbol{\alpha}+\boldsymbol{\beta}} e_\lambda$, using Proposition 6.2.4 coordinatewise. □

Remark 6.5.4. The law of exponents $W^{\boldsymbol{\alpha}} W^{\boldsymbol{\beta}} = W^{\boldsymbol{\alpha}+\boldsymbol{\beta}}$ holds on the *entire* spectral module, with no tail-subspace restriction and no defect projection. Compare this with the one-sided theory, where $C_\alpha^m = {}_0^{\mathrm{C}}D_x^{m\alpha}$ only on the tail $\mathcal{G}_\alpha^{(\geq m)}$ and the complementary low-grade sector carries a nonzero defect.

The absence of a defect is a direct consequence of the absence of a boundary. When there is no vacuum vector, there is no vacuum projection, and the composition identities are exact on both sides.

6.6 Fourier–Laplace viewpoint and elementary spectral language

The diagonal form of the Weyl operators has a natural interpretation in the language of Fourier and Laplace transforms.

Operators as symbol multiplication

Consider the classical picture of constant-coefficient differential equations. On $\mathbb{R}^d$, the Fourier transform converts ∂_j into multiplication by $i\xi_j$, where ξ is the frequency variable. A differential operator $P(\partial) = \sum c_\nu \partial^\nu$ becomes the multiplier $P(i\xi)$, and solving $P(\partial)u = f$ reduces to dividing by the symbol $P(i\xi)$ in Fourier space.

The spectral module $\mathcal{E}_\Lambda^{\mathrm{alg}}$ provides an exact algebraic analogue of this picture for fractional operators. The role of the Fourier variable ξ is played by the spectral parameter λ, and the role of the symbol $P(i\xi)$ is played by the function $\sigma(\lambda)$.

Theorem 6.6.1 (Algebraic solvability of constant-coefficient Weyl equations)**.** *Let $p(z_1, \ldots, z_d)$ be a polynomial in d variables, and let $\boldsymbol{\beta} \in [0, \infty)^d$. The equation*

$$p(W^{\boldsymbol{\beta}})\, u \;=\; f$$

has a unique solution $u \in \mathcal{E}_\Lambda^{\mathrm{alg}}$ for every $f \in \mathcal{E}_\Lambda^{\mathrm{alg}}$ if and only if

$$p(\lambda_1^{\beta_1}, \ldots, \lambda_d^{\beta_d}) \;\neq\; 0 \qquad \text{for all } \lambda \in \Lambda.$$

In that case, $u = M_{\rho^{-1}} f$, where $\rho(\lambda) = p(\lambda_1^{\beta_1}, \ldots, \lambda_d^{\beta_d})$.

Proof. By the law of exponents, each monomial $\prod_j (W_j^{\beta_j})^{\nu_j}$ acts on e_λ by $\prod_j (\lambda_j^{\beta_j})^{\nu_j}$. Hence $p(W^{\boldsymbol{\beta}}) = M_\rho$, and the result follows because M_ρ is invertible on $\mathcal{E}_\Lambda^{\mathrm{alg}}$ if and only if $\rho(\lambda) \neq 0$ for all $\lambda \in \Lambda$. □

Example 6.6.2. Consider the one-dimensional equation $W^{2\alpha}u + u = f$ on $\mathcal{E}_\Lambda^{\mathrm{alg}}$ (a fractional Helmholtz-type equation). The symbol is $\rho(\lambda) = \lambda^{2\alpha} + 1$, which is nonzero for all $\lambda \in \mathbb{C}_+$ (since $\mathrm{Re}(\lambda^{2\alpha}) > 0$). The unique solution is $u = M_{(\lambda^{2\alpha}+1)^{-1}} f$: the solution operator is simply pointwise division by the symbol.

Transform perspective

At the algebraic level, the map

$$\mathcal{F} : \mathcal{E}_\Lambda^{\mathrm{alg}} \longrightarrow \mathbb{C}^\Lambda, \qquad \mathcal{F}\Big(\sum c_\lambda\, e_\lambda\Big) := (c_\lambda)_{\lambda\in\Lambda},$$

sends each element of the spectral module to its coefficient function on Λ. Under $\mathcal{F}$, every Weyl operator M_σ becomes pointwise multiplication by σ on the coefficient side:

$$\mathcal{F}(M_\sigma u)(\lambda) \;=\; \sigma(\lambda)\cdot\mathcal{F}(u)(\lambda).$$

This is the algebraic prototype of the Fourier–Laplace transform model that will be developed analytically in the completion theory of Chapters 8 and 14.

Remark 6.6.3. Compare the solution method of Theorem 6.6.1 with the solution of one-sided equations on the shift model. In the shift model (as will be developed in Chapter 7), solving a constant-coefficient equation involves *both* spectral division in the whole-space variables *and* finite downward recursion in the one-sided grades. In the pure spectral model, there is no recursion—only division. The presence or absence of a boundary determines whether the solution method is purely spectral or requires an additional recursive step.

Looking ahead

This chapter has established the spectral picture for Weyl operators on $\mathbb{R}^d$: exponential characters diagonalize the operators, the algebra is a multiplier algebra, the law of exponents holds without defect, and constant-coefficient equations reduce to pointwise division.

We now have two complementary pictures:

(i) The *shift picture* on $(0,\infty)^r$ (Chapters 4–5): monomial basis, shift operators, vacuum and boundary defects.

(ii) The *spectral picture* on $\mathbb{R}^s$ (this chapter): exponential basis, diagonal multipliers, no vacuum, no defect.

The next chapter combines these two pictures on mixed domains of the form $(0,\infty)^r \times \mathbb{R}^s$, where the first r coordinates carry a boundary and the last s do not. The resulting *hybrid shift-spectral algebra* will be the setting for the completion theory and boundary structure that occupy the second half of the book.

Chapter 7

Hybrid Shift-Spectral Algebra on Mixed Regions

The previous two chapters have each told half of a story. Chapter 5 showed that on the one-sided domain $(0,\infty)^r$, the canonical basis consists of Gamma-normalized monomials, the fractional operators act as coordinate shifts, and the boundary at the origin creates a vacuum and a defect projection. Chapter 6 showed that on the whole space $\mathbb{R}^s$, the natural basis consists of exponential characters, the Weyl operators act as diagonal multipliers, and there is no vacuum and no defect.

The two pictures are not competing descriptions of the same phenomenon. They are complementary descriptions of two genuinely different geometries. This chapter combines them.

The natural setting is the *mixed region*

$$\Omega_{r,s} \;:=\; (0,\infty)^r \times \mathbb{R}^s,$$

in which the first r coordinates are one-sided and the last s coordinates are whole-space. On this domain, we construct a canonical basis whose one-sided component is the Gamma-normalized monomial lattice and whose whole-space component is the exponential character system. The resulting operator algebra is a *hybrid*: shifts in the one-sided directions, diagonal multipliers in the whole-space directions, and complete commutativity between the two blocks. Moreover, every defect phenomenon remains localized in the one-sided block; the spectral block contributes no vacuum and no boundary projection.

The material of this chapter follows the research paper AD05. Throughout, we fix

$$r \ge 1, \qquad s \ge 1, \qquad \boldsymbol{\alpha} = (\alpha_1,\dots,\alpha_r) \in (0,1)^r.$$

7.1 Why mixed domains matter

Many natural problems involve coordinates of both types simultaneously. A time variable $t > 0$ is one-sided: it starts at an initial instant and runs forward, so its fractional calculus involves a base point and creates a boundary. A spatial variable $x \in \mathbb{R}$ is whole-space: it extends in both directions without a distinguished origin, so its fractional calculus is translation-invariant and carries no boundary.

More precisely, consider a domain of the form

$$\Omega_{r,s} = (0,\infty)^r \times \mathbb{R}^s.$$

On the first r coordinates one wants to use Riemann–Liouville integrals and Caputo derivatives, which are one-sided operators anchored at the origin. On the last s coordinates one wants to use Weyl-type operators, which are whole-space convolution operators with no base point.

The algebraic question is then immediate.

> *Is there a single canonical space on which the one-sided operators act as shifts, the whole-space operators act as diagonal multipliers, and the two families commute?*

The answer is yes, and it is obtained by the most natural construction one could hope for: forming the algebraic tensor product of the monomial lattice from Chapter 5 with the exponential module from Chapter 6.

Remark 7.1.1. The mixed domain $\Omega_{r,s}$ is not an artificial construction. In applied settings, r might count the number of "causal" or "evolutionary" variables (time, age, distance from a boundary) while s counts the number of "spatial" or "spectral" variables that extend over all of $\mathbb{R}$. The hybrid algebra captures the mathematical structure that arises whenever one-sided and whole-space behaviors coexist.

7.2 The monomial–exponential hybrid basis

We now define the basis that carries the hybrid algebra. The construction is guided by the principle established in the previous chapters: on one-sided coordinates, use Gamma-normalized monomials; on whole-space coordinates, use exponential characters. The hybrid basis combines these two ingredients by pointwise multiplication.

The spectral set

In Chapter 6, the exponential characters $e_\lambda(x) = e^{\langle \lambda, x\rangle}$ were indexed by a spectral set $\Lambda \subset (\mathbb{C}_+)^s$, where

$$\mathbb{C}_+ \ := \ \{z \in \mathbb{C} : \operatorname{Re} z > 0\}.$$

The restriction to the open right half-plane ensures that the Laplace transforms used in the Weyl calculus converge and that the law of exponents for complex powers holds via the principal branch. We continue to use this convention here.

Definition 7.2.1. A *spectral set* is any nonempty subset $\Lambda \subset (\mathbb{C}_+)^s$. No further structure (topology, countability, or algebraic closure) is assumed.

The hybrid basis

Definition 7.2.2. For each multi-index $\mathbf{k} = (k_1, \ldots, k_r) \in \mathbb{N}_0^r$ and each spectral parameter $\lambda = (\lambda_1, \ldots, \lambda_s) \in \Lambda$, define the *hybrid basis vector*

$$e_{\mathbf{k},\lambda}(x, y) \ := \ \left(\prod_{i=1}^{r} \frac{x_i^{k_i \alpha_i}}{\Gamma(k_i \alpha_i + 1)}\right) e^{\langle \lambda, y\rangle}, \qquad (x, y) \in \Omega_{r,s}.$$

The *canonical hybrid module* is the algebraic direct sum

$$\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda} := \bigoplus_{(\mathbf{k},\lambda)\in\mathbb{N}_0^r\times\Lambda} \mathbb{C}\, e_{\mathbf{k},\lambda}.$$

Every element of $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ is a finite linear combination of the basis vectors $e_{\mathbf{k},\lambda}$. No infinite sums, no convergence questions, and no topology are involved at this stage. (Completions will be introduced in Chapter 8.)

Proposition 7.2.3 (Linear independence)**.** *The family* $\{e_{\mathbf{k},\lambda} : (\mathbf{k},\lambda)\in\mathbb{N}_0^r\times\Lambda\}$ *is linearly independent.*

Proof. Suppose

$$\sum_{\nu=1}^{N}\sum_{\ell=1}^{m_\nu} c_{\nu,\ell}\, e_{\mathbf{k}^{(\nu,\ell)},\,\lambda^{(\nu)}}(x,y) = 0$$

for all $(x,y)\in\Omega_{r,s}$, where $\lambda^{(1)},\dots,\lambda^{(N)}$ are distinct elements of Λ. Grouping by spectral label gives

$$\sum_{\nu=1}^{N} p_\nu(x)\, e^{\langle\lambda^{(\nu)},y\rangle} = 0,$$

where each $p_\nu(x)$ is a finite linear combination of distinct monomials $\prod_i x_i^{k_i\alpha_i}$. Fix any $x\in(0,\infty)^r$. Since the exponentials $e^{\langle\lambda^{(\nu)},y\rangle}$ with distinct $\lambda^{(\nu)}$ are linearly independent as functions of $y\in\mathbb{R}^s$, it follows that $p_\nu(x) = 0$ for every ν. But the monomials with distinct exponent tuples are linearly independent on $(0,\infty)^r$, so all coefficients $c_{\nu,\ell}$ vanish. □

The tensor-product viewpoint

The hybrid basis has a transparent algebraic meaning. Let

$$\mathcal{G}^{\mathrm{alg}}_{\boldsymbol{\alpha}} := \bigoplus_{\mathbf{k}\in\mathbb{N}_0^r} \mathbb{C} \prod_{i=1}^{r} \frac{x_i^{k_i\alpha_i}}{\Gamma(k_i\alpha_i+1)}$$

denote the canonical multi-graded monomial lattice of Chapter 5, and let

$$\mathcal{E}^{\mathrm{alg}}_{\Lambda} := \bigoplus_{\lambda\in\Lambda} \mathbb{C}\, e^{\langle\lambda,y\rangle}$$

denote the exponential spectral module of Chapter 6. Then the map

$$e_{\mathbf{k},\lambda} \longleftrightarrow \left(\prod_{i=1}^{r} \frac{x_i^{k_i\alpha_i}}{\Gamma(k_i\alpha_i+1)}\right)\otimes e^{\langle\lambda,y\rangle}$$

gives a vector-space isomorphism

$$\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda} \cong \mathcal{G}^{\mathrm{alg}}_{\boldsymbol{\alpha}}\otimes\mathcal{E}^{\mathrm{alg}}_{\Lambda}.$$

This is, at the level of vector spaces, simply the algebraic tensor product of the two modules built in Chapters 5 and 6. The content of the present chapter is that the *operators* also decompose in the expected way: shifts on the first factor, multipliers on the second, and full commutativity between the two.

7.3 The hybrid shift-spectral theorem

We now prove the main structural result. The argument has three parts: the one-sided operators act as shifts on the **k**-index, the whole-space operators act diagonally on λ, and the two families commute.

One-sided operators

For each $i \in \{1, \dots, r\}$, define

$$J_i \;:=\; {}_0I^{\alpha_i}_{x_i}, \qquad C_i \;:=\; {}^{\mathrm{C}}_{0}D^{\alpha_i}_{x_i},$$

the partial Riemann–Liouville integral and partial Caputo derivative of order α_i in the variable x_i, as introduced in Chapter 5.

Theorem 7.3.1 (Shift action in the one-sided block)**.** *For every* $(\mathbf{k}, \lambda) \in \mathbb{N}_0^r \times \Lambda$ *and every* $i \in \{1, \dots, r\}$,

$$J_i\, e_{\mathbf{k},\lambda} = e_{\mathbf{k}+\mathbf{e}_i,\lambda},$$

and

$$C_i\, e_{\mathbf{k},\lambda} = \begin{cases} 0, & k_i = 0, \\ e_{\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i \geq 1. \end{cases}$$

In particular, each J_i and each C_i defines a linear endomorphism of $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$.

Proof. Fix $(\mathbf{k}, \lambda)$ and i, and write

$$e_{\mathbf{k},\lambda}(x, y) = \phi(\widehat{x}_i, y)\, \frac{x_i^{k_i\alpha_i}}{\Gamma(k_i\alpha_i + 1)},$$

where

$$\phi(\widehat{x}_i, y) := \left(\prod_{\substack{1 \leq \ell \leq r \\ \ell \neq i}} \frac{x_\ell^{k_\ell\alpha_\ell}}{\Gamma(k_\ell\alpha_\ell + 1)} \right) e^{\langle \lambda, y \rangle}.$$

This factor is independent of x_i. Applying the classical monomial formula for the Riemann–Liouville integral with order α_i and exponent $\beta = k_i\alpha_i$ gives

$$J_i\, e_{\mathbf{k},\lambda} = \phi(\widehat{x}_i, y)\, \frac{x_i^{(k_i+1)\alpha_i}}{\Gamma((k_i + 1)\alpha_i + 1)} = e_{\mathbf{k}+\mathbf{e}_i,\lambda}.$$

For the Caputo derivative, if $k_i = 0$ then $e_{\mathbf{k},\lambda} = \phi(\widehat{x}_i, y)$ is independent of x_i, and the Caputo derivative of a function independent of its differentiation variable vanishes. If $k_i \geq 1$, then $k_i\alpha_i > 0$, and the classical monomial formula for the Caputo derivative gives

$$C_i\, e_{\mathbf{k},\lambda} = \phi(\widehat{x}_i, y)\, \frac{x_i^{(k_i-1)\alpha_i}}{\Gamma((k_i - 1)\alpha_i + 1)} = e_{\mathbf{k}-\mathbf{e}_i,\lambda}. \qquad \square$$

The shift relations in Theorem 7.3.1 are exactly the multi-variable shift relations of Chapter 5, now carrying a passive spectral label λ. For each fixed λ, the spectral fiber

$$\mathcal{H}_{\lambda}^{\mathrm{alg}} := \bigoplus_{\mathbf{k}\in\mathbb{N}_0^r} \mathbb{C}\, e_{\mathbf{k},\lambda}$$

is invariant under every J_i and every C_i, and the action on $\mathcal{H}_{\lambda}^{\mathrm{alg}}$ is precisely the commuting shift algebra of Chapter 5, multiplied by the fixed character $e^{\langle\lambda,y\rangle}$.

Corollary 7.3.2 (Shift commutation). *For all $i, j \in \{1, \dots, r\}$ with $i \neq j$,*

$$J_iJ_j = J_jJ_i, \qquad C_iC_j = C_jC_i, \qquad C_iJ_j = J_jC_i$$

on $\mathcal{H}_{\alpha,\Lambda}^{\mathrm{alg}}$.

Proof. Each identity follows by evaluating both sides on the basis vector $e_{\mathbf{k},\lambda}$ and checking that the results agree. For instance,

$$J_iJ_j\, e_{\mathbf{k},\lambda} = J_i\, e_{\mathbf{k}+\mathbf{e}_j,\lambda} = e_{\mathbf{k}+\mathbf{e}_j+\mathbf{e}_i,\lambda} = J_jJ_i\, e_{\mathbf{k},\lambda},$$

and the mixed identity $C_iJ_j = J_jC_i$ holds because J_j moves only the j-th index while C_i moves only the i-th index. □

Vacuum projections and shift relations

Definition 7.3.3. For each $i \in \{1, \dots, r\}$, define the *coordinate-vacuum projection* $\Pi_i : \mathcal{H}_{\alpha,\Lambda}^{\mathrm{alg}} \to \mathcal{H}_{\alpha,\Lambda}^{\mathrm{alg}}$ by

$$\Pi_i\, e_{\mathbf{k},\lambda} := \begin{cases} e_{\mathbf{k},\lambda}, & k_i = 0, \\ 0, & k_i \geq 1. \end{cases}$$

Theorem 7.3.4 (Fundamental shift relations). *For every $i \in \{1, \dots, r\}$,*

$$C_iJ_i = I, \qquad J_iC_i = I - \Pi_i, \qquad [C_i, J_i] = \Pi_i.$$

Proof. On a basis vector $e_{\mathbf{k},\lambda}$,

$$C_iJ_i\, e_{\mathbf{k},\lambda} = C_i\, e_{\mathbf{k}+\mathbf{e}_i,\lambda} = e_{\mathbf{k},\lambda},$$

so $C_iJ_i = I$. For the reverse composition: if $k_i \geq 1$, then

$$J_iC_i\, e_{\mathbf{k},\lambda} = J_i\, e_{\mathbf{k}-\mathbf{e}_i,\lambda} = e_{\mathbf{k},\lambda} = (I - \Pi_i)\, e_{\mathbf{k},\lambda},$$

while if $k_i = 0$, then $C_i\, e_{\mathbf{k},\lambda} = 0$, so $J_iC_i\, e_{\mathbf{k},\lambda} = 0 = (I - \Pi_i)\, e_{\mathbf{k},\lambda}$. The commutator identity $[C_i, J_i] = C_iJ_i - J_iC_i = \Pi_i$ follows immediately. □

These are exactly the shift-pair relations from Chapters 4–5, now acting on the larger hybrid module. The key observation is that the spectral label λ plays no role in these identities: it is carried passively throughout.

Whole-space operators

We now turn to the Weyl variables. Recall from Chapter 6 that a *generalized Weyl integral* in the j-th whole-space coordinate is defined by

$$(I_{W,j}^{a_j}u)(x,y) \;:=\; \int_0^\infty a_j(t)\,u(x,\,y-t\,\varepsilon_j)\,dt,$$

where $a_j \in L^1_{\mathrm{loc}}([0,\infty))$ has a Laplace transform

$$\widehat{a_j}(z) = \int_0^\infty a_j(t)\,e^{-zt}\,dt, \qquad z \in \mathbb{C}_+,$$

and ε_j is the j-th standard basis vector of $\mathbb{R}^s$. More generally, a *generalized Weyl derivative* of order $m_j \in \mathbb{N}_0$ is

$$D_{W,j}^{a_j,m_j}u \;:=\; \partial_{y_j}^{m_j}(I_{W,j}^{a_j}u).$$

The coordinatewise operators compose as $I_{W,a} := I_{W,1}^{a_1}\cdots I_{W,s}^{a_s}$ and $D_{W,a,m} := D_{W,1}^{a_1,m_1}\cdots D_{W,s}^{a_s,m_s}$.

Theorem 7.3.5 (Spectral action in the whole-space block). *Let $a = (a_1,\ldots,a_s)$ be an admissible kernel tuple and let $m = (m_1,\ldots,m_s) \in \mathbb{N}_0^s$. Then for every $(\mathbf{k},\lambda) \in \mathbb{N}_0^r \times \Lambda$,*

$$I_{W,a}\,e_{\mathbf{k},\lambda} = \widehat{a}(\lambda)\,e_{\mathbf{k},\lambda}, \qquad \widehat{a}(\lambda) := \prod_{j=1}^{s}\widehat{a_j}(\lambda_j),$$

and

$$D_{W,a,m}\,e_{\mathbf{k},\lambda} = \sigma_{a,m}(\lambda)\,e_{\mathbf{k},\lambda}, \qquad \sigma_{a,m}(\lambda) := \prod_{j=1}^{s}\lambda_j^{m_j}\,\widehat{a_j}(\lambda_j).$$

In particular, the whole-space operators leave the one-sided index $\mathbf{k}$ unchanged and act purely by scalar multiplication on the spectral label.

Proof. Write $e_{\mathbf{k},\lambda}(x,y) = \psi_{\mathbf{k}}(x)\,e^{\langle\lambda,y\rangle}$, where $\psi_{\mathbf{k}}(x) = \prod_{i=1}^r x_i^{k_i\alpha_i}/\Gamma(k_i\alpha_i+1)$. Since $\psi_{\mathbf{k}}(x)$ is independent of y, the calculation reduces to the one-variable computation already performed in Chapter 6. For each j,

$$I_{W,j}^{a_j}\,e_{\mathbf{k},\lambda} = \psi_{\mathbf{k}}(x)\int_0^\infty a_j(t)\,e^{\langle\lambda,\,y-t\varepsilon_j\rangle}\,dt = \psi_{\mathbf{k}}(x)\,e^{\langle\lambda,y\rangle}\int_0^\infty a_j(t)\,e^{-\lambda_j t}\,dt = \widehat{a_j}(\lambda_j)\,e_{\mathbf{k},\lambda}.$$

Composing in all s coordinates multiplies the eigenvalues, giving the formula for $I_{W,a}$. Differentiating m_j times in y_j produces the additional factor $\lambda_j^{m_j}$. □

Commutativity between the two blocks

Corollary 7.3.6 (Mixed commutativity). *Every one-sided operator commutes with every whole-space operator on $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$. Precisely: for all $i \in \{1,\ldots,r\}$, all admissible kernel tuples a, and all $m \in \mathbb{N}_0^s$,*

$$J_i\,I_{W,a} = I_{W,a}\,J_i, \qquad C_i\,I_{W,a} = I_{W,a}\,C_i,$$

and similarly with $D_{W,a,m}$ in place of $I_{W,a}$.

Proof. Each identity is checked on the basis. For example,

$$J_i\, I_{W,a}\, e_{\mathbf{k},\lambda} = J_i\big(\widehat{a}(\lambda)\, e_{\mathbf{k},\lambda}\big) = \widehat{a}(\lambda)\, e_{\mathbf{k}+\mathbf{e}_i,\lambda} = I_{W,a}\, J_i\, e_{\mathbf{k},\lambda}.$$

The remaining identities are proved identically: the Weyl operators change only the scalar coefficient (which depends on λ), while the shift operators change only the lattice index $\mathbf{k}$. ☐

Spectral multipliers and the multiplier algebra

The diagonal action of the Weyl operators motivates the following general definition.

Definition 7.3.7. Let $\sigma : \Lambda \to \mathbb{C}$ be any function. The *spectral multiplier* $M_\sigma : \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda} \to \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ is defined by

$$M_\sigma\left(\sum_{(\mathbf{k},\lambda)\in F} c_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}\right) := \sum_{(\mathbf{k},\lambda)\in F} c_{\mathbf{k},\lambda}\, \sigma(\lambda)\, e_{\mathbf{k},\lambda},$$

where $F \subset \mathbb{N}_0^r \times \Lambda$ is finite. Denote by $\mathscr{M}(\Lambda) := \{M_\sigma : \sigma : \Lambda \to \mathbb{C}\}$ the algebra of all spectral multipliers.

Theorem 7.3.8 (Multiplier algebra)**.**

(i) *$\mathscr{M}(\Lambda)$ is a commutative subalgebra of* $\mathrm{End}(\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda})$*, and the map $\sigma \mapsto M_\sigma$ is an algebra isomorphism from the algebra of scalar functions on Λ (pointwise operations) onto $\mathscr{M}(\Lambda)$.*

(ii) *Every multiplier commutes with every J_i and every C_i.*

(iii) *The Weyl operators are multipliers: $I_{W,a} = M_{\widehat{a}}$ and $D_{W,a,m} = M_{\sigma_{a,m}}$.*

Proof. For part (i): on a basis vector, $M_\sigma M_\tau\, e_{\mathbf{k},\lambda} = \sigma(\lambda)\tau(\lambda)\, e_{\mathbf{k},\lambda} = M_{\sigma\tau}\, e_{\mathbf{k},\lambda}$, so $\mathscr{M}(\Lambda)$ is closed under multiplication and the map is a homomorphism. Injectivity follows from $M_\sigma\, e_{\mathbf{k},\lambda} = \sigma(\lambda)\, e_{\mathbf{k},\lambda}$.

For part (ii): $J_i\, M_\sigma\, e_{\mathbf{k},\lambda} = \sigma(\lambda)\, e_{\mathbf{k}+\mathbf{e}_i,\lambda} = M_\sigma\, J_i\, e_{\mathbf{k},\lambda}$, and similarly for C_i.

Part (iii) is simply a restatement of Theorem 7.3.5 in the language of multipliers. ☐

Standard Weyl derivatives

The most important special case arises from the fractional power kernels that define standard Weyl derivatives.

Definition 7.3.9. For $j \in \{1, \ldots, s\}$ and $\beta \geq 0$, the *standard Weyl derivative of order β* in the j-th coordinate is

$$W_j^\beta := \begin{cases} I, & \beta = 0, \\ \partial_{y_j}^\beta, & \beta \in \mathbb{N}, \\ D_{W,j}^{g_{m-\beta},\, m}, & \beta \notin \mathbb{N},\ m := \lceil \beta \rceil, \end{cases}$$

where $g_\nu(t) := t^{\nu-1}/\Gamma(\nu)$ for $t > 0$. For a multi-order $\boldsymbol{\beta} = (\beta_1, \dots, \beta_s) \in [0,\infty)^s$, define $W^{\boldsymbol{\beta}} := W_1^{\beta_1} \cdots W_s^{\beta_s}$.

Proposition 7.3.10 (Eigenvalue formula for Weyl derivatives). *For every $\boldsymbol{\beta} \in [0,\infty)^s$ and every $(\mathbf{k},\lambda) \in \mathbb{N}_0^r \times \Lambda$,*

$$W^{\boldsymbol{\beta}}\, e_{\mathbf{k},\lambda} = \lambda^{\boldsymbol{\beta}}\, e_{\mathbf{k},\lambda}, \qquad \lambda^{\boldsymbol{\beta}} := \prod_{j=1}^{s} \lambda_j^{\beta_j}.$$

Thus $W^{\boldsymbol{\beta}} = M_{\lambda\mapsto\lambda^{\boldsymbol{\beta}}}$ is a spectral multiplier, and the law of exponents holds:

$$W^{\boldsymbol{\beta}}\, W^{\boldsymbol{\gamma}} = W^{\boldsymbol{\beta}+\boldsymbol{\gamma}} = W^{\boldsymbol{\gamma}}\, W^{\boldsymbol{\beta}} \qquad \textit{for all } \boldsymbol{\beta}, \boldsymbol{\gamma} \in [0,\infty)^s.$$

Proof. For the case $\beta_j \notin \mathbb{N}$, write $m_j = \lceil \beta_j \rceil$. The Laplace transform of $g_{m_j-\beta_j}$ is $\widehat{g_{m_j-\beta_j}}(\lambda_j) = \lambda_j^{-(m_j-\beta_j)}$ (a standard identity recalled in Chapter 6). Hence

$$W_j^{\beta_j}\, e_{\mathbf{k},\lambda} = \lambda_j^{m_j}\, \lambda_j^{-(m_j-\beta_j)}\, e_{\mathbf{k},\lambda} = \lambda_j^{\beta_j}\, e_{\mathbf{k},\lambda}.$$

The integer cases $\beta_j = 0$ and $\beta_j \in \mathbb{N}$ are immediate from the definition. Composing over all j and using the law of exponents $\lambda_j^{\beta_j}\lambda_j^{\gamma_j} = \lambda_j^{\beta_j+\gamma_j}$ for the principal branch on $\mathbb{C}_+$ gives the semigroup property. □

The abstract hybrid algebra

The results above show that $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ supports exactly two types of operator: shifts in the one-sided directions and multipliers in the spectral directions. The following theorem makes this precise by exhibiting a coordinate-free algebraic model.

Theorem 7.3.11 (Transport to the standard model). *Let $c_{00}(\mathbb{N}_0^r \times \Lambda)$ denote the vector space of finitely supported functions on $\mathbb{N}_0^r \times \Lambda$, with standard basis $\{u_{\mathbf{k},\lambda}\}$. Define coordinate shifts S_i^+, S_i^- and diagonal multipliers $\widetilde{M}_\sigma$ on this space in the obvious way:*

$$\begin{aligned} S_i^+\, u_{\mathbf{k},\lambda} &:= u_{\mathbf{k}+\mathbf{e}_i,\lambda}, \\ S_i^-\, u_{\mathbf{k},\lambda} &:= \begin{cases} 0, & k_i = 0, \\ u_{\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i \geq 1, \end{cases} \\ \widetilde{M}_\sigma\, u_{\mathbf{k},\lambda} &:= \sigma(\lambda)\, u_{\mathbf{k},\lambda}. \end{aligned}$$

Then the map $U : \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda} \to c_{00}(\mathbb{N}_0^r \times \Lambda)$ defined by $U(e_{\mathbf{k},\lambda}) := u_{\mathbf{k},\lambda}$ is a vector space isomorphism satisfying

$$UJ_iU^{-1} = S_i^+, \qquad UC_iU^{-1} = S_i^-, \qquad UM_\sigma U^{-1} = \widetilde{M}_\sigma.$$

Proof. Since both families of vectors form bases, U is an isomorphism. The intertwining relations follow from the basis-level actions computed in Theorems 7.3.1 and 7.3.5. □

Remark 7.3.12 (Extreme cases). The hybrid algebra interpolates continuously between the two earlier models.

(i) If the spectral set is a single point, $\Lambda = \{\lambda_0\}$, the spectral block is one-dimensional and $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ collapses to the shift algebra of Chapter 5 multiplied by the fixed character $e^{\langle \lambda_0, y\rangle}$.

(ii) If the one-sided dimension is formally set to $r = 0$ (no one-sided variables), the shift block disappears and $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ reduces to the spectral module of Chapter 6.

Example 7.3.13. Consider the simplest mixed case: $r = s = 1$, with $\alpha \in (0,1)$ and $\Lambda = \{\lambda\} \subset \mathbb{C}_+$ a single spectral parameter. The domain is $(0,\infty) \times \mathbb{R}$. The hybrid basis is

$$e_{k,\lambda}(x,y) = \frac{x^{k\alpha}}{\Gamma(k\alpha+1)}\, e^{\lambda y}, \qquad k \in \mathbb{N}_0.$$

The forward shift $J = {}_0I^{\alpha}_x$ moves k to $k+1$, the backward shift $C = {}^{\mathrm{C}}_0D^{\alpha}_x$ moves k to $k-1$ (and annihilates $k = 0$), and every Weyl derivative $W^{\beta} = D^{g_{m-\beta},m}_{W,y}$ simply multiplies by λ^{β}. The algebra is a one-dimensional unilateral shift tensored with a one-dimensional multiplier—the absolute simplest version of the hybrid picture.

7.4 Defect localization

One of the most important structural conclusions of the hybrid theory is that all defect phenomena are localized in the one-sided block. The whole-space block creates no vacuum, no boundary layer, and no projection anomaly. This section makes the statement precise.

Higher powers and boundary-layer projections

Definition 7.4.1. For $\mathbf{m} = (m_1,\dots,m_r) \in \mathbb{N}_0^r$, define

$$J^{\mathbf{m}} := J_1^{m_1}\cdots J_r^{m_r}, \qquad C^{\mathbf{m}} := C_1^{m_1}\cdots C_r^{m_r}.$$

These are well defined because the coordinate operators commute (Corollary 7.3.2).

Definition 7.4.2. The *boundary-layer projection* $\Pi_{<\mathbf{m}}$ is defined by

$$\Pi_{<\mathbf{m}}\, e_{\mathbf{k},\lambda} := \begin{cases} e_{\mathbf{k},\lambda}, & \mathbf{k} \not\geq \mathbf{m}, \\ 0, & \mathbf{k} \geq \mathbf{m}, \end{cases}$$

where $\mathbf{k} \geq \mathbf{m}$ means $k_i \geq m_i$ for every i. The range of $\Pi_{<\mathbf{m}}$ is the *boundary-layer sector*: the span of all basis vectors whose one-sided multi-index has at least one coordinate below the threshold.

Theorem 7.4.3 (Higher-order shift relations)**.** *For every* $\mathbf{m} \in \mathbb{N}_0^r$*,*

$$C^{\mathbf{m}} J^{\mathbf{m}} = I \qquad \text{on all of } \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda},$$

and

$$J^{\mathbf{m}} C^{\mathbf{m}} = I - \Pi_{<\mathbf{m}}.$$

Proof. By repeated application of Theorem 7.3.1, $J^{\mathbf{m}} e_{\mathbf{k},\lambda} = e_{\mathbf{k}+\mathbf{m},\lambda}$, and

$$C^{\mathbf{m}} e_{\mathbf{k},\lambda} = \begin{cases} e_{\mathbf{k}-\mathbf{m},\lambda}, & \mathbf{k} \geq \mathbf{m}, \\ 0, & \mathbf{k} \ngeq \mathbf{m}. \end{cases}$$

Hence $C^{\mathbf{m}} J^{\mathbf{m}} e_{\mathbf{k},\lambda} = C^{\mathbf{m}} e_{\mathbf{k}+\mathbf{m},\lambda} = e_{\mathbf{k},\lambda}$, and

$$J^{\mathbf{m}} C^{\mathbf{m}} e_{\mathbf{k},\lambda} = \begin{cases} e_{\mathbf{k},\lambda}, & \mathbf{k} \geq \mathbf{m}, \\ 0, & \mathbf{k} \ngeq \mathbf{m}, \end{cases} = (I - \Pi_{<\mathbf{m}})\, e_{\mathbf{k},\lambda}. \qquad \square$$

The identity $C^{\mathbf{m}} J^{\mathbf{m}} = I$ says that integrating and then differentiating always recovers the original function. The identity $J^{\mathbf{m}} C^{\mathbf{m}} = I - \Pi_{<\mathbf{m}}$ says that differentiating and then integrating recovers the function only up to the boundary-layer component—the "initial conditions" that were lost in differentiation. This is the higher-order hybrid version of the fundamental defect identity from Chapter 4.

The spectral block contributes no defect

Theorem 7.4.4 (Defect localization)**.** *For every* $\mathbf{m} \in \mathbb{N}_0^r$ *and every scalar function* $\sigma : \Lambda \to \mathbb{C}$,

$$\Pi_{<\mathbf{m}}\, M_\sigma = M_\sigma\, \Pi_{<\mathbf{m}}.$$

Consequently, $[C^{\mathbf{m}} J^{\mathbf{m}}, M_\sigma] = 0$ *and* $[J^{\mathbf{m}} C^{\mathbf{m}}, M_\sigma] = 0$.

Proof. On a basis vector,

$$\Pi_{<\mathbf{m}}\, M_\sigma\, e_{\mathbf{k},\lambda} = \Pi_{<\mathbf{m}}\big(\sigma(\lambda)\, e_{\mathbf{k},\lambda}\big) = \begin{cases} \sigma(\lambda)\, e_{\mathbf{k},\lambda}, & \mathbf{k} \ngeq \mathbf{m}, \\ 0, & \mathbf{k} \geq \mathbf{m}, \end{cases}$$

which is exactly $M_\sigma\, \Pi_{<\mathbf{m}}\, e_{\mathbf{k},\lambda}$. The commutator statements follow from the higher-order shift relations of Theorem 7.4.3. □

The theorem says that the boundary-layer projection—the projection onto the subspace where initial data live—does not interact with the spectral multipliers at all. The defect is entirely a property of the one-sided grading. Changing the spectral label λ changes the "frequency content" in the whole-space directions, but it has no effect on whether a vector lies in the boundary layer.

Remark 7.4.5. If Λ is infinite, the boundary-layer sector $\Pi_{<\mathbf{m}} \mathcal{H}^{\mathrm{alg}}_{\alpha,\Lambda}$ has infinite rank, because the sector is indexed by all pairs $(\mathbf{k}, \lambda)$ with $\mathbf{k} \ngeq \mathbf{m}$ and $\lambda \in \Lambda$. But the "source" of the defect remains finite-dimensional in spirit: it comes from the finitely many lattice positions $\mathbf{k} \ngeq \mathbf{m}$ in $\mathbb{N}_0^r$, each replicated once for every $\lambda \in \Lambda$. The spectral block merely contributes multiplicity, not mechanism.

Uniqueness of the hybrid basis

The canonical hybrid basis is not merely one convenient choice among many. Under a natural structural assumption, it is the only one.

Theorem 7.4.6 (Uniqueness)**.** *Let* $\{f_{\mathbf{k},\lambda}\}$ *be a family of functions on* $\Omega_{r,s}$*, indexed by* $(\mathbf{k},\lambda) \in \mathbb{N}_0^r \times \Lambda$*, of the form*

$$f_{\mathbf{k},\lambda}(x,y) = c_{\mathbf{k},\lambda}\left(\prod_{i=1}^{r} x_i^{\beta_i(\mathbf{k})}\right)e^{\langle\lambda,y\rangle}$$

with $c_{\mathbf{k},\lambda} \neq 0$ *and* $\beta_i(\mathbf{k}) \geq 0$ *for all* i*. Suppose that for every* $(\mathbf{k},\lambda)$ *and every* i*,*

$$J_i\, f_{\mathbf{k},\lambda} = f_{\mathbf{k}+\mathbf{e}_i,\lambda}, \qquad C_i\, f_{\mathbf{k},\lambda} = \begin{cases} 0, & k_i = 0, \\ f_{\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i \geq 1. \end{cases}$$

Then for every $\lambda \in \Lambda$ *there exists a unique nonzero scalar* c_λ *such that*

$$f_{\mathbf{k},\lambda} = c_\lambda\, e_{\mathbf{k},\lambda}$$

for all $\mathbf{k} \in \mathbb{N}_0^r$*.*

Proof. Fix $\lambda \in \Lambda$. The character $e^{\langle\lambda,y\rangle}$ factors out and is everywhere nonzero, so the shift relations become relations among the monomial families $g_{\mathbf{k}}(x) = c_{\mathbf{k},\lambda}\prod_i x_i^{\beta_i(\mathbf{k})}$. The forward shift condition $J_i\, g_{\mathbf{k}} = g_{\mathbf{k}+\mathbf{e}_i}$ forces $\beta_i(\mathbf{k}+\mathbf{e}_i) = \beta_i(\mathbf{k}) + \alpha_i$ and gives a coefficient recurrence. The vacuum annihilation $C_i\, g_{\mathbf{k}} = 0$ when $k_i = 0$ forces $\beta_i(\mathbf{k}) = 0$ whenever $k_i = 0$, for otherwise the Caputo formula would yield a nonzero result. Combining the forward recurrence with the initial condition yields $\beta_i(\mathbf{k}) = k_i\alpha_i$ for all i.

The coefficient recurrence then simplifies: defining $d_{\mathbf{k},\lambda} = c_{\mathbf{k},\lambda}\prod_i \Gamma(k_i\alpha_i + 1)$, one finds $d_{\mathbf{k}+\mathbf{e}_i,\lambda} = d_{\mathbf{k},\lambda}$ for every i. Since $\mathbb{N}_0^r$ is connected by successive additions of standard basis vectors, $d_{\mathbf{k},\lambda}$ is constant, equal to some $c_\lambda \neq 0$. Therefore $c_{\mathbf{k},\lambda} = c_\lambda / \prod_i \Gamma(k_i\alpha_i + 1)$, which gives $f_{\mathbf{k},\lambda} = c_\lambda\, e_{\mathbf{k},\lambda}$. ☐

The only freedom left is an independent scalar on each spectral fiber. This is structurally unavoidable, because the shift operators do not mix different values of λ and the multipliers do not connect distinct fibers.

7.5 Constant-coefficient mixed equations

The hybrid algebra provides a particularly transparent algebraic solution method for constant-coefficient fractional equations that involve both one-sided and whole-space operators. The idea is simple: the spectral block is handled by division, the shift block by finite downward recursion, and the two mechanisms do not interfere.

Mixed lower-triangular operators

Definition 7.5.1. Let $\mathbf{M} \in \mathbb{N}_0^r$, and let $\{\sigma_{\mathbf{m}} : \Lambda \to \mathbb{C}\}$ be a finite family of scalar functions, indexed by $\mathbf{0} \leq \mathbf{m} \leq \mathbf{M}$. A *mixed lower-triangular operator* is

$$T := \sum_{\mathbf{0}\leq\mathbf{m}\leq\mathbf{M}} C^{\mathbf{m}}\, M_{\sigma_{\mathbf{m}}}.$$

The term $\mathbf{m} = \mathbf{0}$ contributes a pure multiplier $M_{\sigma_{\mathbf{0}}}$; the terms with $\mathbf{m} > \mathbf{0}$ are lower-triangular in the sense that each $C^{\mathbf{m}}$ lowers the total grade $|\mathbf{k}| = k_1 + \cdots + k_r$.

The key observation is that the "lower-triangular part" is locally nilpotent.

Lemma 7.5.2 (Local nilpotence)**.** *Let* $N := \sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} C^{\mathbf{m}}\, M_{\tau_{\mathbf{m}}}$ *for arbitrary scalar functions* $\tau_{\mathbf{m}}$*. Then* N *is locally nilpotent on* $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$*: for every* $u \in \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ *there exists* $q \in \mathbb{N}$ *such that* $N^q u = 0$.

Proof. Let $u = \sum_{(\mathbf{k},\lambda)\in F} c_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$ with F finite, and set $L = \max\{|\mathbf{k}| : (\mathbf{k},\lambda) \in F\}$. Every nonzero summand of N lowers the total grade by at least 1. Hence every nonzero term in $N^q u$ has total grade at most $L - q$. For $q > L$ this is impossible, so $N^q u = 0$. □

Theorem 7.5.3 (Invertibility of mixed operators)**.** *Let* T *be as in Definition 7.5.1, and assume that*

$$\sigma_{\mathbf{0}}(\lambda) \neq 0 \qquad \text{for all } \lambda \in \Lambda.$$

Then T *is invertible on* $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$*. More precisely, writing* $N := \sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} C^{\mathbf{m}}\, M_{\sigma_{\mathbf{m}}}$ *so that* $T = M_{\sigma_{\mathbf{0}}} + N$*, one has*

$$T^{-1} = \sum_{q=0}^{\infty} \left(-M_{\sigma_{\mathbf{0}}^{-1}}\, N\right)^q M_{\sigma_{\mathbf{0}}^{-1}},$$

where the series is finite on every vector.

Proof. Since $\sigma_{\mathbf{0}}$ is nowhere zero on Λ, the multiplier $M_{\sigma_{\mathbf{0}}}$ is invertible with inverse $M_{\sigma_{\mathbf{0}}^{-1}}$. By Theorem 7.3.8(ii), this inverse commutes with every shift operator, so

$$T = M_{\sigma_{\mathbf{0}}}\left(I + M_{\sigma_{\mathbf{0}}^{-1}} N\right).$$

By Lemma 7.5.2, the operator $K := M_{\sigma_{\mathbf{0}}^{-1}} N$ is locally nilpotent. Therefore, on every vector, the Neumann series $(I + K)^{-1} = \sum_{q=0}^{\infty} (-K)^q$ terminates after finitely many steps. Multiplying by $M_{\sigma_{\mathbf{0}}^{-1}}$ gives the stated formula. □

Example 7.5.4. Consider the equation $Tu = f$ on $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ with $r = s = 1$ and

$$T = W^{\beta} + a\, C\, W^{\gamma},$$

where $\beta, \gamma \ge 0$, $a \in \mathbb{C}$, and C is the one-sided Caputo derivative. Here $W^{\beta} = M_{\lambda\mapsto\lambda^{\beta}}$, so the operator is

$$T = M_{\sigma_0} + C\, M_{\sigma_1},$$

with $\sigma_0(\lambda) = \lambda^{\beta}$ and $\sigma_1(\lambda) = a\lambda^{\gamma}$. If $\lambda^{\beta} \neq 0$ for all $\lambda \in \Lambda$ (which holds whenever $\Lambda \subset \mathbb{C}_{+}$), then T is invertible by Theorem 7.5.3. On the basis, the inverse acts by first dividing by λ^{β} (spectral step) and then applying a finite recursion in k (shift step). The spectral step handles the whole-space behavior; the recursive step handles the one-sided behavior. The two steps do not interfere.

Remark 7.5.5. The invertibility theorem reveals the algebraic anatomy of mixed constant-coefficient fractional equations. The leading term $M_{\sigma_{\mathbf{0}}}$ carries the "spectral content" of the equation: it can be inverted by pointwise division, exactly as in the pure whole-space setting of Chapter 6. The lower-order terms involving $C^{\mathbf{m}}$ carry the "one-sided content": they produce finite downward recursion in the lattice grades, exactly as in the pure shift setting of Chapter 5. The hybrid algebra cleanly separates these two mechanisms.

7.6 Why AD05 is a bridge paper

This chapter occupies a special position in the overall story. Let us pause to summarize what has been achieved and where the theory goes next.

What the chapter has established

The results of this chapter may be summarized in a single structural statement:

> *On the mixed domain $(0,\infty)^r \times \mathbb{R}^s$, the canonical hybrid module $\mathcal{H}^{\mathrm{alg}}_{\alpha,\Lambda}$ carries an exact operator algebra that is the tensor product of a commuting unilateral shift algebra in the one-sided directions and a diagonal multiplier algebra in the whole-space directions. The higher-order defects are entirely localized in the one-sided block. The basis is unique under a natural monomial–character ansatz. Mixed constant-coefficient equations are solved by combining spectral division with locally nilpotent downward recursion.*

This confirms the geometric principle previewed at the end of Chapter 6:

Geometry	*Natural operators*	*Algebraic picture*
Boundary (one-sided)	Caputo, Riemann–Liouville	Shift algebra
No boundary (whole-space)	Weyl	Spectral algebra
Mixed	Both	Hybrid

What is still missing

The entire chapter has been algebraic. Every element of $\mathcal{H}^{\mathrm{alg}}_{\alpha,\Lambda}$ is a finite linear combination of basis vectors. This is enough to reveal the operator-algebraic structure, but it is not enough for analysis.

First, important generating vectors—such as the Mittag–Leffler function, which generalizes the exponential—require infinite series. In the algebraic direct sum, such vectors do not exist.

Second, the operators J_i, C_i, and M_σ are endomorphisms of $\mathcal{H}^{\mathrm{alg}}_{\alpha,\Lambda}$ but have not been assigned norms; one cannot speak of bounded versus unbounded operators, or of continuous dependence of solutions on data.

Third, there is no transform model. In the classical integer-order case, the Fourier or Laplace transform converts differential operators into multiplication operators. The hybrid algebra already achieves a partial version of this (the spectral block is diagonal), but a full analytic transform model requires a completed space.

All three of these issues will be resolved in the next chapter, where the algebraic direct sum is replaced by a weighted Banach space of coefficient sequences.

Looking ahead

The passage from algebra to analysis is not a routine formality. Completion changes the nature of the theory: the operators acquire norms, the generating series become genuine vectors, and a holomorphic transform model emerges. Chapter 8 will develop this completion theory in detail.

After that, the boundary structure will be examined more closely. The defect localization theorem of this chapter shows *where* defects occur (in the one-sided block), but it does not yet address the deeper question of *what happens when one enlarges the space to record boundary information more explicitly*. That question will lead to the boundary-augmented spaces and maximal commuting sectors of Chapter 9.

Remark 7.6.1 (Continuous–discrete parallelism). The hybrid picture of this chapter has a precise discrete counterpart, developed in Chapter 11. There, the domain $(0, \infty)^r \times \mathbb{R}^s$ is replaced by $\mathbb{N}_0^r \times \mathbb{Z}^s$, the Gamma-normalized monomials are replaced by normalized rising factorials, and the exponential characters are replaced by lattice characters. The resulting discrete hybrid algebra is a structural mirror of the continuous one: the same shift relations, the same spectral diagonalization, the same defect localization, and the same tensor-product structure. The reader should keep this parallelism in mind throughout the rest of the book.

Chapter 8

Banach Completions and Analytic Realizations

The previous five chapters constructed a complete algebraic picture of fractional operators on mixed domains. On the canonical hybrid module $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$, the one-sided operators act as commuting coordinate shifts, the whole-space operators act as diagonal spectral multipliers, and all defects are localized in the one-sided block. Every calculation took place in the algebraic direct sum—the space of finite linear combinations of the hybrid basis vectors—and every proof used only finite-dimensional reasoning.

This chapter replaces the algebraic direct sum by a weighted Banach space of coefficient sequences. The passage from algebra to analysis is the central theme. It resolves three deficiencies of the algebraic theory at once: it creates a space in which the important generating functions actually live; it equips the operators with norms so that one can distinguish bounded from unbounded, stable from unstable; and it produces a concrete transform model that converts the one-sided shift operators into multiplication operators on holomorphic functions.

The guiding principle throughout is that the exact algebraic relations from Chapter 7 should *persist* on the completion whenever the weight is compatible with the shift structure. The completion therefore does not replace the algebraic theory; it extends it into the world of Banach-space operator theory.

The material of this chapter follows the research paper AD06. Throughout, we continue to work with the fixed data

$$r \geq 1, \quad s \geq 1, \quad \boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_r) \in (0,1)^r, \quad \Lambda \subset (\mathbb{C}_+)^s.$$

8.1 Why the algebraic direct sum is too small

The algebraic hybrid module $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ was defined in Chapter 7 as the span of all finite linear combinations $\sum c_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$. Its algebraic structure is rich and exact: every commutation relation, every defect identity, every uniqueness theorem holds on this space. Yet the space is too small for three independent reasons.

Generating eigenvectors do not exist. Recall from Chapter 4 (Proposition 4.6.1) that the formal series

$$E_\alpha(\zeta\, x^\alpha) \;=\; \sum_{n=0}^{\infty} \zeta^n\, e_n(x), \qquad e_n(x) = \frac{x^{n\alpha}}{\Gamma(n\alpha+1)},$$

is the natural eigenfunction of the Caputo derivative: applying C_α term by term gives $C_\alpha\, E_\alpha(\zeta x^\alpha) = \zeta\, E_\alpha(\zeta x^\alpha)$. In one variable this is a Mittag–Leffler function—the fractional analogue of the exponential $e^{\zeta x}$. In the full hybrid setting, the corresponding object is the hybrid Mittag–Leffler vector

$$E_{\zeta,\lambda} := \sum_{\mathbf{k}\in\mathbb{N}_0^r} \zeta^{\mathbf{k}}\, e_{\mathbf{k},\lambda},$$

which would be a joint eigenvector of the commuting Caputo tuple and of every spectral multiplier. But this is an infinite series and therefore does *not* belong to $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$. The most important single class of vectors in the theory simply does not exist in the algebraic module.

There are no norms. On $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$, the operators J_i, C_i, and M_σ are well-defined endomorphisms, but they have no associated operator norms. One cannot ask whether J_i is bounded, whether the Neumann series for a mixed operator converges, or whether the solution of a fractional equation depends continuously on its data. All of these are fundamentally analytic questions that require a normed space.

There is no transform model. In classical calculus, the Fourier or Laplace transform converts differential operators into multiplication operators. The hybrid algebra achieves a partial version of this (the spectral block is already diagonal), but a full analytic transform—mapping coefficient sequences to holomorphic functions—requires a completed space in which the correspondence becomes an isomorphism.

All three deficiencies are resolved by a single idea: view each element of $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ not as a function on the domain but as a sequence of coefficients with respect to the canonical basis, and complete in a weighted ℓ^p norm.

Remark 8.1.1. The construction has a natural finite-dimensional precedent. In linear algebra, one identifies $\mathbb{R}^n$ with n-tuples of real numbers by choosing a basis. Changing the norm on $\mathbb{R}^n$ (from the Euclidean norm to a weighted norm, say) does not change the vector space; it changes the geometry of the space and the boundedness properties of the operators. The present chapter does the same thing in infinite dimensions: the canonical hybrid basis identifies $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ with finitely supported sequences on $\mathbb{N}_0^r \times \Lambda$, and the weighted norm gives the space a Banach structure in which the operators acquire quantitative meaning.

8.2 Weighted coefficient spaces and Banach completions

We now define the central object of the chapter.

Definition 8.2.1. A *weight* on the index set $I := \mathbb{N}_0^r \times \Lambda$ is any function

$$\omega : \mathbb{N}_0^r \times \Lambda \longrightarrow (0,\infty).$$

Positivity is required: every index pair receives a strictly positive weight.

Definition 8.2.2. Fix $1 \le p < \infty$ and a weight ω. For a finite-support element

$$u = \sum_{(\mathbf{k},\lambda)\in F} a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda} \in \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}, \qquad F \subset I \text{ finite},$$

define the *weighted coefficient norm*

$$\|u\|_{p,\omega} := \left(\sum_{(\mathbf{k},\lambda)\in F} |a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k},\lambda)^p \right)^{1/p}.$$

The *weighted Banach completion* $X^p_\omega := X^p_\omega(\boldsymbol{\alpha},\Lambda)$ is the completion of $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ with respect to this norm.

The completion admits a simple concrete description.

Proposition 8.2.3. *The space X^p_ω is canonically isometric to the weighted sequence space*

$$\ell^p_\omega(I) := \left\{ a = (a_{\mathbf{k},\lambda})_{(\mathbf{k},\lambda)\in I} : \sum_{(\mathbf{k},\lambda)\in I} |a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k},\lambda)^p < \infty \right\},$$

via the map that sends each element to its coefficient family. Under this identification, $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ corresponds to the finitely supported families, which are dense in $\ell^p_\omega(I)$.

Proof. The map $u = \sum a_{\mathbf{k},\lambda} \, e_{\mathbf{k},\lambda} \mapsto (a_{\mathbf{k},\lambda})$ is a linear isometry from $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ into $\ell^p_\omega(I)$. Its image is exactly the set of finitely supported families, which is dense in $\ell^p_\omega(I)$ for $1 \le p < \infty$. Hence the unique isometric extension maps X^p_ω onto $\ell^p_\omega(I)$. □

Remark 8.2.4. An element of X^p_ω is, a priori, a coefficient family—not a pointwise-defined function on $\Omega_{r,s}$. Whether the formal series $\sum a_{\mathbf{k},\lambda} \, e_{\mathbf{k},\lambda}$ converges pointwise depends on additional hypotheses. For now, the Banach completion is best viewed as an operator-theoretic object. Pointwise and holomorphic realizations will emerge later under geometric weight assumptions (Section 8.6).

Remark 8.2.5. When $p = 2$, the space X^2_ω is a Hilbert space. The normalized vectors

$$f_{\mathbf{k},\lambda} := \omega(\mathbf{k},\lambda) \, e_{\mathbf{k},\lambda}$$

form an orthonormal basis. In this picture the operators J_i and C_i become weighted unilateral shifts along the lattice $\mathbb{N}_0^r$, with multiplicity indexed by the spectral parameter λ. The Hilbert-space setting is the cleanest for spectral arguments, but the theory works for all $1 \le p < \infty$.

Example 8.2.6. In one variable ($r = 1$, $s = 0$, no spectral parameter), the index set is simply $\mathbb{N}_0$, the weight is a positive sequence $\omega(n)$, and X^p_ω is the weighted ℓ^p space with norm $\|(a_n)\|^p = \sum |a_n|^p \omega(n)^p$. The forward shift J_α sends e_n to e_{n+1}, and its operator norm on X^p_ω is $\|J_\alpha\| = \sup_{n\ge 0} \omega(n+1)/\omega(n)$. This is the one-variable prototype of the general theory.

8.3 Shift-admissible weights and bounded extensions

Not every weight is compatible with the shift operators. If the weight ratios along a coordinate direction are unbounded, then the corresponding shift operator acts as an unbounded map and cannot be extended to the completion. The correct compatibility condition is the following.

Definition 8.3.1. A weight ω on $I = \mathbb{N}_0^r \times \Lambda$ is called *shift-admissible* if, for every $i \in \{1, \dots, r\}$, both ratios

$$L_i := \sup_{(\mathbf{k},\lambda)\in I} \frac{\omega(\mathbf{k}+\mathbf{e}_i,\lambda)}{\omega(\mathbf{k},\lambda)} < \infty$$

and

$$R_i := \sup_{\substack{(\mathbf{k},\lambda)\in I \\ k_i \ge 1}} \frac{\omega(\mathbf{k}-\mathbf{e}_i,\lambda)}{\omega(\mathbf{k},\lambda)} < \infty$$

are finite.

The constant L_i measures the maximum cost (in norm) of a single forward step in the i-th lattice direction. The constant R_i measures the maximum cost of a single backward step. Both must be uniformly controlled across the entire lattice and all spectral fibers.

Example 8.3.2. The constant weight $\omega \equiv 1$ is shift-admissible with $L_i = R_i = 1$ for every i. More interesting are the geometric weights of Section 8.5, which have constant ratios $L_i = \rho_i$ and $R_i = \rho_i^{-1}$. A weight of the form $\omega(\mathbf{k},\lambda) = (1+|\mathbf{k}|)^\beta$ for fixed $\beta > 0$ is *not* shift-admissible when $r \ge 2$, because the ratio $\omega(\mathbf{k}+\mathbf{e}_i)/\omega(\mathbf{k})$ depends on $|\mathbf{k}|$ and is unbounded as $|\mathbf{k}| \to \infty$.

Theorem 8.3.3 (Bounded shift extensions). *Let ω be shift-admissible and let $1 \le p < \infty$. Then each J_i and each C_i extends uniquely from the algebraic core to a bounded operator on X_ω^p, with exact operator norms*

$$\|J_i\| = L_i, \qquad \|C_i\| = R_i.$$

Proof. For a finite-support vector $u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$, the basis action $J_i\, e_{\mathbf{k},\lambda} = e_{\mathbf{k}+\mathbf{e}_i,\lambda}$ gives

$$\|J_i u\|_{p,\omega}^p = \sum |a_{\mathbf{k},\lambda}|^p\, \omega(\mathbf{k}+\mathbf{e}_i,\lambda)^p \le L_i^p \sum |a_{\mathbf{k},\lambda}|^p\, \omega(\mathbf{k},\lambda)^p = L_i^p \|u\|_{p,\omega}^p.$$

Hence J_i is bounded on the dense core with $\|J_i\| \le L_i$, and extends uniquely. For the reverse inequality, test on the single basis vector $e_{\mathbf{k},\lambda}$:

$$\frac{\|J_i\, e_{\mathbf{k},\lambda}\|_{p,\omega}}{\|e_{\mathbf{k},\lambda}\|_{p,\omega}} = \frac{\omega(\mathbf{k}+\mathbf{e}_i,\lambda)}{\omega(\mathbf{k},\lambda)}.$$

Taking the supremum over all $(\mathbf{k},\lambda)$ shows $\|J_i\| \ge L_i$.

For C_i, the basis action gives $C_i u = \sum_{k_i \ge 1} a_{\mathbf{k},\lambda}\, e_{\mathbf{k}-\mathbf{e}_i,\lambda}$, so

$$\|C_i u\|_{p,\omega}^p = \sum_{k_i \ge 1} |a_{\mathbf{k},\lambda}|^p\, \omega(\mathbf{k}-\mathbf{e}_i,\lambda)^p \le R_i^p \|u\|_{p,\omega}^p.$$

Testing on $e_{\mathbf{k},\lambda}$ with $k_i \ge 1$ gives $\|C_i\| \ge R_i$. □

The vacuum and boundary-layer projections are automatically bounded, because restricting a sum to a subset of its terms can only decrease the norm.

Proposition 8.3.4. *Each projection Π_i and each $\Pi_{<\mathbf{m}}$ extends uniquely to a contractive projection on X_ω^p.*

Proof. For any finite-support vector u, the norm $\|\Pi_i u\|_{p,\omega}^p$ is a partial sum of the terms in $\|u\|_{p,\omega}^p$, hence at most $\|u\|_{p,\omega}^p$. The same holds for $\Pi_{<\mathbf{m}}$. □

The central consequence is that all algebraic identities survive completion.

Corollary 8.3.5 (Algebraic identities on the completion). *For every $i \in \{1, \dots, r\}$,*

$$C_i J_i = I, \qquad J_i C_i = I - \Pi_i, \qquad [C_i, J_i] = \Pi_i$$

on all of X^p_ω. For all $i \neq j$,

$$J_i J_j = J_j J_i, \qquad C_i C_j = C_j C_i, \qquad C_i J_j = J_j C_i.$$

More generally, for every $\mathbf{m} \in \mathbb{N}_0^r$,

$$C^{\mathbf{m}} J^{\mathbf{m}} = I, \qquad J^{\mathbf{m}} C^{\mathbf{m}} = I - \Pi_{<\mathbf{m}},$$

where $J^{\mathbf{m}} := J_1^{m_1} \cdots J_r^{m_r}$ and $C^{\mathbf{m}} := C_1^{m_1} \cdots C_r^{m_r}$.

Proof. Each identity holds on the algebraic core by Chapter 7. Since every operator involved is bounded on X^p_ω, and the algebraic core is dense, the identities extend to the whole space by continuity. □

The submultiplicativity of the operator norm gives the following useful bounds on higher powers.

Proposition 8.3.6. *For every $\mathbf{m} \in \mathbb{N}_0^r$,*

$$\|J^{\mathbf{m}}\| \leq L^{\mathbf{m}} := L_1^{m_1} \cdots L_r^{m_r}, \qquad \|C^{\mathbf{m}}\| \leq R^{\mathbf{m}} := R_1^{m_1} \cdots R_r^{m_r}.$$

Proof. Since the coordinate operators commute, their bounded extensions also commute. Hence $\|J^{\mathbf{m}}\| \leq \|J_1\|^{m_1} \cdots \|J_r\|^{m_r} = L^{\mathbf{m}}$, and similarly for $C^{\mathbf{m}}$. □

Remark 8.3.7. The product $\|J_i\| \cdot \|C_i\| = L_i R_i$ measures the "round-trip cost" of one forward step followed by one backward step in the i-th direction. The identity $C_i J_i = I$ implies $\|C_i\| \cdot \|J_i\| \geq \|I\| = 1$, so the product satisfies the universal lower bound

$$L_i R_i \;\geq\; 1.$$

This inequality will be promoted to a sharp optimality theorem in Chapter 14, where geometric weights are shown to be the unique minimizers achieving $L_i R_i = 1$.

8.4 Spectral multipliers and closed operators

The spectral block acts diagonally on the label λ. On the completion, bounded symbols yield bounded operators, while arbitrary (unbounded) symbols give rise to closed operators on their maximal domains. We develop both cases.

Bounded multipliers

Theorem 8.4.1 (Isometric multiplier calculus). *For every bounded function $\sigma : \Lambda \to \mathbb{C}$, the multiplier M_σ extends uniquely to a bounded operator on X^p_ω with exact norm*

$$\|M_\sigma\| \;=\; \sup_{\lambda \in \Lambda} |\sigma(\lambda)| \;=:\; \|\sigma\|_\infty.$$

The map $\sigma \mapsto M_\sigma$ is an isometric unital algebra homomorphism from $\ell^\infty(\Lambda)$ into $\mathcal{B}(X^p_\omega)$.

Proof. For a finite-support vector u,

$$\|M_\sigma u\|^p_{p,\omega} = \sum |a_{\mathbf{k},\lambda}|^p \, |\sigma(\lambda)|^p \, \omega(\mathbf{k},\lambda)^p \le \|\sigma\|^p_\infty \|u\|^p_{p,\omega},$$

so $\|M_\sigma\| \le \|\sigma\|_\infty$. For the reverse inequality, fix any $(\mathbf{k}_0, \lambda_0)$ and test on $e_{\mathbf{k}_0,\lambda_0}$:

$$\frac{\|M_\sigma \, e_{\mathbf{k}_0,\lambda_0}\|}{\|e_{\mathbf{k}_0,\lambda_0}\|} = |\sigma(\lambda_0)|.$$

Taking the supremum over $\lambda_0 \in \Lambda$ gives the lower bound.

The homomorphism properties follow from the diagonal action: $M_\sigma M_\tau \, e_{\mathbf{k},\lambda} = \sigma(\lambda)\tau(\lambda)\, e_{\mathbf{k},\lambda} = M_{\sigma\tau}\, e_{\mathbf{k},\lambda}$. □

Corollary 8.4.2. *Every bounded multiplier commutes with every J_i, every C_i, every Π_i, and every $\Pi_{<\mathbf{m}}$ on X^p_ω.*

Proof. The commutation relations hold on the algebraic core because the shift operators move only the index $\mathbf{k}$ while M_σ changes only the scalar weight attached to λ. All operators are bounded, so the relations extend. □

Remark 8.4.3. If $U \subset \mathbb{C}^s$ is any open set containing Λ, the map

$$H^\infty(U) \longrightarrow \mathcal{B}(X^p_\omega), \qquad f \longmapsto M_{f|_\Lambda},$$

is a contractive unital algebra homomorphism with $\|M_{f|_\Lambda}\| = \sup_{\lambda\in\Lambda} |f(\lambda)| \le \|f\|_{H^\infty(U)}$. This is a bounded holomorphic functional calculus for the spectral block.

Closed multipliers for arbitrary symbols

When σ is unbounded on Λ, the multiplier M_σ is no longer defined on all of X^p_ω. However, it admits a natural realization as a closed operator on its maximal domain. This is the same phenomenon that occurs in classical spectral theory: the multiplication operator $f(x) \mapsto x f(x)$ on $L^2(\mathbb{R})$ is unbounded but closed.

Definition 8.4.4. For an arbitrary function $\sigma : \Lambda \to \mathbb{C}$, define the *maximal domain*

$$\mathcal{D}(M_\sigma) \; := \; \left\{ u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda} \in X^p_\omega : \; \sum |\sigma(\lambda)\, a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k},\lambda)^p < \infty \right\}.$$

For $u \in \mathcal{D}(M_\sigma)$, set $M_\sigma u := \sum \sigma(\lambda)\, a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$.

Theorem 8.4.5. *For every $\sigma : \Lambda \to \mathbb{C}$, the operator $(M_\sigma, \mathcal{D}(M_\sigma))$ is densely defined and closed on X^p_ω.*

Proof. Dense domain. Every finite-support vector belongs to $\mathcal{D}(M_\sigma)$, because multiplying a finite family of scalars by σ produces another finite family. Since $\mathcal{H}^{\mathrm{alg}}_{\alpha,\Lambda}$ is dense in X^p_ω, the domain is dense.

Closedness. Let (u_n) be a sequence in $\mathcal{D}(M_\sigma)$ with $u_n \to u$ and $M_\sigma u_n \to v$ in X^p_ω. Write $u_n = \sum a^{(n)}_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$, $u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$, and $v = \sum b_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$.

The key observation is that for each fixed $(\mathbf{k}, \lambda)$, the coordinate functional $u \mapsto a_{\mathbf{k},\lambda}$ is continuous on X^p_ω, since

$$|a_{\mathbf{k},\lambda}| \;\leq\; \omega(\mathbf{k}, \lambda)^{-1} \, \|u\|_{p,\omega}.$$

Therefore $a^{(n)}_{\mathbf{k},\lambda} \to a_{\mathbf{k},\lambda}$ and $\sigma(\lambda) a^{(n)}_{\mathbf{k},\lambda} \to b_{\mathbf{k},\lambda}$ for every $(\mathbf{k}, \lambda)$. It follows that $b_{\mathbf{k},\lambda} = \sigma(\lambda)\, a_{\mathbf{k},\lambda}$. Since $v \in X^p_\omega$, the family $(\sigma(\lambda)\, a_{\mathbf{k},\lambda})$ is p-summable with weight ω, so $u \in \mathcal{D}(M_\sigma)$ and $M_\sigma u = v$. □

Definition 8.4.6. For $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_s) \in [0, \infty)^s$, the *standard Weyl operator on* X^p_ω is the closed multiplier

$$W^{\boldsymbol{\beta}} \;:=\; M_{\lambda \mapsto \lambda^{\boldsymbol{\beta}}},$$

with maximal domain, where $\lambda^{\boldsymbol{\beta}} = \prod_{j=1}^{s} \lambda_j^{\beta_j}$ uses the principal branch on $\mathbb{C}_+$.

Corollary 8.4.7. *The Weyl operator $W^{\boldsymbol{\beta}}$ is densely defined and closed on X^p_ω. It is bounded if and only if the symbol $\lambda \mapsto \lambda^{\boldsymbol{\beta}}$ is bounded on Λ, in which case $\|W^{\boldsymbol{\beta}}\| = \sup_\lambda |\lambda^{\boldsymbol{\beta}}|$. Moreover, the bounded shifts J_i and C_i preserve the domain $\mathcal{D}(W^{\boldsymbol{\beta}})$ and commute with $W^{\boldsymbol{\beta}}$ on that domain:*

$$J_i W^{\boldsymbol{\beta}} = W^{\boldsymbol{\beta}} J_i, \qquad C_i W^{\boldsymbol{\beta}} = W^{\boldsymbol{\beta}} C_i \qquad \text{on } \mathcal{D}(W^{\boldsymbol{\beta}}).$$

Proof. Closedness and density follow from Theorem 8.4.5. For domain invariance: if $u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda} \in \mathcal{D}(W^{\boldsymbol{\beta}})$, then

$$\sum |\lambda^{\boldsymbol{\beta}}\, a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k} + \mathbf{e}_i, \lambda)^p \leq L_i^p \sum |\lambda^{\boldsymbol{\beta}}\, a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k}, \lambda)^p < \infty,$$

so $J_i u \in \mathcal{D}(W^{\boldsymbol{\beta}})$. The argument for C_i uses R_i in place of L_i. Commutativity follows from the algebraic identities on the core, extended by closedness. □

Example 8.4.8. Consider $s = 1$ and $\Lambda = \{n : n \in \mathbb{N}\} \subset \mathbb{C}_+$. The Weyl symbol $\sigma(\lambda) = \lambda^\beta$ for $\beta > 0$ is unbounded on Λ, so W^β is a closed but unbounded operator. Its domain consists of those coefficient families that satisfy the additional summability condition $\sum |n^\beta\, a_{k,n}|^p \, \omega(k, n)^p < \infty$. If instead Λ is a compact subset of $\mathbb{C}_+$, then every power symbol λ^β is bounded and W^β is a bounded operator.

8.5 Geometric weights and genuine generating vectors

Among all shift-admissible weights, one family stands out for its simplicity and structural importance: the *geometric* weights. These are the weights for which the forward and backward ratios are constant across the lattice, and they are the weights under which the Mittag–Leffler eigenvectors become actual elements of the completed space.

Definition 8.5.1. Fix $\boldsymbol{\rho} = (\rho_1, \ldots, \rho_r) \in (0, \infty)^r$ and a positive function $\eta : \Lambda \to (0, \infty)$. The *geometric weight* is

$$\omega_{\boldsymbol{\rho},\eta}(\mathbf{k}, \lambda) \;:=\; \eta(\lambda)\, \boldsymbol{\rho}^{\mathbf{k}}, \qquad \boldsymbol{\rho}^{\mathbf{k}} := \rho_1^{k_1} \cdots \rho_r^{k_r}.$$

We write $X^p_{\boldsymbol{\rho},\eta} := X^p_{\omega_{\boldsymbol{\rho},\eta}}$.

Proposition 8.5.2 (Exact norms for geometric weights)**.** *The geometric weight is shift-admissible, and*

$$\|J_i\| = \rho_i, \qquad \|C_i\| = \rho_i^{-1} \qquad (1 \le i \le r).$$

Consequently, $\|J^{\mathbf{m}}\| = \boldsymbol{\rho}^{\mathbf{m}}$ *and* $\|C^{\mathbf{m}}\| = \boldsymbol{\rho}^{-\mathbf{m}}$ *for every* $\mathbf{m} \in \mathbb{N}_0^r$*. In particular,*

$$\|J_i\| \cdot \|C_i\| = 1 \qquad (1 \le i \le r).$$

Proof. The weight ratios are independent of $(\mathbf{k}, \lambda)$:

$$\frac{\omega_{\rho,\eta}(\mathbf{k} + \mathbf{e}_i, \lambda)}{\omega_{\rho,\eta}(\mathbf{k}, \lambda)} = \rho_i, \qquad \frac{\omega_{\rho,\eta}(\mathbf{k} - \mathbf{e}_i, \lambda)}{\omega_{\rho,\eta}(\mathbf{k}, \lambda)} = \rho_i^{-1} \quad (k_i \ge 1).$$

The result follows from Theorem 8.3.3. For the higher powers, the submultiplicative bound of Proposition 8.3.6 is achieved with equality because the constant ratio $\rho_i^{m_i}$ is attained on every single basis vector. □

The balanced product $\|J_i\| \cdot \|C_i\| = 1$ is the smallest value consistent with the identity $C_i J_i = I$. No weight can achieve a product strictly less than 1 (Remark 8.3.7), and geometric weights are exactly those that achieve equality. This optimality is not obvious from the definition; it will be proved as a sharp theorem in Chapter 14.

Mittag–Leffler vectors in the completion

Definition 8.5.3. Fix $\lambda \in \Lambda$ and $\boldsymbol{\zeta} = (\zeta_1, \ldots, \zeta_r) \in \mathbb{C}^r$ with $|\zeta_i| < \rho_i^{-1}$ for every i. The *hybrid Mittag–Leffler vector* is

$$E_{\zeta,\lambda} := \sum_{\mathbf{k} \in \mathbb{N}_0^r} \boldsymbol{\zeta}^{\mathbf{k}}\, e_{\mathbf{k},\lambda}, \qquad \boldsymbol{\zeta}^{\mathbf{k}} := \zeta_1^{k_1} \cdots \zeta_r^{k_r}.$$

The constraint $|\zeta_i| < \rho_i^{-1}$, equivalently $|\zeta_i|\rho_i < 1$, means that each parameter ζ_i lies in the open disk of radius $\|C_i\| = \rho_i^{-1}$.

Proposition 8.5.4 (Convergence)**.** *Under the condition* $|\zeta_i| < \rho_i^{-1}$ *for all* i*, the vector* $E_{\zeta,\lambda}$ *belongs to* $X^p_{\rho,\eta}$*, with*

$$\|E_{\zeta,\lambda}\|^p_{p,\omega} = \eta(\lambda)^p \prod_{i=1}^{r} \frac{1}{1 - |\zeta_i|^p \rho_i^p}.$$

Proof. Since all terms lie in the single spectral fiber λ,

$$\begin{aligned}\|E_{\zeta,\lambda}\|^p_{p,\omega} &= \sum_{\mathbf{k} \in \mathbb{N}_0^r} |\boldsymbol{\zeta}^{\mathbf{k}}|^p\, \eta(\lambda)^p\, \boldsymbol{\rho}^{p\mathbf{k}} \\ &= \eta(\lambda)^p \prod_{i=1}^{r} \sum_{n=0}^{\infty} (|\zeta_i|\rho_i)^{pn} = \eta(\lambda)^p \prod_{i=1}^{r} \frac{1}{1 - |\zeta_i|^p \rho_i^p},\end{aligned}$$

a product of convergent geometric series. □

Theorem 8.5.5 (Mittag–Leffler eigenvectors)**.** *Let* $|\zeta_i| < \rho_i^{-1}$ *for every* i*. Then:*

(i) $E_{\zeta,\lambda}$ *is an eigenvector of each backward shift:*

$$C_i\, E_{\zeta,\lambda} = \zeta_i\, E_{\zeta,\lambda} \qquad (1 \le i \le r).$$

(ii) $E_{\zeta,\lambda}$ *is an eigenvector of every bounded spectral multiplier:*

$$M_\sigma\, E_{\zeta,\lambda} = \sigma(\lambda)\, E_{\zeta,\lambda}.$$

(iii) $E_{\zeta,\lambda}$ *lies in the domain of every Weyl operator and satisfies*

$$W^{\boldsymbol{\beta}}\, E_{\zeta,\lambda} = \lambda^{\boldsymbol{\beta}}\, E_{\zeta,\lambda} \qquad (\boldsymbol{\beta} \in [0,\infty)^s).$$

Proof. Since C_i is bounded on $X^p_{\rho,\eta}$ and the defining series converges in that space, we may apply C_i term by term:

$$C_i\, E_{\zeta,\lambda} = \sum_{\mathbf{k}\in\mathbb{N}_0^r} \boldsymbol{\zeta}^{\mathbf{k}}\, C_i\, e_{\mathbf{k},\lambda} = \sum_{k_i \ge 1} \boldsymbol{\zeta}^{\mathbf{k}}\, e_{\mathbf{k}-\mathbf{e}_i,\lambda}.$$

Reindexing with $\mathbf{m} = \mathbf{k} - \mathbf{e}_i$ gives

$$C_i\, E_{\zeta,\lambda} = \zeta_i \sum_{\mathbf{m}\in\mathbb{N}_0^r} \boldsymbol{\zeta}^{\mathbf{m}}\, e_{\mathbf{m},\lambda} = \zeta_i\, E_{\zeta,\lambda}.$$

For (ii), the vector is supported on the single spectral fiber λ, so M_σ multiplies every coefficient by $\sigma(\lambda)$.

For (iii), the same single-fiber argument gives $W^{\boldsymbol{\beta}}\, E_{\zeta,\lambda} = \lambda^{\boldsymbol{\beta}}\, E_{\zeta,\lambda}$. The right-hand side belongs to $X^p_{\rho,\eta}$ because it is a scalar multiple of a vector already in the space, so $E_{\zeta,\lambda} \in \mathcal{D}(W^{\boldsymbol{\beta}})$. □

Pointwise realization

Corollary 8.5.6. *For every* $(x, y) \in \Omega_{r,s}$*,*

$$E_{\zeta,\lambda}(x, y) \; = \; \left(\prod_{i=1}^{r} E_{\alpha_i}(\zeta_i\, x_i^{\alpha_i})\right) e^{\langle\lambda, y\rangle},$$

where $E_\alpha(z) = \sum_{n=0}^{\infty} z^n/\Gamma(n\alpha + 1)$ *is the classical Mittag–Leffler function.*

Proof. Substituting the definition of $e_{\mathbf{k},\lambda}$ into the series and rearranging:

$$\sum_{\mathbf{k}\in\mathbb{N}_0^r} \boldsymbol{\zeta}^{\mathbf{k}}\, e_{\mathbf{k},\lambda}(x, y) = \left(\prod_{i=1}^{r} \sum_{n=0}^{\infty} \frac{(\zeta_i x_i^{\alpha_i})^n}{\Gamma(n\alpha_i + 1)}\right) e^{\langle\lambda, y\rangle}.$$ □

Remark 8.5.7. Before completion, the Mittag–Leffler eigenvectors were only formal objects—infinite series outside the algebraic module. After completion, they are genuine Banach-space vectors on which the operators act rigorously. The formal eigenvalue relation of Chapter 4 (Remark 4.6.2) is now a theorem.

Example 8.5.8. In one variable ($r = 1$, $s = 0$), with weight parameter $\rho > 0$, the generating eigenvector $E_\zeta = \sum_{n=0}^{\infty} \zeta^n\, e_n$ has norm $\|E_\zeta\|_p = (1 - |\zeta|^p \rho^p)^{-1/p}$ and satisfies $C_\alpha\, E_\zeta = \zeta\, E_\zeta$. As

$|\zeta| \to \rho^{-1}$ from below, the norm diverges—the eigenvector just barely fits into the space for each $|\zeta| < \rho^{-1}$. The open disk $\{|\zeta| < \rho^{-1}\}$ is the set of eigenvalues of C_α on X^p_ρ. Changing ρ changes the size of this eigenvalue disk, illustrating concretely how the choice of weight determines the spectral content of the completion.

8.6 The first transform model

The geometric completion yields a concrete analytic realization: a fiberwise holomorphic model on a polydisk. This transform converts the one-sided shift operators into multiplication and quotient operators on power series, while the spectral variable remains a discrete fiber parameter.

Definition 8.6.1. Fix a geometric weight $\omega_{\rho,\eta}$. The *reciprocal polydisk* is

$$\mathbb{D}_\rho \;:=\; \{z = (z_1, \dots, z_r) \in \mathbb{C}^r : |z_i| < \rho_i \text{ for all } i\}.$$

Definition 8.6.2 (Fiberwise power-series transform)**.** For $u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda} \in X^p_{\rho,\eta}$ and each fixed $\lambda \in \Lambda$, define

$$(\mathscr{Z}u)(z,\lambda) \;:=\; \sum_{\mathbf{k}\in\mathbb{N}_0^r} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}}, \qquad z \in \mathbb{D}_\rho.$$

Proposition 8.6.3. *For every $u \in X^p_{\rho,\eta}$ and every $\lambda \in \Lambda$:*

(i) The series defining $(\mathscr{Z}u)(\cdot,\lambda)$ converges absolutely on $\mathbb{D}_\rho$ and defines a holomorphic function.

(ii) The transform $\mathscr{Z}$ is injective.

Proof. Fix u, λ, and $z \in \mathbb{D}_\rho$. For $p = 1$, the bound $|z^{\mathbf{k}}| \le \boldsymbol{\rho}^{\mathbf{k}}$ gives

$$\sum_{\mathbf{k}} |a_{\mathbf{k},\lambda}\, z^{\mathbf{k}}| \le \eta(\lambda)^{-1} \sum_{\mathbf{k}} |a_{\mathbf{k},\lambda}|\, \eta(\lambda)\, \boldsymbol{\rho}^{\mathbf{k}} \le \eta(\lambda)^{-1} \|u\|_{1,\omega} < \infty.$$

For $1 < p < \infty$, Hölder's inequality with conjugate exponent $q = p/(p-1)$ yields a similar bound with an additional factor from a convergent geometric series in $(|z_i|/\rho_i)^q$. Absolute convergence on compact subsets gives holomorphy by the Weierstrass theorem.

For injectivity: if $\mathscr{Z}u = 0$, then for each λ the Taylor coefficients at the origin all vanish, so $a_{\mathbf{k},\lambda} = 0$ for every $(\mathbf{k},\lambda)$. □

The transform model represents the operators in a simple form.

Theorem 8.6.4 (Operator intertwining)**.** *Let $u \in X^p_{\rho,\eta}$, let $i \in \{1, \dots, r\}$, and let $\sigma \in \ell^\infty(\Lambda)$. Then:*

(i) Forward shift becomes multiplication: $(\mathscr{Z}(J_i u))(z,\lambda) = z_i\, (\mathscr{Z}u)(z,\lambda)$.

(ii) Spectral multiplier passes through: $(\mathscr{Z}(M_\sigma u))(z,\lambda) = \sigma(\lambda)\, (\mathscr{Z}u)(z,\lambda)$.

(iii) Backward shift becomes a quotient operator: *for $z_i \neq 0$,*

$$(\mathscr{Z}(C_i u))(z, \lambda) = \frac{(\mathscr{Z}u)(z, \lambda) - (\mathscr{Z}u)(z|_{z_i=0}, \lambda)}{z_i},$$

where $z|_{z_i=0}$ denotes the point obtained by replacing z_i with 0. The right-hand side extends holomorphically across $z_i = 0$.

Proof. Write $\varphi(z, \lambda) := (\mathscr{Z}u)(z, \lambda) = \sum_{\mathbf{k}} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}}$.

Part (i). Since J_i shifts the $\mathbf{k}$-index by $\mathbf{e}_i$, the coefficient of $e_{\mathbf{m},\lambda}$ in $J_i u$ is $a_{\mathbf{m}-\mathbf{e}_i,\lambda}$ when $m_i \geq 1$ and 0 when $m_i = 0$. Therefore

$$(\mathscr{Z}(J_i u))(z, \lambda) = \sum_{\mathbf{m}:\, m_i \geq 1} a_{\mathbf{m}-\mathbf{e}_i,\lambda}\, z^{\mathbf{m}} = z_i \sum_{\mathbf{k}} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}} = z_i\, \varphi(z, \lambda).$$

Part (ii) follows directly from the diagonal action of M_σ.

Part (iii). The backward shift C_i annihilates all terms with $k_i = 0$ and shifts the remaining terms down by $\mathbf{e}_i$:

$$(\mathscr{Z}(C_i u))(z, \lambda) = \sum_{k_i \geq 1} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}-\mathbf{e}_i} = \frac{1}{z_i}\left(\sum_{\mathbf{k}} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}} - \sum_{k_i=0} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}}\right).$$

The second sum is $\varphi(z|_{z_i=0}, \lambda)$—the restriction of φ to the hyperplane $z_i = 0$. Since the numerator vanishes at $z_i = 0$, the quotient extends holomorphically. □

Remark 8.6.5. In the transform picture, the geometric completion becomes a mixed Hardy-type space: its elements are families of holomorphic functions on $\mathbb{D}_\rho$, parametrized by the discrete spectral variable λ. The forward shift J_i is the Toeplitz-type multiplication by z_i—a standard object in Hardy-space theory. The backward shift C_i is the "backward quotient" operator, which subtracts the constant term along the i-th coordinate and divides by z_i. The spectral multipliers act fiberwise in λ.

This is the first concrete appearance of a Fourier–Laplace-type model in the hybrid fractional theory. The discrete analogue (a fiberwise Z-transform) appears in Chapter 12, and the unified abstract framework in Chapter 14.

Example 8.6.6. Under $\mathscr{Z}$, the Mittag–Leffler eigenvector $E_{\zeta,\lambda} = \sum \zeta^{\mathbf{k}}\, e_{\mathbf{k},\lambda}$ maps to

$$(\mathscr{Z}\, E_{\zeta,\lambda})(z, \lambda) = \sum_{\mathbf{k}} \zeta^{\mathbf{k}}\, z^{\mathbf{k}} = \prod_{i=1}^{r} \frac{1}{1 - \zeta_i z_i},$$

a product of Cauchy kernels. The eigenvector equation $C_i\, E_{\zeta,\lambda} = \zeta_i\, E_{\zeta,\lambda}$ becomes, in the transform picture, the elementary identity

$$\frac{1}{z_i}\left(\frac{1}{\prod_j (1 - \zeta_j z_j)} - \frac{1}{\prod_{j \neq i}(1 - \zeta_j z_j)}\right) = \frac{\zeta_i}{\prod_j (1 - \zeta_j z_j)},$$

which is the partial-fraction identity for the backward quotient of a Cauchy kernel. In the transform model, the abstract eigenvector theory reduces to elementary algebra.

8.7 Why completion changes the theory

This chapter marks a turning point in the book. The passage from algebra to analysis is not a routine formality; it changes the nature of the theory in several fundamental ways. Before summarizing these changes, we develop the most dramatic illustration: the inversion theory for mixed constant-coefficient operators, which works completely differently on the completion than on the algebraic core.

Banach-space inversion

In Chapter 7 (Theorem 7.5.3), mixed lower-triangular operators were inverted on the algebraic core by exploiting *local nilpotence*: on any finite-support vector, the backward shift eventually reaches zero, so the Neumann series terminates after finitely many steps. The only hypothesis was that the leading symbol $\sigma_{\mathbf{0}}(\lambda)$ be nonvanishing.

On the completion, this argument breaks down. An element of X^p_ω may have infinitely many nonzero coefficients, and the backward shifts no longer reach zero in finitely many steps. The Neumann series becomes a genuine infinite series in operator norm, and its convergence requires quantitative control.

Theorem 8.7.1 (Banach-space invertibility)**.** *Let ω be shift-admissible, let $\sigma_{\mathbf{m}} \in \ell^\infty(\Lambda)$ for each $\mathbf{0} \le \mathbf{m} \le \mathbf{M}$, and define the bounded operator*

$$T := M_{\sigma_{\mathbf{0}}} + \sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} C^{\mathbf{m}} M_{\sigma_{\mathbf{m}}}$$

on X^p_ω. Assume:

- *(a)* (Uniform ellipticity) $m_0 := \inf_{\lambda\in\Lambda} |\sigma_{\mathbf{0}}(\lambda)| > 0$.
- *(b)* (Perturbative smallness) $\displaystyle\sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} R^{\mathbf{m}} \|\sigma_{\mathbf{m}}\|_\infty < m_0$.

Then T is boundedly invertible on X^p_ω, with inverse

$$T^{-1} = \sum_{q=0}^{\infty} (-K)^q M_{\sigma_{\mathbf{0}}^{-1}},$$

where $K := M_{\sigma_{\mathbf{0}}^{-1}} \sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} C^{\mathbf{m}} M_{\sigma_{\mathbf{m}}}$ and the series converges in operator norm.

Proof. The uniform ellipticity ensures that $\sigma_{\mathbf{0}}^{-1}$ is a bounded function on Λ with $\|\sigma_{\mathbf{0}}^{-1}\|_\infty = m_0^{-1}$. Hence $M_{\sigma_{\mathbf{0}}}$ is invertible on X^p_ω. We factor $T = M_{\sigma_{\mathbf{0}}}(I + K)$ and estimate

$$\|K\| \le m_0^{-1} \sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} \|C^{\mathbf{m}}\| \, \|\sigma_{\mathbf{m}}\|_\infty \le m_0^{-1} \sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} R^{\mathbf{m}} \|\sigma_{\mathbf{m}}\|_\infty < 1,$$

using the commutation of multipliers with shifts (Corollary 8.4.2) and the norm bound $\|C^{\mathbf{m}}\| \le R^{\mathbf{m}}$ (Proposition 8.3.6). By the classical Neumann lemma, $I + K$ is invertible with $(I + K)^{-1} = \sum_{q=0}^{\infty} (-K)^q$ converging in operator norm. Composing with $M_{\sigma_{\mathbf{0}}^{-1}}$ gives the formula for T^{-1}. □

Remark 8.7.2. Compare with Theorem 7.5.3 of Chapter 7. In the algebraic setting, the only condition was $\sigma_{\mathbf{0}}(\lambda) \neq 0$ for all λ—a purely qualitative condition. On the completion, two quantitative conditions replace it: the leading symbol must be *bounded away from zero*, and the lower-order terms must be *small in operator norm* relative to the elliptic gap. This is a genuinely new analytic result that does not follow by density from the algebraic inversion theorem.

Example 8.7.3. Take $r = s = 1$, a geometric weight with parameter ρ, and the operator $T = W^{\beta} + a\, C\, W^{\gamma}$ from Example 7.5.4 of Chapter 7. On the geometric completion $X^{p}_{\rho,\eta}$, the Banach-space conditions become: (a) $\inf_{\lambda} |\lambda^{\beta}| > 0$, and (b) $\rho^{-1}|a| \sup_{\lambda} |\lambda^{\gamma}| < \inf_{\lambda} |\lambda^{\beta}|$. When Λ is bounded away from the origin in $\mathbb{C}_{+}$, condition (a) is automatic. Condition (b) is a smallness requirement on the coupling constant a—one that can be satisfied for any fixed a by choosing ρ large enough (thereby shrinking $\|C\| = \rho^{-1}$), at the cost of reducing the radius of the transform polydisk.

Summary of the transformation

The changes produced by completion can be organized into four themes:

(i) *The operators have norms.* The forward and backward shifts are bounded, with exact norms determined by the weight ratios. Spectral multipliers are bounded for bounded symbols and closed for arbitrary symbols. One can now discuss convergence, continuity, and spectral theory.

(ii) *The generating eigenvectors exist.* The hybrid Mittag–Leffler vectors are genuine Banach-space vectors, and the formal eigenvalue relations of Chapter 4 become rigorous identities.

(iii) *A transform model has emerged.* The fiberwise power-series transform converts forward shifts to multiplication, backward shifts to backward quotient operators, and spectral multipliers to pointwise multiplication in λ.

(iv) *Inversion requires new hypotheses.* Local nilpotence is replaced by a quantitative Neumann argument involving uniform ellipticity and perturbative smallness, as demonstrated in Theorem 8.7.1 above.

What has not changed. The fundamental algebraic identities—the shift relations, the vacuum projections, the higher-order defect formulas, the defect localization, the commutativity between the two blocks—all survive completion without modification. The algebraic skeleton of the theory is preserved; what completion adds is the analytic flesh.

Looking ahead. The Banach completions of this chapter provide the functional-analytic framework in which the remaining questions of the theory can be formulated precisely. The next chapter addresses the boundary structure: what happens when one enlarges the ambient space to record additional boundary information? The answer will involve ordered trace words, extended Caputo operators, and the classification of maximal commuting sectors—questions that require the operator norms and closed-operator theory developed here.

Later, in Chapter 14, the geometric weight theory will be revisited from a more abstract perspective. The product $\|J_i\| \cdot \|C_i\| = 1$ for geometric weights will be recognized as the balanced case of a universal inequality, the transform model will be unified with its discrete counterpart, and the semigroup-generation properties of diagonal spectral multipliers will be characterized.

Remark 8.7.4 (The continuous–discrete parallel). The completion program of this chapter has an exact discrete counterpart, developed in Chapter 12. There, the domain $(0,\infty)^{r} \times \mathbb{R}^{s}$ is replaced by $\mathbb{N}_0^r \times \mathbb{Z}^s$, the Gamma-normalized monomials by rising factorials, the exponential characters by lattice characters, and the fiberwise power-series transform by a fiberwise Z-transform. The structure of the discrete completion theory—weighted coefficient norms, shift-admissibility, geometric weights, generating vectors, and Banach-space inversion—mirrors the continuous theory presented here. This structural parallelism is not accidental; in Chapter 14, both theories will be recognized as instances of a single abstract coefficient-space framework.

Chapter 9

Boundary Augmentation and Maximal Commuting Sectors

In the canonical hybrid module and its Banach completions, the partial Caputo operators form a commuting tuple. This commutativity was proved on the algebraic core in Chapter 7 and extended to the completed space in Chapter 8. It is one of the most attractive features of the canonical model: all r backward shifts commute with one another, all r forward shifts commute, and mixed compositions $C_iJ_j = J_jC_i$ hold for $i \neq j$.

This chapter asks what happens when one *enlarges* the canonical space. The canonical completion records the residual one-sided grades and the spectral labels, but it does not record the *order* in which boundary extractions have occurred. If one adjoins additional "boundary-trace layers" that carry this ordering information, does commutativity survive?

The answer is no—but the failure is sharply localized. The extended Caputo operators fail to commute only at very specific boundary strata: those where at least two one-sided coordinates are simultaneously at vacuum (grade zero) inside a defect layer. The whole-space spectral block plays no role at all in the obstruction. This chapter classifies, completely, the largest graded invariant sector on which the extended tuple remains commuting.

The material follows the research paper AD07. Throughout, we work with the fixed data

$$r \geq 1, \quad s \geq 1, \quad \boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r) \in (0,1)^r, \quad \Lambda \subset (\mathbb{C}_+)^s,$$

and with a geometric weight $\boldsymbol{\rho} = (\rho_1, \dots, \rho_r) \in (0, \infty)^r$ and a positive spectral function $\eta : \Lambda \to (0, \infty)$.

9.1 Why go beyond the canonical completion

On the canonical Banach completion $X^p_{\rho,\eta}$ of Chapter 8, the partial Caputo derivative C_i acts as a bounded backward shift: it lowers the i-th one-sided grade by one and annihilates vectors with $k_i = 0$. The annihilation at $k_i = 0$ is the vacuum phenomenon studied throughout the preceding chapters.

But what does the Caputo derivative "really do" when it reaches the boundary? In the canonical model, it simply kills the vector. This is the correct behavior for the canonical space, but one can ask a broader question: is there a natural larger space in which the boundary event is *recorded* rather than erased?

Remark 9.1.1. Consider the two-variable domain $(0, \infty)^2$. A vector $e_{\mathbf{k},\lambda}$ with $\mathbf{k} = (0, 0)$ sits at the corner where both coordinate walls $\{x_1 = 0\}$ and $\{x_2 = 0\}$ meet. If one applies C_1 first, the extraction touches the wall $\{x_1 = 0\}$ before $\{x_2 = 0\}$. If one applies C_2 first, the order is reversed.

In the canonical model both compositions give zero, so the order is invisible. But in a larger space that records boundary history, the two orderings produce distinct states: one carries the trace word (1) and the other the trace word (2). This distinction is the seed of noncommutativity.

There is also a structural motivation. The canonical model is designed to be the "best possible" commuting space. To make this claim precise, one must embed the canonical model in a natural ambient space and prove that it is *maximal* with respect to commutativity. The boundary-augmented space of this chapter provides exactly that ambient environment.

9.2 Ordered words and boundary-trace layers

The boundary-trace layers are parametrized by ordered words recording which one-sided coordinates have already produced boundary extractions, and in what order.

Definition 9.2.1. An *ordered trace word* (or simply *word*) in the alphabet $\{1,\dots,r\}$ is a finite sequence $w = (i_1,\dots,i_m)$ of distinct elements of $\{1,\dots,r\}$. The empty word is denoted $\varnothing$. The set of all such words is

$$\mathcal{W}_r \ := \ \{\varnothing\} \cup \big\{(i_1,\dots,i_m) : 1 \le m \le r,\ i_\nu \in \{1,\dots,r\} \text{ distinct}\big\}.$$

For each word $w = (i_1,\dots,i_m)$, define:

(i) the *length* $\ell(w) := m$ (with $\ell(\varnothing) := 0$),

(ii) the *support* $\operatorname{supp}(w) := \{i_1,\dots,i_m\}$,

(iii) the *free set* $\mathrm{F}(w) := \{1,\dots,r\} \setminus \operatorname{supp}(w)$,

(iv) the *free rank* $q(w) := |\mathrm{F}(w)| = r - \ell(w)$.

If $i \in \mathrm{F}(w)$, write wi for the word obtained by appending the letter i to the end of w.

Example 9.2.2. For $r = 3$, the words include: $\varnothing$ (length 0, free rank 3); (1), (2), (3) (length 1, free rank 2); $(1,2)$, $(1,3)$, $(2,1)$, $(2,3)$, $(3,1)$, $(3,2)$ (length 2, free rank 1); and six words of length 3 with free rank 0. Note that $(1,2)$ and $(2,1)$ have the same support $\{1,2\}$ but different orderings; they represent distinct boundary histories.

The key observation is that the free rank $q(w)$ counts how many one-sided coordinates still have unresolved boundary status. Once the word has consumed all r letters, no free coordinate remains and no further boundary event can occur.

Definition 9.2.3. For each word $w \in \mathcal{W}_r$, the *algebraic trace block* is

$$\mathcal{T}_w^{\mathrm{alg}} \ := \bigoplus_{\mathbf{k}\in\mathbb{N}_0^{\mathrm{F}(w)},\ \lambda\in\Lambda} \mathbb{C}\, t_{w,\mathbf{k},\lambda},$$

where $t_{w,\mathbf{k},\lambda}$ is a purely formal basis symbol. The index $\mathbf{k} \in \mathbb{N}_0^{\mathrm{F}(w)}$ records the *residual grades* in the free one-sided coordinates, and $\lambda \in \Lambda$ records the spectral label.

The symbol $t_{w,\mathbf{k},\lambda}$ should be read as an ordered trace state: the word w says which coordinates have already reached the boundary (and in what order), the residual multi-index $\mathbf{k}$ records the remaining grades in the coordinates that have not yet reached the boundary, and λ is the spectral fiber.

Definition 9.2.4. The *boundary-augmented algebraic ambient space* is

$$\mathfrak{H}^{\mathrm{alg}} := \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda} \oplus \bigoplus_{w\in\mathcal{W}_r} \mathcal{T}_w^{\mathrm{alg}},$$

where $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ is the canonical hybrid module of Chapter 7.

The ambient space thus consists of the original canonical module plus one formal trace block for each ordered word. As a vector space it is simply a larger direct sum over a bigger index set.

The boundary-augmented Banach space

We place the trace layers into the same weighted Banach framework as the canonical completion.

Definition 9.2.5. Fix $1 \le p < \infty$ and a *trace-weight parameter* $\tau > 0$. For a finite-support vector

$$u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda} + \sum_{w\in\mathcal{W}_r} \sum b_{w,\mathbf{k},\lambda}\, t_{w,\mathbf{k},\lambda} \in \mathfrak{H}^{\mathrm{alg}},$$

define the norm

$$\|u\|^p_{p,\boldsymbol{\rho},\tau,\eta} := \sum_{\mathbf{k},\lambda} |a_{\mathbf{k},\lambda}|^p\, \eta(\lambda)^p\, \boldsymbol{\rho}^{p\mathbf{k}} + \sum_{w\in\mathcal{W}_r} \sum_{\mathbf{k},\lambda} |b_{w,\mathbf{k},\lambda}|^p\, \tau^{p(\ell(w)+1)}\, \eta(\lambda)^p\, \boldsymbol{\rho}_w^{p\mathbf{k}},$$

where $\boldsymbol{\rho}_w^{\mathbf{k}} := \prod_{i\in\mathrm{F}(w)} \rho_i^{k_i}$ (empty product equal to 1). The *boundary-augmented completion* is

$$\mathfrak{X}^p_{\boldsymbol{\rho},\tau,\eta} := \overline{\mathfrak{H}^{\mathrm{alg}}} \quad \text{under } \|\cdot\|_{p,\boldsymbol{\rho},\tau,\eta}.$$

For each $w \in \mathcal{W}_r$, let T_w^p denote the closure of $\mathcal{T}_w^{\mathrm{alg}}$ in the ambient norm.

The factor $\tau^{\ell(w)+1}$ measures the cost of entering a trace layer. The exponent $\ell(w)+1$ ensures that even the empty-word trace block $T_\varnothing^p$ sits one weight level above the canonical completion. The parameter τ controls the relative strength of the trace layers: small τ makes them lightweight, large τ makes them dominant.

Proposition 9.2.6. *There is an isometric direct-sum decomposition*

$$\mathfrak{X}^p_{\boldsymbol{\rho},\tau,\eta} = X^p_{\boldsymbol{\rho},\eta} \oplus \bigoplus_{w\in\mathcal{W}_r} T_w^p.$$

In particular, the canonical completion $X^p_{\rho,\eta}$ *embeds isometrically as a complemented subspace.*

Proof. The basis families $\{e_{\mathbf{k},\lambda}\}$ and $\{t_{w,\mathbf{k},\lambda}\}$ are pairwise disjoint, and the ambient norm is the ℓ^p-sum of their individual norms. Since $\mathcal{W}_r$ is finite (it has at most $\sum_{m=0}^{r} r!/(r-m)!$ elements), the ambient space is a finite ℓ^p-direct sum of block completions. The first summand is $X^p_{\rho,\eta}$ by construction. □

9.3 Extended Caputo operators

We now define the extended Caputo operators on the ambient space. On the canonical block, they act exactly as before. On the trace blocks, they act by grade-lowering as long as the relevant coordinate has positive grade; when the grade reaches zero, they record the boundary event by appending a new letter to the trace word.

Definition 9.3.1. For each $i \in \{1,\dots,r\}$, define the *extended Caputo operator* $C_i : \mathfrak{H}^{\mathrm{alg}} \longrightarrow \mathfrak{H}^{\mathrm{alg}}$ on basis vectors as follows.

On the canonical block:

$$C_i\, e_{\mathbf{k},\lambda} \;:=\; \begin{cases} 0, & k_i = 0, \\ e_{\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i \geq 1. \end{cases}$$

This is exactly the canonical backward shift C_i from Chapters 7–8.

On a trace basis vector $t_{w,\mathbf{k},\lambda}$*:*

$$C_i\, t_{w,\mathbf{k},\lambda} \;:=\; \begin{cases} 0, & i \notin \mathrm{F}(w), \\ t_{w,\,\mathbf{k}-\mathbf{e}_i,\,\lambda}, & i \in \mathrm{F}(w),\; k_i \geq 1, \\ t_{wi,\,\mathbf{k}^{\widehat{i}},\,\lambda}, & i \in \mathrm{F}(w),\; k_i = 0, \end{cases}$$

where $\mathbf{k}^{\widehat{i}} \in \mathbb{N}_0^{\mathrm{F}(wi)}$ denotes the residual multi-index obtained by removing the i-th component.

The three cases have clear interpretations:

(i) If i is not free (i.e., i already appears in the word w), then the i-th boundary has already been recorded, and C_i has nothing left to do—it annihilates the vector.

(ii) If i is free and the i-th residual grade is still positive, then C_i lowers that grade by one, exactly as in the canonical shift model.

(iii) If i is free and the i-th residual grade is zero, then C_i records the boundary extraction: it appends the letter i to the word w, producing the longer word wi, and removes i from the set of free coordinates.

The third case is the sole source of noncommutativity.

Definition 9.3.2. For a bounded function $\sigma : \Lambda \to \mathbb{C}$, the *ambient spectral multiplier* acts by

$$M_\sigma\, e_{\mathbf{k},\lambda} := \sigma(\lambda)\, e_{\mathbf{k},\lambda}, \qquad M_\sigma\, t_{w,\mathbf{k},\lambda} := \sigma(\lambda)\, t_{w,\mathbf{k},\lambda}.$$

Theorem 9.3.3 (Boundedness)**.** *Each C_i extends to a bounded operator on $\mathfrak{X}^p_{\rho,\tau,\eta}$ with norm*

$$\|C_i\| = \max\{\rho_i^{-1},\ \tau\}.$$

Each bounded multiplier M_σ extends to a bounded operator with $\|M_\sigma\| = \|\sigma\|_\infty$. Moreover, $C_i M_\sigma = M_\sigma C_i$ on the ambient space.

Proof. On each basis vector, the ratio $\|C_i\, \xi\|/\|\xi\|$ takes one of three values: 0 (when C_i annihilates), ρ_i^{-1} (when it lowers a positive grade), or τ (when it appends a trace letter, since the norm factor changes from $\tau^{\ell(w)+1}$ to $\tau^{\ell(w)+2}$ while $\rho_i^{k_i}$ drops out with $k_i = 0$). Hence $\|C_i\| = \max\{\rho_i^{-1}, \tau\}$.

The multiplier estimate $\|M_\sigma\| = \|\sigma\|_\infty$ follows by the same diagonal argument as in the canonical setting, because M_σ changes only the scalar attached to λ on every basis vector.

Commutativity $C_i M_\sigma = M_\sigma C_i$ holds because C_i changes only the word and the one-sided grades, while M_σ changes only the spectral scalar. Both operators are bounded, so the identity extends from the core. ☐

9.4 The commutator formula

We now compute the commutator $[C_i, C_j]$ explicitly. The result is that the commutator vanishes everywhere except at a very specific locus: simultaneous vacuum in two free one-sided coordinates inside a trace block.

Theorem 9.4.1 (Explicit commutator)**.** *Let $i, j \in \{1, \ldots, r\}$ with $i \neq j$.*

On the canonical block:

$$[C_i, C_j]\, e_{\mathbf{k},\lambda} = 0 \qquad \textit{for all } (\mathbf{k}, \lambda) \in \mathbb{N}_0^r \times \Lambda.$$

On a trace basis vector $t_{w,\mathbf{k},\lambda}$:

$$[C_i, C_j]\, t_{w,\mathbf{k},\lambda} \;=\; \begin{cases} t_{wji,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda} \;-\; t_{wij,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda}, & i, j \in \mathrm{F}(w) \textit{ and } k_i = k_j = 0, \\ 0, & \textit{otherwise}, \end{cases}$$

where $\mathbf{k}^{\widehat{\{i,j\}}}$ denotes the residual multi-index with both the i-th and j-th components removed.

Proof. Canonical block. On $e_{\mathbf{k},\lambda}$, both C_i and C_j are the canonical backward shifts C_i and C_j, which commute on $\mathcal{H}^{\mathrm{alg}}_{\alpha,\Lambda}$ by Corollary 7.3.2 of Chapter 7.

Trace block: setup. Fix a trace basis vector $t_{w,\mathbf{k},\lambda}$. If either i or j is not in $\mathrm{F}(w)$, then the corresponding operator annihilates the vector, making both compositions zero.

Assume $i, j \in \mathrm{F}(w)$. There are four cases depending on whether k_i and k_j are positive or zero.

Case $k_i \geq 1$, $k_j \geq 1$. Both operators simply lower residual grades in different coordinates:

$$C_i C_j\, t_{w,\mathbf{k},\lambda} = t_{w,\,\mathbf{k}-\mathbf{e}_i-\mathbf{e}_j,\,\lambda} = C_j C_i\, t_{w,\mathbf{k},\lambda}.$$

Case $k_i = 0$, $k_j \geq 1$ (and symmetrically $k_i \geq 1$, $k_j = 0$). Applying C_j first lowers k_j; then C_i finds $k_i = 0$ and appends i to the word. Applying C_i first appends i; then C_j lowers the j-grade inside the new word wi. Both compositions give the same result $t_{wi,(\mathbf{k}-\mathbf{e}_j)^{\hat{i}},\lambda}$, so the commutator vanishes.

Case $k_i = 0$, $k_j = 0$. Applying C_j first appends j (producing word wj), then C_i appends i (producing word wji):

$$C_i C_j\, t_{w,\mathbf{k},\lambda} = t_{wji,\mathbf{k}^{\widehat{\{i,j\}}},\lambda}.$$

Reversing the order gives

$$C_j C_i\, t_{w,\mathbf{k},\lambda} = t_{wij,\mathbf{k}^{\widehat{\{i,j\}}},\lambda}.$$

Since the words wji and wij are distinct (they have the same support but different orderings), these are distinct basis vectors. Their difference is the stated commutator. □

Remark 9.4.2. The commutator formula has a clear geometric reading. When two free one-sided coordinates are simultaneously at vacuum, the next boundary extraction must choose one of them first. Choosing j first and then i produces the trace word wji; choosing i first and then j produces wij. The commutator is exactly the difference between these two ordered boundary histories.

Note what does *not* appear in the commutator: the spectral label λ, the spectral multipliers, the residual grades of other free coordinates. The obstruction is a purely ordered-boundary phenomenon, localized at simultaneous vacuum in two free one-sided directions.

Example 9.4.3. For $r = 2$, consider the empty-word trace block. The basis vector $t_{\varnothing,(0,0),\lambda}$ has both free coordinates at grade zero. Then

$$C_1 C_2\, t_{\varnothing,(0,0),\lambda} = t_{(2,1),\varnothing,\lambda}, \qquad C_2 C_1\, t_{\varnothing,(0,0),\lambda} = t_{(1,2),\varnothing,\lambda},$$

and the commutator is $t_{(2,1),\varnothing,\lambda} - t_{(1,2),\varnothing,\lambda} \neq 0$. This is the simplest instance of noncommutativity in the boundary-augmented theory.

9.5 The maximal closed graded invariant commuting sector

We now classify the largest part of the ambient space on which commutativity is preserved. The key notion is that of a *graded sector*: a closed subspace that is spanned by a collection of individual basis lines.

Definition 9.5.1. A closed subspace $U \subset \mathfrak{X}^p_{\rho,\tau,\eta}$ is called *closed graded* if there exists a subset S of the full basis

$$\{e_{\mathbf{k},\lambda}\} \cup \{t_{w,\mathbf{k},\lambda}\}$$

such that U is the closure of $\operatorname{span}(S)$.

The graded hypothesis is essential. Without it, one can always build accidental finite-dimensional commuting subspaces by taking linear combinations that cancel the commutator, and no interesting maximality theorem survives. The earlier chapters were formulated entirely in terms of canonical basis components, so the graded framework is the natural setting for the maximality question.

Definition 9.5.2. The *thin trace sector* is the algebraic span

$$\mathfrak{T}^{\mathrm{alg}}_{\leq 1} := \bigoplus_{\substack{w \in \mathcal{W}_r \\ q(w) \leq 1}} \mathcal{T}^{\mathrm{alg}}_w,$$

consisting of all trace blocks whose free rank is at most 1. The *maximal commuting algebraic sector* is

$$\mathfrak{K}^{\mathrm{alg}} := \mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda} \oplus \mathfrak{T}^{\mathrm{alg}}_{\leq 1},$$

and its Banach closure is

$$\mathfrak{K}^p_{\boldsymbol{\rho},\tau,\eta} := X^p_{\boldsymbol{\rho},\eta} \oplus \bigoplus_{\substack{w \in \mathcal{W}_r \\ q(w) \leq 1}} T^p_w.$$

The thin trace sector contains precisely those trace blocks in which at most one free one-sided coordinate remains. When $q(w) = 0$, no free coordinate exists and every C_i annihilates. When $q(w) = 1$, only one free coordinate i_0 remains; the other operators C_j with $j \neq i_0$ annihilate, so no commutator can form.

Proposition 9.5.3 (Invariance). *The sector $\mathfrak{K}^{\mathrm{alg}}$ is invariant under every C_i and every bounded multiplier M_σ.*

Proof. The canonical block $\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ is invariant under all C_i and all M_σ by construction. For a trace basis vector $t_{w,\mathbf{k},\lambda}$ with $q(w) \leq 1$: if $q(w) = 0$, every C_i annihilates; if $q(w) = 1$ with unique free coordinate i_0, then C_{i_0} either lowers the residual grade (staying in the same trace block) or appends i_0, moving to a word of free rank 0—still inside the thin sector. Multipliers preserve every basis vector's block membership. □

Proposition 9.5.4 (Commutativity). *The tuple $(C_1, \dots, C_r)$ is commuting on $\mathfrak{K}^{\mathrm{alg}}$.*

Proof. On the canonical block, commutativity holds by Theorem 9.4.1(canonical block). On a trace basis vector $t_{w,\mathbf{k},\lambda}$ with $q(w) \leq 1$: for any pair of distinct indices $i \neq j$, at most one can belong to $\mathrm{F}(w)$, so at least one of C_i or C_j annihilates the vector. Hence both compositions C_iC_j and C_jC_i give zero, and the commutator vanishes. □

The maximality theorem says that this is the best one can do.

Theorem 9.5.5 (Maximal commuting graded sector). *Let $U \subset \mathfrak{X}^p_{\boldsymbol{\rho},\tau,\eta}$ be a closed graded subspace such that:*

(i) $X^p_{\boldsymbol{\rho},\eta} \subset U$,

(ii) $C_iU \subset U$ *for all* $i \in \{1, \dots, r\}$,

(iii) the tuple $(C_1, \dots, C_r)$ *is commuting on* U.

Then $U \subset \mathfrak{K}^p_{\boldsymbol{\rho},\tau,\eta}$.

Consequently, $\mathfrak{K}^p_{\boldsymbol{\rho},\tau,\eta}$ *is the* unique maximal *closed graded invariant sector containing the canonical completion and carrying a commuting Caputo tuple.*

Proof. Write $U = \overline{U^{\mathrm{alg}}}$ for a graded algebraic subspace U^{alg}. We must show that U^{alg} contains no trace basis vector $t_{w,\mathbf{k},\lambda}$ with $q(w) \geq 2$.

Suppose for contradiction that such a vector lies in U^{alg}. Since U^{alg} is graded and invariant under every C_i, repeated application of the operators C_i along the free coordinates lowers all positive residual grades to zero. (Each application of C_i to a vector with $k_i \geq 1$ lowers that grade by one without appending a letter.) Hence the vector $t_{w,\mathbf{0},\lambda}$—with all free coordinates at grade zero—also belongs to U^{alg}.

Since $q(w) \geq 2$, there exist distinct $i, j \in \mathrm{F}(w)$. Applying the commutator formula (Theorem 9.4.1):

$$C_iC_j\, t_{w,\mathbf{0},\lambda} = t_{wji,\varnothing,\lambda}, \qquad C_jC_i\, t_{w,\mathbf{0},\lambda} = t_{wij,\varnothing,\lambda}.$$

These are distinct basis vectors (the words wji and wij differ in order). Hence

$$C_iC_j\, t_{w,\mathbf{0},\lambda} \neq C_jC_i\, t_{w,\mathbf{0},\lambda},$$

contradicting the assumption that the tuple commutes on U^{alg}.

Therefore U^{alg} contains no trace basis vector with $q(w) \geq 2$, and $U^{\mathrm{alg}} \subset \mathfrak{K}^{\mathrm{alg}}$. Taking closures gives $U \subset \mathfrak{K}^{p}_{\rho,\tau,\eta}$.

Maximality follows: $\mathfrak{K}^{p}_{\rho,\tau,\eta}$ itself satisfies (i)–(iii) by Propositions 9.5.3 and 9.5.4, and every other sector satisfying (i)–(iii) is contained in it. □

9.6 Why noncommutativity is boundary-generated

The results of this chapter paint a sharp picture of where commutativity holds and where it fails. Let us distill the main lessons.

The spectral block is innocent. The Weyl operators and the spectral multipliers act diagonally on the spectral label λ throughout the entire ambient space. They commute with every C_i on every block—canonical and trace alike (Theorem 9.3.3). The commutator formula (Theorem 9.4.1) never involves λ. The whole-space directions contribute multiplicity but create no obstruction to commutativity.

Noncommutativity arises from ordered boundary creation. The commutator $[C_i, C_j]$ is nonzero only when two distinct free one-sided coordinates are simultaneously at grade zero inside a trace block. In that situation, the next boundary extraction must choose which coordinate to record first, and the two possible orderings produce genuinely different trace words. The commutator is the difference of these two ordered histories.

The critical threshold is free rank 2. The maximal commuting sector consists of the canonical completion plus all trace blocks with free rank at most 1. When $q(w) \leq 1$, there is at most one free coordinate active at any time, so no ordering ambiguity can arise. When $q(w) \geq 2$, two or more free coordinates can simultaneously reach vacuum, and the ordering ambiguity creates a nonzero commutator.

Example 9.6.1 ($r = 1$: no obstruction)**.** When there is only one one-sided coordinate ($r = 1$), every word has free rank 0 or 1. The maximal commuting sector is the entire ambient space: $\mathfrak{K}^{p}_{\rho,\tau,\eta} = \mathfrak{X}^{p}_{\rho,\tau,\eta}$. In one one-sided dimension, there are no mixed boundary events and no noncommutativity.

Example 9.6.2 ($r = 2$: the first obstruction)**.** With $r = 2$, the words are $\varnothing$, (1), (2), $(1, 2)$, $(2, 1)$, with free ranks $2, 1, 1, 0, 0$ respectively. The maximal commuting sector is

$$\mathfrak{R}^{p}_{\rho,\tau,\eta} = X^{p}_{\rho,\eta} \oplus T^{p}_{(1)} \oplus T^{p}_{(2)} \oplus T^{p}_{(1,2)} \oplus T^{p}_{(2,1)}.$$

The only excluded block is $T^{p}_{\varnothing}$, the empty-word trace block with free rank 2. This is the smallest case in which noncommutativity appears: the basic commutator $[C_1, C_2]\, t_{\varnothing,(0,0),\lambda} = t_{(2,1),\varnothing,\lambda} - t_{(1,2),\varnothing,\lambda}$ lives in $T^{p}_{\varnothing}$.

9.7 Restriction as a commuting reduction

The maximal commuting sector $\mathfrak{R}^{p}_{\rho,\tau,\eta}$ is obtained by *restriction*: one starts with the full ambient space and discards the trace blocks whose free rank is too large. This is a natural operation: take a noncommutative ambient space, identify where commutativity fails, and restrict to the largest sector where it holds.

Remark 9.7.1. Restriction is not the only way to recover commutativity from a noncommutative ambient space. An alternative strategy is to take a *quotient*: instead of discarding the problematic vectors, one identifies them—collapses the ordering information that causes the commutator. This quotient approach will be developed in Chapter 15, where the “ordering-defect space” generated by differences of the form $t_{wji,\ldots} - t_{wij,\ldots}$ is shown to be the range of the commutator ideal, and the quotient by this ideal yields a universal commuting reduction.

The two strategies—restriction and quotient—are genuinely different. Restriction preserves the ambient-space structure on a smaller domain; the quotient preserves the full domain but coarsens the space by identifying ordered histories with the same support. In Chapter 15, both reductions will be compared explicitly.

Remark 9.7.2. The maximality theorem gives a precise sense in which the canonical completion $X^{p}_{\rho,\eta}$ is special. It is not just one convenient commuting space among many; it is the core around which the unique maximal graded commuting sector is built. Every graded commuting extension of the canonical completion within the boundary-augmented ambient space must be contained in $\mathfrak{R}^{p}_{\rho,\tau,\eta}$. The canonical model is therefore, in a precise operator-theoretic sense, the natural starting point for any commuting theory of mixed-domain fractional operators.

Looking ahead

This chapter concludes the first part of the continuous theory. Chapters 7–9 form a tightly linked sequence: the hybrid algebra (Chapter 7), its Banach completion (Chapter 8), and its boundary-augmented maximal commuting structure (this chapter).

The next phase of the book develops the discrete counterpart. Chapter 10 introduces the discrete operators (fractional sums, Caputo differences, lattice characters), and Chapters 11–13 mirror the continuous sequence exactly: the discrete hybrid algebra, its weighted Banach completion and Z-transform model, and its boundary-augmented maximal commuting sectors. The structural parallelism between the continuous and discrete theories will be a recurring theme, ultimately unified in the abstract coefficient-space framework of Chapter 14.

Remark 9.7.3 (The discrete mirror). The boundary-augmented theory of this chapter has a precise discrete counterpart in Chapter 13. There, the ordered trace words, the extended discrete Caputo operators, the explicit commutator formula, and the maximal commuting sector classification all carry over with the same combinatorial structure. The reason is that the mechanism of noncommutativity—ordered boundary creation at simultaneous vacuum—is a property of the coefficient-space lattice $\mathbb{N}_0^r$, not of the specific function-space realization.

Chapter 10

Preparation for Discrete Fractional Calculus

The preceding chapters built a complete continuous theory: a canonical hybrid algebra on $(0,\infty)^r \times \mathbb{R}^s$ (Chapter 7), its Banach completion (Chapter 8), and the classification of maximal commuting sectors in boundary-augmented spaces (Chapter 9). The entire construction rested on two ingredients: Gamma-normalized monomials on the one-sided coordinates and exponential characters on the whole-space coordinates.

This chapter begins the second half of the book by introducing the *discrete* counterparts of those ingredients. The continuous half-line $(0,\infty)$ is replaced by the discrete half-line $\mathbb{N} = \{1,2,3,\dots\}$, the full line $\mathbb{R}$ is replaced by the integer lattice $\mathbb{Z}$, and the mixed domain $(0,\infty)^r \times \mathbb{R}^s$ is replaced by the mixed lattice $\mathbb{N}^r \times \mathbb{Z}^s$. The operators change accordingly: integrals become sums, derivatives become differences, exponentials become characters of the form λ^m.

The resulting discrete theory is not merely analogous to the continuous one—it is structurally identical. The same shift relations, the same spectral diagonalization, the same defect localization, and the same boundary mechanism will reappear. The purpose of this preparatory chapter is to introduce the discrete building blocks so that the parallel can be seen cleanly.

The reader who already knows discrete fractional calculus may skim Sections 10.1–10.2 and proceed to the dictionary in Section 10.5, which organizes the continuous–discrete correspondence that governs the rest of the book. For background on discrete fractional calculus, see Goodrich and Peterson [5] and Atıcı and Eloe [1].

10.1 Difference operators and rising factorials

In continuous calculus, the monomial x^n is the building block of Taylor series, and the normalized monomial $x^n/n!$ is the building block of the canonical shift model. In discrete calculus, the role of the monomial is played by the *rising factorial*.

The backward difference

Definition 10.1.1. For a function $u : \mathbb{Z} \to \mathbb{C}$, the *backward difference operator* (nabla) is

$$(\nabla u)(n) \;:=\; u(n) - u(n-1).$$

The backward difference is the discrete analogue of the derivative d/dx. Just as the derivative lowers the degree of a polynomial by one, the backward difference lowers the "degree" of a discrete polynomial by one—provided one uses the right notion of discrete polynomial.

Rising factorials

Definition 10.1.2. For $\beta \geq 0$ and $n \in \mathbb{N}_0$, the *rising factorial* is

$$n^{\overline{\beta}} := \begin{cases} 1, & n = 0 \text{ and } \beta = 0, \\ 0, & n = 0 \text{ and } \beta > 0, \\ \dfrac{\Gamma(n+\beta)}{\Gamma(n)}, & n \in \mathbb{N}. \end{cases}$$

The *normalized rising-factorial monomial* is

$$h_\beta(n) := \frac{n^{\overline{\beta}}}{\Gamma(\beta+1)}, \qquad n \in \mathbb{N}_0.$$

Remark 10.1.3. When $\beta = k \in \mathbb{N}_0$ is a nonnegative integer and $n \in \mathbb{N}$,

$$n^{\overline{k}} = n(n+1)(n+2)\cdots(n+k-1) = \frac{(n+k-1)!}{(n-1)!},$$

with $n^{\overline{0}} = 1$. The rising factorial is also called the *Pochhammer symbol* in parts of the literature. For noninteger β, the formula $n^{\overline{\beta}} = \Gamma(n+\beta)/\Gamma(n)$ extends the definition smoothly via the Gamma function, just as $x^\beta = e^{\beta \log x}$ extends integer powers of x to fractional powers.

Example 10.1.4. For the first few values: $h_0(n) = 1$ for all $n \geq 0$ (with $h_0(0) = 1$), $h_1(n) = n$, and $h_2(n) = n(n+1)/2$. These are the discrete analogues of the continuous normalized monomials $1, x, x^2/2$.

The key property of the rising factorial is the following.

Proposition 10.1.5 (Nabla rule for rising factorials). *For every $\beta > 0$ and every $n \in \mathbb{N}$,*

$$\nabla h_\beta(n) = h_{\beta-1}(n).$$

For $\beta = 0$, one has $\nabla h_0(n) = 0$ for $n \geq 1$.

Proof. For $\beta > 0$ and $n \in \mathbb{N}$,

$$\nabla h_\beta(n) = h_\beta(n) - h_\beta(n-1) = \frac{n^{\overline{\beta}} - (n-1)^{\overline{\beta}}}{\Gamma(\beta+1)}.$$

Using the identity $n^{\overline{\beta}} - (n-1)^{\overline{\beta}} = \beta\,(n-1)^{\overline{\beta-1}} + (n-1)^{\overline{\beta}} - (n-1)^{\overline{\beta}} = \beta\, n^{\overline{\beta-1}}$ (which follows from the functional equation of the Gamma function applied to $\Gamma(n+\beta)/\Gamma(n)$), one obtains

$$\nabla h_\beta(n) = \frac{\beta\, n^{\overline{\beta-1}}}{\Gamma(\beta+1)} = \frac{n^{\overline{\beta-1}}}{\Gamma(\beta)} = h_{\beta-1}(n).$$

For $\beta = 0$, $h_0(n) = 1$ for all $n \geq 1$, so $\nabla h_0(n) = 1 - 1 = 0$. □

This is the exact discrete analogue of the continuous identity $\frac{d}{dx}\left[x^\beta/\Gamma(\beta+1)\right] = x^{\beta-1}/\Gamma(\beta)$.

The normalized rising factorial h_β is the “correct” discrete monomial: the one on which the backward difference acts by lowering the exponent by exactly 1.

Remark 10.1.6. Notice that $h_\beta(0) = 0$ for every $\beta > 0$, while $h_0(0) = 1$. This is the discrete analogue of the fact that $x^\beta/\Gamma(\beta+1)$ vanishes at $x = 0$ for $\beta > 0$ while $e_0(0) = 1$. The vanishing at the boundary $n = 0$ is what will produce the vacuum annihilation in the discrete shift model.

The asymptotic connection to continuous monomials

Proposition 10.1.7. *For every fixed* $\beta \geq 0$,

$$h_\beta(n) \sim \frac{n^\beta}{\Gamma(\beta+1)} \qquad \text{as } n \to \infty.$$

Proof. By the ratio asymptotic of the Gamma function ($\Gamma(n+\beta)/\Gamma(n) \sim n^\beta$ as $n \to \infty$), one has $n^{\overline{\beta}} \sim n^\beta$, and dividing by $\Gamma(\beta+1)$ gives the result. □

Thus, for large n, the discrete monomial $h_\beta(n)$ behaves like its continuous counterpart $x^\beta/\Gamma(\beta+1)$ evaluated at $x = n$. The two theories are asymptotically the same; the differences appear at small n and, crucially, at the boundary $n = 0$.

10.2 Discrete fractional sums and discrete Caputo differences

Just as the Riemann–Liouville integral generalizes iterated integration to fractional order, the *nabla fractional sum* generalizes iterated summation to fractional order. And just as the Caputo derivative replaces the ordinary derivative in the one-sided continuous theory, the *Caputo nabla difference* replaces the ordinary backward difference in the one-sided discrete theory.

Definition 10.2.1 (Left nabla fractional sum)**.** For $\mu > 0$ and a function $u : \mathbb{N}_0 \to \mathbb{C}$, define

$$(\nabla_0^{-\mu} u)(n) := \frac{1}{\Gamma(\mu)} \sum_{q=1}^{n} (n-q+1)^{\overline{\mu-1}}\, u(q), \qquad n \in \mathbb{N}.$$

We set $(\nabla_0^{-\mu} u)(0) := 0$.

This is the discrete analogue of the Riemann–Liouville integral ${}_0I_x^\mu u(x) = \frac{1}{\Gamma(\mu)} \int_0^x (x-t)^{\mu-1} u(t)\, dt$. The kernel $(x-t)^{\mu-1}/\Gamma(\mu)$ is replaced by the discrete kernel $(n-q+1)^{\overline{\mu-1}}/\Gamma(\mu)$, and the integral is replaced by a sum.

Definition 10.2.2 (Caputo nabla fractional difference)**.** For $0 < \mu < 1$ and a function $u : \mathbb{N}_0 \to \mathbb{C}$, define

$$({}^{\mathrm{C}}\nabla_0^{\mu} u)(n) := (\nabla_0^{-(1-\mu)} \nabla u)(n), \qquad n \in \mathbb{N},$$

where $\nabla u(n) = u(n) - u(n-1)$ is the backward difference, extended to $n = 1$ by taking $u(0)$ as the initial value.

The Caputo nabla difference is defined, like its continuous counterpart, by first differencing (the discrete analogue of differentiating) and then summing (the discrete analogue of integrating) with a fractional kernel. The order of operations matters: differencing first ensures that the initial value $u(0)$ participates in the formula, just as the Caputo derivative incorporates the initial data $u(0)$.

Action on rising-factorial monomials

The central computation is the following, which is the discrete analogue of the monomial formulas for J_α and C_α from Chapter 2.

Proposition 10.2.3 (Shift formulas for discrete monomials)**.** *Let* $0 < \mu < 1$.

(i) For every $\beta \geq 0$ *and* $n \in \mathbb{N}$,

$$\nabla_0^{-\mu} h_\beta(n) \;=\; h_{\beta+\mu}(n).$$

(ii) For every $\beta > 0$ *and* $n \in \mathbb{N}$,

$${}^{\mathrm{C}}\nabla_0^{\mu} h_\beta(n) \;=\; h_{\beta-\mu}(n).$$

(iii) The Caputo difference annihilates the constant:

$${}^{\mathrm{C}}\nabla_0^{\mu} h_0(n) \;=\; 0 \qquad (n \in \mathbb{N}).$$

Proof. These are standard identities in the nabla fractional calculus; see, for example, Goodrich and Peterson [5]. Part (i) follows from the convolution identity for rising factorials. Part (ii) follows from applying (i) with exponent $1-\mu$ to the nabla rule $\nabla h_\beta = h_{\beta-1}$. Part (iii) follows because $\nabla h_0 = 0$, so the Caputo difference integrates zero. □

Remark 10.2.4. Compare with the continuous formulas from Chapter 2:

$${}_0I_x^{\mu}\,\frac{x^\beta}{\Gamma(\beta+1)} = \frac{x^{\beta+\mu}}{\Gamma(\beta+\mu+1)}, \qquad {}_0^{\mathrm{C}}D_x^{\mu}\,\frac{x^\beta}{\Gamma(\beta+1)} = \frac{x^{\beta-\mu}}{\Gamma(\beta-\mu+1)}.$$

The discrete formulas are identical in form, with $h_\beta(n)$ replacing $x^\beta/\Gamma(\beta+1)$. This structural identity is the reason the discrete theory mirrors the continuous one so precisely.

Setting $\beta = k\alpha$ for an integer k and $\mu = \alpha$, we obtain exactly the shift formulas that will define the canonical discrete model in the next chapter:

$$\nabla_0^{-\alpha} h_{k\alpha} = h_{(k+1)\alpha}, \qquad {}^{\mathrm{C}}\nabla_0^{\alpha} h_{k\alpha} = h_{(k-1)\alpha} \quad (k \geq 1), \qquad {}^{\mathrm{C}}\nabla_0^{\alpha} h_0 = 0.$$

The fractional sum shifts the exponent up by α, the Caputo difference shifts it down by α, and the constant $h_0 = 1$ is annihilated. These are the exact discrete analogues of the forward shift, backward shift, and vacuum annihilation from Chapter 4.

The partial operators in several variables

For the multivariable discrete theory (Chapter 11), one works on $\mathbb{N}^r$ and applies the nabla fractional sum and Caputo nabla difference coordinatewise, exactly as in the continuous multivariable theory of Chapter 5. For each $i \in \{1, \ldots, r\}$, one defines

$$J_i := \nabla_{0,i}^{-\alpha_i}, \qquad C_i := {}^{\mathrm{C}}\nabla_{0,i}^{\alpha_i},$$

where the subscript $0, i$ indicates that the operator acts in the i-th coordinate with base point 0. All other coordinates are held fixed.

10.3 Lattice characters and the discrete spectral picture

In the continuous theory, the whole-space block on $\mathbb{R}^s$ was carried by the exponential characters $e_\lambda(y) = e^{\langle \lambda, y\rangle}$. In the discrete theory, the whole-space block on $\mathbb{Z}^s$ is carried by *lattice characters*: functions of the form $m \mapsto \lambda^m$.

Definition 10.3.1. For $\lambda = (\lambda_1, \dots, \lambda_s) \in (\mathbb{C}^\times)^s$ and $m = (m_1, \dots, m_s) \in \mathbb{Z}^s$, define the *lattice character*

$$\chi_\lambda(m) \;:=\; \lambda^m \;:=\; \lambda_1^{m_1} \cdots \lambda_s^{m_s}.$$

These are the group characters of the abelian group $\mathbb{Z}^s$: they satisfy $\chi_\lambda(m+m') = \chi_\lambda(m)\chi_\lambda(m')$. In the continuous theory, the exponentials $e^{\langle \lambda, y\rangle}$ are eigenvectors of the partial derivative ∂_{y_j} with eigenvalue λ_j. In the discrete theory, the characters λ^m are eigenvectors of the backward difference ∇_{m_j} with eigenvalue $1 - \lambda_j^{-1}$.

Proposition 10.3.2. *For every $\lambda \in (\mathbb{C}^\times)^s$ and every $j \in \{1, \dots, s\}$,*

$$\nabla_{m_j} \lambda^m \;=\; (1 - \lambda_j^{-1})\, \lambda^m.$$

Proof. By definition, $\nabla_{m_j} \lambda^m = \lambda^m - \lambda^{m-\varepsilon_j} = \lambda^m(1 - \lambda_j^{-1})$. □

The discrete spectral set

In the continuous theory, the spectral set Λ was required to lie in $(\mathbb{C}_+)^s = \{\operatorname{Re}\lambda_j > 0\}$ to ensure convergence of the Laplace transforms in the Weyl calculus. In the discrete theory, the analogous condition is $|\lambda_j| > 1$.

Definition 10.3.3. The *exterior disk* is $\mathbb{E} := \{z \in \mathbb{C} : |z| > 1\}$. A *discrete spectral set* is any nonempty subset $\Lambda \subset \mathbb{E}^s$.

The condition $|\lambda_j| > 1$ ensures two things. First, the discrete Weyl kernel series $\sum_{q=0}^{\infty} a(q)\lambda_j^{-q}$ converges absolutely whenever $a \in \ell^1(\mathbb{N}_0)$. Second, the quantity $1 - \lambda_j^{-1}$ lies in the open right half-plane $\mathbb{C}_+$, so that fractional powers $(1 - \lambda_j^{-1})^\beta$ can be defined via the principal branch. This is exactly the condition needed for the discrete Weyl spectral calculus.

Discrete Weyl operators

The generalized discrete Weyl operators on $\mathbb{Z}^s$ are defined by causal convolution, just as in the continuous case.

Definition 10.3.4. Let $a = (a_1, \dots, a_s)$ with each $a_j \in \ell^1(\mathbb{N}_0)$. For a function u on $\mathbb{N}_0^r \times \mathbb{Z}^s$ and each $j \in \{1, \dots, s\}$, define

$$(I_{W,j}^{a_j} u)(n, m) \;:=\; \sum_{q=0}^{\infty} a_j(q)\, u(n,\, m - q\varepsilon_j).$$

The *generalized discrete Weyl derivative* of order $\ell_j \in \mathbb{N}_0$ is $D_{W,j}^{a_j,\ell_j} := \nabla_{m_j}^{\ell_j} \circ I_{W,j}^{a_j}$.

The action on lattice characters is diagonal: $I_{W,j}^{a_j}(\lambda^m) = \widehat{a_j}(\lambda_j^{-1})\,\lambda^m$, where $\widehat{a_j}(\zeta) = \sum_{q=0}^{\infty} a_j(q)\,\zeta^q$ is the generating function of a_j. This is the discrete analogue of the Laplace-transform eigenvalue from Chapter 6.

The most important special case is the *standard discrete Weyl fractional difference*.

Definition 10.3.5. For $\beta \geq 0$, define the binomial kernel

$$w_\beta(q) \;:=\; (-1)^q \binom{\beta}{q}, \qquad q \in \mathbb{N}_0.$$

The *standard discrete Weyl fractional difference of order* β in the j-th coordinate is

$$(W_j^\beta u)(n, m) \;:=\; \sum_{q=0}^{\infty} w_\beta(q)\, u(n,\, m - q\varepsilon_j).$$

For $\boldsymbol{\beta} = (\beta_1, \dots, \beta_s) \in [0, \infty)^s$, define $W^{\boldsymbol{\beta}} := W_1^{\beta_1} \cdots W_s^{\beta_s}$.

Proposition 10.3.6 (Spectral eigenvalue for discrete Weyl). *For $\lambda \in \mathbb{E}^s$ and $\boldsymbol{\beta} \in [0, \infty)^s$,*

$$W^{\boldsymbol{\beta}} \lambda^m \;=\; \omega_{\boldsymbol{\beta}}(\lambda)\, \lambda^m, \qquad \omega_{\boldsymbol{\beta}}(\lambda) := \prod_{j=1}^{s} (1 - \lambda_j^{-1})^{\beta_j}.$$

Proof. The generating function of the binomial kernel w_β is $(1-z)^\beta$ for $|z| < 1$. Since $|\lambda_j| > 1$, the evaluation at $z = \lambda_j^{-1}$ gives $\widehat{w_{\beta_j}}(\lambda_j^{-1}) = (1 - \lambda_j^{-1})^{\beta_j}$. Composing in all s coordinates multiplies the eigenvalues. □

Remark 10.3.7. In the continuous theory, the Weyl eigenvalue was $\lambda^{\boldsymbol{\beta}} = \prod_j \lambda_j^{\beta_j}$. In the discrete theory, it is $\omega_{\boldsymbol{\beta}}(\lambda) = \prod_j (1 - \lambda_j^{-1})^{\beta_j}$. The role of λ_j is played by $1 - \lambda_j^{-1}$: this is the eigenvalue of ∇_{m_j} on λ^m, just as λ_j is the eigenvalue of ∂_{y_j} on $e^{\lambda_j y_j}$. The spectral algebra is diagonal in both cases, but the symbol changes.

10.4 Generating functions and a preview of the Z-transform

In the continuous theory, the passage from the algebraic module to the Banach completion produced a fiberwise holomorphic model (Section 8.6 of Chapter 8): each coefficient family

$(a_{\mathbf{k},\lambda})$ was mapped to a power series $\sum a_{\mathbf{k},\lambda} z^{\mathbf{k}}$ holomorphic on a polydisk. In the discrete theory, the same construction yields a *Z-transform model*.

At the preparatory level, we need only the following observation.

Definition 10.4.1. Let $(a_k)_{k\geq 0}$ be a sequence of complex numbers. Its *(one-sided) Z-transform* is the formal power series

$$\mathscr{Z}(a)(z) \;:=\; \sum_{k=0}^{\infty} a_k \, z^k.$$

When the series converges on a disk $|z| < R$, it defines a holomorphic function there.

Remark 10.4.2. In the engineering literature, the Z-transform is usually defined with z^{-k} rather than z^k. We use the "power series" convention because it is consistent with the fiberwise transform of Chapter 8 and produces a cleaner interplay with the forward shift J (which becomes multiplication by z) and the backward shift C (which becomes the quotient operator).

The Z-transform is relevant because the discrete canonical basis vectors $h_{k\alpha}$ are the "coefficient functions" of a discrete analogue of the Mittag–Leffler generating series. Recall from Chapter 4 that in the continuous theory,

$$E_\alpha(\zeta \, x^\alpha) = \sum_{k=0}^{\infty} \zeta^k \, e_k(x), \qquad e_k(x) = \frac{x^{k\alpha}}{\Gamma(k\alpha + 1)},$$

is the formal eigenvector of the Caputo derivative. In the discrete theory, the analogous object is

$$\sum_{k=0}^{\infty} \zeta^k \, h_{k\alpha}(n),$$

which is a formal eigenvector of the Caputo nabla difference: ${}^{C}\nabla_0^\alpha$ applied term by term gives $\zeta \sum_{k=0}^{\infty} \zeta^k \, h_{k\alpha}(n)$. The Z-transform maps the coefficient sequence $(\zeta^k)_{k\geq 0}$ to the Cauchy kernel $1/(1-\zeta z)$, just as in the continuous case.

In the full multi-variable hybrid setting (Chapter 12), the Z-transform will become a fiberwise power-series transform on a polydisk, exactly parallel to the holomorphic model of Chapter 8. Under this transform, the forward shift J_i becomes multiplication by z_i, the backward shift C_i becomes a backward quotient operator, and the spectral multipliers act pointwise in λ. The detailed construction will be given in Chapter 12.

10.5 A continuous and discrete dictionary

The most efficient way to prepare for the discrete chapters is to organize the continuous–discrete correspondence into an explicit dictionary. The following table summarizes the replacements that govern the passage from Chapters 7–9 (continuous theory) to Chapters 11–13 (discrete theory).

Concept	*Continuous*	*Discrete*
One-sided domain	$(0,\infty)^r$	$\mathbb{N}^r$
Whole-space domain	$\mathbb{R}^s$	$\mathbb{Z}^s$
Mixed domain	$(0,\infty)^r \times \mathbb{R}^s$	$\mathbb{N}^r \times \mathbb{Z}^s$
Normalized monomial	$\dfrac{x^{k\alpha}}{\Gamma(k\alpha+1)}$	$h_{k\alpha}(n) = \dfrac{n^{\overline{k\alpha}}}{\Gamma(k\alpha+1)}$
Whole-space character	$e^{\langle\lambda,y\rangle}$	λ^m
Spectral set condition	$\Lambda \subset (\mathbb{C}_+)^s$	$\Lambda \subset \mathbb{E}^s = \{\lvert\lambda_j\rvert > 1\}$
Forward shift operator	${}_0I^{\alpha_i}_{x_i}$	$\nabla^{-\alpha_i}_{0,i}$
Backward shift operator	${}^{\mathrm{C}}_{0}D^{\alpha_i}_{x_i}$	${}^{\mathrm{C}}\nabla^{\alpha_i}_{0,i}$
Spectral eigenvalue	$\lambda_j^{\beta_j}$	$(1-\lambda_j^{-1})^{\beta_j}$
Transform model	Fiberwise power series (polydisk)	Fiberwise Z-transform (polydisk)
Generating eigenvector	Mittag–Leffler function	Discrete Mittag–Leffler series
Vacuum vector	$e_{\mathbf{0},\lambda} = e^{\langle\lambda,y\rangle}$	$e_{\mathbf{0},\lambda} = \lambda^m$
Vacuum annihilation	$C_i\, e_{\mathbf{0},\lambda} = 0$	$C_i\, e_{\mathbf{0},\lambda} = 0$

Several entries deserve comment.

Why $\mathbb{N}$ rather than $\mathbb{N}_0$. The one-sided domain uses $\mathbb{N} = \{1,2,3,\dots\}$ rather than $\mathbb{N}_0$ because the nabla fractional sum is naturally defined for $n \geq 1$, with the base point at $n = 0$ playing the role of the boundary. In the canonical basis, $h_{k\alpha}(0) = 0$ for $k \geq 1$ and $h_0(0) = 1$, so the boundary value $n = 0$ contributes only to the vacuum component. (Some authors work on $\mathbb{N}_0^r$, treating $n = 0$ as an included boundary; the distinction is a matter of convention and does not affect the algebraic structure.)

The spectral set condition. The condition $|\lambda_j| > 1$ ensures that $1 - \lambda_j^{-1}$ lies in $\mathbb{C}_+$, so the principal branch of $(1-\lambda_j^{-1})^{\beta}$ is well defined. This is the precise analogue of the continuous condition $\operatorname{Re}\lambda_j > 0$.

What is structurally identical. The shift relations ($J_i\, e_{\mathbf{k},\lambda} = e_{\mathbf{k}+\mathbf{e}_i,\lambda}$, $C_i\, e_{\mathbf{k},\lambda} = e_{\mathbf{k}-\mathbf{e}_i,\lambda}$ or 0), the commutation of all operators, the defect identity $J_iC_i = I - \Pi_i$, the defect localization, the uniqueness of the canonical basis, the invertibility of mixed lower-triangular operators, the weighted completion theory, the generating eigenvectors, the transform model, the boundary-augmented ambient space, the commutator formula, and the maximal commuting sector classification—all of these carry over from the continuous to the discrete setting with essentially the same proofs. The reason is that all of these results depend only on the *coefficient-space structure* ($\mathbb{N}_0^r \times \Lambda$ with shifts and multipliers), not on the specific function-space realization (monomials versus rising factorials, exponentials versus lattice characters).

What changes. The function-space realization changes (monomials become rising factorials, exponentials become lattice characters), the spectral eigenvalue changes (from $\lambda_j^{\beta_j}$ to $(1-\lambda_j^{-1})^{\beta_j}$), and the spectral-set condition changes (from $\operatorname{Re}\lambda_j > 0$ to $|\lambda_j| > 1$). These are genuine differences, but they do not affect the algebraic or analytic structure of the hybrid algebra.

Remark 10.5.1. The fact that the continuous and discrete theories have identical coefficient-space structures suggests that they should be unified in a single abstract framework. This unification is the subject of Chapter 14, where the common coefficient-space model is isolated and the transform theory, optimal weights, and semigroup generation are developed in a manner that encompasses both the continuous and discrete cases simultaneously.

Looking ahead

This chapter has introduced the discrete building blocks: rising factorials, nabla fractional sums, Caputo nabla differences, lattice characters, and discrete Weyl operators. The reader has seen that each continuous ingredient has a precise discrete counterpart, and that the fundamental shift formulas are structurally identical.

The next three chapters develop the discrete theory in parallel with the continuous one. Chapter 11 constructs the discrete hybrid shift-spectral algebra on $\mathbb{N}^r \times \mathbb{Z}^s$ (the discrete AD05). Chapter 12 develops the weighted Banach completions, discrete Mittag–Leffler eigenvectors, and the fiberwise Z-transform model (the discrete AD06). Chapter 13 classifies the maximal commuting sectors in the discrete boundary-augmented space (the discrete AD07). Throughout, the reader should keep the continuous–discrete dictionary of this chapter in mind: every theorem in the discrete chapters has a continuous counterpart, and the proofs follow the same logic.

Chapter 11

Discrete Hybrid Shift-Spectral Algebra

Chapter 10 introduced the discrete building blocks: normalized rising factorials, nabla fractional sums, Caputo nabla differences, lattice characters, and discrete Weyl operators. This chapter assembles them into a discrete hybrid shift-spectral algebra on the mixed lattice

$$\Omega_{r,s}^{\mathrm{disc}} := \mathbb{N}^r \times \mathbb{Z}^s.$$

The construction is the exact discrete counterpart of the continuous hybrid algebra developed in Chapter 7.

On the one-sided block $\mathbb{N}^r$, the discrete Riemann–Liouville sums and Caputo differences act as commuting forward and backward shifts on the grade index. On the whole-space block $\mathbb{Z}^s$, the discrete Weyl operators act diagonally on the spectral label. The two blocks commute, the defects are entirely localized in the one-sided coordinates, the canonical basis is unique up to fiberwise scalars, and mixed constant-coefficient difference equations reduce to spectral division plus finite downward recursion.

Every theorem in this chapter has a continuous counterpart in Chapter 7, and every proof follows the same logic. The reader should keep the continuous–discrete dictionary of Chapter 10 at hand throughout.

The material follows the research paper AD08. Throughout, we fix

$$r \geq 1, \quad s \geq 1, \quad \boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r) \in (0,1)^r, \quad \Lambda \subset \mathbb{E}^s, \quad \mathbb{E} := \{z \in \mathbb{C} : |z| > 1\}.$$

11.1 The discrete mixed domain and the factorial-character basis

The mixed lattice

The continuous mixed domain $(0,\infty)^r \times \mathbb{R}^s$ is replaced by its discrete counterpart $\Omega_{r,s}^{\mathrm{disc}} = \mathbb{N}^r \times \mathbb{Z}^s$. Points are denoted (n, m) with $n = (n_1, \dots, n_r) \in \mathbb{N}^r$ and $m = (m_1, \dots, m_s) \in \mathbb{Z}^s$. The first r coordinates are one-sided (bounded below by the boundary at $n_i = 0$), and the last s coordinates are whole-space (extending in both directions without a boundary).

The canonical discrete hybrid basis

Definition 11.1.1. For $\mathbf{k} = (k_1, \dots, k_r) \in \mathbb{N}_0^r$ and $\lambda = (\lambda_1, \dots, \lambda_s) \in \Lambda$, define the *discrete hybrid basis vector*

$$e_{\mathbf{k},\lambda}(n, m) \;:=\; \left(\prod_{i=1}^{r} h_{k_i\alpha_i}(n_i)\right) \lambda^m, \qquad (n, m) \in \Omega_{r,s}^{\mathrm{disc}},$$

where $h_\beta(n) = n^{\overline{\beta}}/\Gamma(\beta + 1)$ is the normalized rising-factorial monomial of Chapter 10 and $\lambda^m = \lambda_1^{m_1} \cdots \lambda_s^{m_s}$ is the lattice character.

The *canonical discrete hybrid module* is the algebraic direct sum

$$\mathcal{H}_{\boldsymbol{\alpha},\Lambda}^{\mathrm{disc,alg}} \;:=\; \bigoplus_{(\mathbf{k},\lambda)\in\mathbb{N}_0^r\times\Lambda} \mathbb{C}\, e_{\mathbf{k},\lambda}.$$

Each basis vector is the product of r factorial monomials (one per one-sided coordinate) and a lattice character (encoding all s whole-space coordinates). This is the exact discrete analogue of the continuous hybrid basis $e_{\mathbf{k},\lambda}(x, y) = \prod_i x_i^{k_i\alpha_i}/\Gamma(k_i\alpha_i + 1) \cdot e^{\langle\lambda,y\rangle}$ from Chapter 7.

Proposition 11.1.2 (Linear independence)**.** *The family $\{e_{\mathbf{k},\lambda} : (\mathbf{k}, \lambda) \in \mathbb{N}_0^r \times \Lambda\}$ is linearly independent.*

Proof. Suppose $\sum_\nu \sum_\ell c_{\nu,\ell}\, e_{\mathbf{k}^{(\nu,\ell)},\lambda^{(\nu)}}(n, m) = 0$ for all (n, m), where the $\lambda^{(\nu)}$ are distinct. Grouping by spectral label gives $\sum_\nu p_\nu(n)\, (\lambda^{(\nu)})^m = 0$. Since distinct characters of $\mathbb{Z}^s$ are linearly independent, $p_\nu(n) = 0$ for all ν and all $n \in \mathbb{N}^r$. Each p_ν is a finite linear combination of products of rising-factorial monomials with distinct exponent tuples. The asymptotic formula $h_\beta(n) \sim n^\beta/\Gamma(\beta + 1)$ as $n \to \infty$ (Proposition 10.1.7) implies that such a combination cannot vanish identically on $\mathbb{N}^r$ unless all coefficients are zero. □

The tensor-product viewpoint

Let $\mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{disc,alg}} := \bigoplus_{\mathbf{k}} \mathbb{C} \prod_i h_{k_i\alpha_i}(n_i)$ denote the discrete factorial lattice and $\mathcal{E}_\Lambda^{\mathrm{disc,alg}} := \bigoplus_\lambda \mathbb{C}\, \lambda^m$ the discrete spectral module. Then

$$\mathcal{H}_{\boldsymbol{\alpha},\Lambda}^{\mathrm{disc,alg}} \;\cong\; \mathcal{G}_{\boldsymbol{\alpha}}^{\mathrm{disc,alg}} \otimes \mathcal{E}_\Lambda^{\mathrm{disc,alg}}$$

algebraically, via $e_{\mathbf{k},\lambda} \leftrightarrow \prod_i h_{k_i\alpha_i}(n_i) \otimes \lambda^m$. The content of this chapter is that the operators also decompose in the expected way: shifts on the first factor, multipliers on the second, and complete commutativity between the two.

11.2 The discrete shift-spectral theorem

We now prove the main structural result. Define, for each $i \in \{1, \dots, r\}$,

$$J_i \;:=\; \nabla_{0,i}^{-\alpha_i}, \qquad C_i \;:=\; {}^{\mathrm{C}}\nabla_{0,i}^{\alpha_i},$$

the partial nabla fractional sum and partial Caputo nabla difference of order α_i in the i-th coordinate.

Shift action in the one-sided block

Theorem 11.2.1 (Discrete shift action)**.** *For every* $(\mathbf{k},\lambda) \in \mathbb{N}_0^r \times \Lambda$ *and every* $i \in \{1,\dots,r\}$*,*

$$J_i\, e_{\mathbf{k},\lambda} = e_{\mathbf{k}+\mathbf{e}_i,\lambda},$$

and

$$C_i\, e_{\mathbf{k},\lambda} = \begin{cases} 0, & k_i = 0, \\ e_{\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i \geq 1. \end{cases}$$

In particular, each J_i *and* C_i *defines a linear endomorphism of* $\mathcal{H}_{\boldsymbol{\alpha},\Lambda}^{\text{disc,alg}}$*.*

Proof. Write $e_{\mathbf{k},\lambda}(n,m) = \phi(\widehat{n}_i,m)\, h_{k_i\alpha_i}(n_i)$, where $\phi(\widehat{n}_i,m) = \left(\prod_{\ell\neq i} h_{k_\ell\alpha_\ell}(n_\ell)\right)\lambda^m$ is independent of n_i. The discrete monomial shift formulas from Chapter 10 (Proposition 10.2.3) give:

$$J_i\, e_{\mathbf{k},\lambda} = \phi(\widehat{n}_i,m)\, h_{(k_i+1)\alpha_i}(n_i) = e_{\mathbf{k}+\mathbf{e}_i,\lambda}.$$

If $k_i = 0$, then $e_{\mathbf{k},\lambda} = \phi(\widehat{n}_i,m)\, h_0(n_i) = \phi(\widehat{n}_i,m)$, and the Caputo difference of a function independent of n_i vanishes. If $k_i \geq 1$, the Caputo formula gives $C_i\, e_{\mathbf{k},\lambda} = \phi(\widehat{n}_i,m)\, h_{(k_i-1)\alpha_i}(n_i) = e_{\mathbf{k}-\mathbf{e}_i,\lambda}$. □

Corollary 11.2.2 (Discrete shift commutation)**.** *For all* $i \neq j$*,* $J_iJ_j = J_jJ_i$*,* $C_iC_j = C_jC_i$*, and* $C_iJ_j = J_jC_i$ *on* $\mathcal{H}_{\boldsymbol{\alpha},\Lambda}^{\text{disc,alg}}$*.*

Proof. Each identity is checked on the basis: J_i raises only the i-th index, C_i lowers only the i-th index, so operators in different coordinates do not interfere. □

Definition 11.2.3. The *discrete coordinate-vacuum projection* Π_i is defined by $\Pi_i\, e_{\mathbf{k},\lambda} := e_{\mathbf{k},\lambda}$ if $k_i = 0$ and 0 otherwise.

Theorem 11.2.4 (Discrete fundamental shift relations)**.** *For every* $i \in \{1,\dots,r\}$*,*

$$C_iJ_i = I, \qquad J_iC_i = I - \Pi_i, \qquad [C_i, J_i] = \Pi_i.$$

Proof. On a basis vector $e_{\mathbf{k},\lambda}$: $C_iJ_i\, e_{\mathbf{k},\lambda} = C_i\, e_{\mathbf{k}+\mathbf{e}_i,\lambda} = e_{\mathbf{k},\lambda}$, giving $C_iJ_i = I$. For J_iC_i: if $k_i \geq 1$, $J_iC_i\, e_{\mathbf{k},\lambda} = J_i\, e_{\mathbf{k}-\mathbf{e}_i,\lambda} = e_{\mathbf{k},\lambda}$; if $k_i = 0$, $C_i\, e_{\mathbf{k},\lambda} = 0$. Hence $J_iC_i = I - \Pi_i$. □

These are identical to the continuous identities of Theorem 7.3.4 in Chapter 7—the same equations, the same proof structure, applied to a different function-space realization.

Spectral action in the whole-space block

Theorem 11.2.5 (Discrete spectral action)**.** *Let* $a = (a_1,\dots,a_s)$ *with each* $a_j \in \ell^1(\mathbb{N}_0)$ *be an admissible kernel tuple, and let* $\ell = (\ell_1,\dots,\ell_s) \in \mathbb{N}_0^s$*. Then for every* $(\mathbf{k},\lambda)$*,*

$$I_{W,a}\, e_{\mathbf{k},\lambda} = \widehat{a}(\lambda^{-1})\, e_{\mathbf{k},\lambda}, \qquad D_{W,a,\ell}\, e_{\mathbf{k},\lambda} = \sigma_{a,\ell}(\lambda)\, e_{\mathbf{k},\lambda},$$

where $\widehat{a}(\lambda^{-1}) = \prod_j \widehat{a_j}(\lambda_j^{-1})$ *and* $\sigma_{a,\ell}(\lambda) = \prod_j (1-\lambda_j^{-1})^{\ell_j}\,\widehat{a_j}(\lambda_j^{-1})$*. In particular, the whole-space operators leave the one-sided index* $\mathbf{k}$ *unchanged.*

Proof. Write $e_{\mathbf{k},\lambda}(n,m) = \psi_{\mathbf{k}}(n)\,\lambda^m$. Since $\psi_{\mathbf{k}}(n)$ is independent of m, the Weyl computation reduces to the one-variable character eigenvalue (Lemma from Chapter 10): $I^{a_j}_{W,j}(\lambda^m) = \widehat{a}_j(\lambda_j^{-1})\lambda^m$, and $\nabla^{\ell_j}_{m_j}(\lambda^m) = (1-\lambda_j^{-1})^{\ell_j}\lambda^m$. Composing in all s coordinates multiplies the eigenvalues. □

Corollary 11.2.6 (Mixed commutativity). *Every one-sided operator (J_i or C_i) commutes with every whole-space operator ($I_{W,a}$ or $D_{W,a,\ell}$) on $\mathcal{H}^{\mathrm{disc,alg}}_{\boldsymbol{\alpha},\Lambda}$.*

Proof. Both identities are checked on basis vectors. The one-sided operators change only $\mathbf{k}$; the whole-space operators change only the scalar eigenvalue attached to λ. □

Discrete spectral multipliers

Definition 11.2.7. For $\sigma : \Lambda \to \mathbb{C}$, the *discrete spectral multiplier* M_σ acts by $M_\sigma\, e_{\mathbf{k},\lambda} := \sigma(\lambda)\, e_{\mathbf{k},\lambda}$.

Theorem 11.2.8 (Discrete multiplier algebra). *The algebra $\mathscr{M}(\Lambda) = \{M_\sigma\}$ is a commutative subalgebra of* $\mathrm{End}(\mathcal{H}^{\mathrm{disc,alg}}_{\boldsymbol{\alpha},\Lambda})$*, the map $\sigma \mapsto M_\sigma$ is an algebra isomorphism, and every multiplier commutes with every J_i and every C_i.*

Proof. Identical to the continuous case (Theorem 7.3.8 of Chapter 7). □

The standard discrete Weyl fractional differences are multipliers.

Proposition 11.2.9 (Discrete Weyl eigenvalue). *For $\boldsymbol{\beta} \in [0,\infty)^s$,*

$$W^{\boldsymbol{\beta}}\, e_{\mathbf{k},\lambda} = \omega_{\boldsymbol{\beta}}(\lambda)\, e_{\mathbf{k},\lambda}, \qquad \omega_{\boldsymbol{\beta}}(\lambda) := \prod_{j=1}^{s}(1-\lambda_j^{-1})^{\beta_j}.$$

Thus $W^{\boldsymbol{\beta}} = M_{\omega_{\boldsymbol{\beta}}}$, and the law of exponents holds: $W^{\boldsymbol{\beta}}W^{\boldsymbol{\gamma}} = W^{\boldsymbol{\beta}+\boldsymbol{\gamma}}$.

Proof. The generating function of the binomial kernel w_β is $(1-z)^\beta$ for $|z|<1$. Since $|\lambda_j|>1$, evaluation at $z = \lambda_j^{-1}$ gives the eigenvalue. The law of exponents follows from $(1-\lambda_j^{-1})^{\beta_j}(1-\lambda_j^{-1})^{\gamma_j} = (1-\lambda_j^{-1})^{\beta_j+\gamma_j}$ on $\mathbb{C}_+$ via the principal branch. □

The abstract discrete hybrid algebra

Theorem 11.2.10 (Discrete transport theorem). *The map $U : \mathcal{H}^{\mathrm{disc,alg}}_{\boldsymbol{\alpha},\Lambda} \to c_{00}(\mathbb{N}_0^r \times \Lambda)$ defined by $U(e_{\mathbf{k},\lambda}) := u_{\mathbf{k},\lambda}$ satisfies*

$$UJ_iU^{-1} = S_i^+, \qquad UC_iU^{-1} = S_i^-, \qquad UM_\sigma U^{-1} = \widetilde{M}_\sigma,$$

where S_i^+, S_i^-, and $\widetilde{M}_\sigma$ are the abstract coordinate shifts and diagonal multipliers on the finitely supported scalar space.

Proof. This is immediate from the basis-level actions. □

Remark 11.2.11 (The key observation). Theorem 11.2.10 says that the abstract operator algebra on $\mathcal{H}^{\mathrm{disc,alg}}_{\boldsymbol{\alpha},\Lambda}$ is *exactly the same* as the abstract operator algebra on the continuous hybrid module

$\mathcal{H}^{\mathrm{alg}}_{\boldsymbol{\alpha},\Lambda}$ from Chapter 7 (Theorem 7.3.11). Both are realized as coordinate shifts on $\mathbb{N}_0^r$ tensored with diagonal multipliers on Λ. The continuous and discrete function-space realizations differ, but the coefficient-space algebra is identical. This observation is the structural reason for the continuous–discrete parallelism, and it will be formalized in the abstract framework of Chapter 14.

11.3 Combining the one-sided and whole-space blocks

The results of the previous section show that the discrete hybrid module carries two complementary operator structures. We now record the higher-power identities and the defect localization theorem that combine them.

Definition 11.3.1. For $\mathbf{m} = (m_1, \dots, m_r) \in \mathbb{N}_0^r$, define $J^{\mathbf{m}} := J_1^{m_1} \cdots J_r^{m_r}$ and $C^{\mathbf{m}} := C_1^{m_1} \cdots C_r^{m_r}$. The *discrete boundary-layer projection* $\Pi_{<\mathbf{m}}$ is defined by $\Pi_{<\mathbf{m}}\, e_{\mathbf{k},\lambda} := e_{\mathbf{k},\lambda}$ if $\mathbf{k} \not\geq \mathbf{m}$, and 0 otherwise.

Theorem 11.3.2 (Discrete higher-order shift relations)**.** *For every* $\mathbf{m} \in \mathbb{N}_0^r$*,*

$$C^{\mathbf{m}} J^{\mathbf{m}} = I, \qquad J^{\mathbf{m}} C^{\mathbf{m}} = I - \Pi_{<\mathbf{m}}.$$

Proof. By repeated application of Theorem 11.2.1: $J^{\mathbf{m}}\, e_{\mathbf{k},\lambda} = e_{\mathbf{k}+\mathbf{m},\lambda}$ and $C^{\mathbf{m}}\, e_{\mathbf{k},\lambda} = e_{\mathbf{k}-\mathbf{m},\lambda}$ when $\mathbf{k} \geq \mathbf{m}$, 0 otherwise. The identities follow by composing. □

Theorem 11.3.3 (Discrete defect localization)**.** *For every* $\mathbf{m} \in \mathbb{N}_0^r$ *and every* $\sigma : \Lambda \to \mathbb{C}$*,*

$$\Pi_{<\mathbf{m}} M_\sigma = M_\sigma \Pi_{<\mathbf{m}}.$$

The higher-order defect sector is entirely contained in the one-sided block: the whole-space Weyl variables contribute no vacuum projection and no boundary defect.

Proof. On basis vectors, $\Pi_{<\mathbf{m}}$ tests only the condition $\mathbf{k} \not\geq \mathbf{m}$, while M_σ multiplies by $\sigma(\lambda)$. The two operations commute because they affect different parts of the index pair. □

Theorem 11.3.4 (Discrete uniqueness)**.** *Let* $\{f_{\mathbf{k},\lambda}\}$ *be a family of functions on* $\Omega^{\mathrm{disc}}_{r,s}$ *of the form*

$$f_{\mathbf{k},\lambda}(n, m) = c_{\mathbf{k},\lambda} \left(\prod_{i=1}^{r} h_{\beta_i(\mathbf{k})}(n_i) \right) \lambda^m$$

with $c_{\mathbf{k},\lambda} \neq 0$ *and* $\beta_i(\mathbf{k}) \geq 0$*. If the shift relations* $J_i\, f_{\mathbf{k},\lambda} = f_{\mathbf{k}+\mathbf{e}_i,\lambda}$ *and* $C_i\, f_{\mathbf{k},\lambda} = f_{\mathbf{k}-\mathbf{e}_i,\lambda}$ *(or* 0 *when* $k_i = 0$*) hold for all* i*, then* $f_{\mathbf{k},\lambda} = c_\lambda\, e_{\mathbf{k},\lambda}$ *for a unique nonzero scalar* c_λ *depending only on* λ*.*

Proof. The proof is identical to the continuous uniqueness theorem (Theorem 7.4.6 of Chapter 7). The forward shift condition forces $\beta_i(\mathbf{k}) = k_i \alpha_i$ by induction from the vacuum condition $\beta_i(\mathbf{k}) = 0$ when $k_i = 0$. The coefficient recurrence then gives a constant c_λ per spectral fiber. □

Remark 11.3.5. Mixed constant-coefficient discrete equations on the canonical module are solved by the same mechanism as in the continuous theory (Section 7.5 of Chapter 7). A mixed operator $T = \sum_{\mathbf{0} \leq \mathbf{m} \leq \mathbf{M}} C^{\mathbf{m}} M_{\sigma_{\mathbf{m}}}$ is invertible whenever $\sigma_{\mathbf{0}}(\lambda) \neq 0$ for all $\lambda \in \Lambda$, because the lower-order part $N = \sum_{\mathbf{m} > \mathbf{0}} C^{\mathbf{m}} M_{\sigma_{\mathbf{m}}}$ is locally nilpotent on the algebraic core. The inverse is given by the same finite Neumann series as in the continuous case. Specifically, if the operator includes discrete Weyl derivatives, the condition becomes nonvanishing of $\sum_\nu a_{\mathbf{0},\nu} \prod_j (1-\lambda_j^{-1})^{\beta_j^{(\nu)}}$—the discrete spectral symbol. The solution is obtained by spectral division in λ followed by finite downward recursion in the one-sided grades.

11.4 Why AD08 is a mirror, not a loose analogy

The reader may have noticed that every theorem in this chapter is "the same" as its continuous counterpart. This is not a coincidence or a vague resemblance; it is a precise structural identity. Let us make this explicit.

What is identical

The following features of the discrete theory are literally the same as in the continuous theory, at the level of the abstract operator algebra:

(i) The index set for the hybrid basis is $\mathbb{N}_0^r \times \Lambda$, with $\mathbf{k} \in \mathbb{N}_0^r$ the one-sided grades and $\lambda \in \Lambda$ the spectral labels—the same in both theories.

(ii) The forward shift sends $e_{\mathbf{k},\lambda} \mapsto e_{\mathbf{k}+\mathbf{e}_i,\lambda}$ and the backward shift sends $e_{\mathbf{k},\lambda} \mapsto e_{\mathbf{k}-\mathbf{e}_i,\lambda}$ (or 0 at vacuum)—the same formulas in both theories.

(iii) The spectral multipliers act by $e_{\mathbf{k},\lambda} \mapsto \sigma(\lambda)\, e_{\mathbf{k},\lambda}$—the same in both theories.

(iv) The fundamental shift relations $C_iJ_i = I$, $J_iC_i = I - \Pi_i$, the higher-order identities, the defect localization, the uniqueness of the basis, and the invertibility of mixed operators—all carry over verbatim.

What differs

The differences lie entirely in the *function-space realization*:

(i) The one-sided monomial changes from $x^{k\alpha}/\Gamma(k\alpha+1)$ to $h_{k\alpha}(n) = n^{\overline{k\alpha}}/\Gamma(k\alpha+1)$.

(ii) The whole-space character changes from $e^{\langle \lambda, y\rangle}$ to λ^m.

(iii) The spectral eigenvalue of the Weyl operator changes from $\lambda_j^{\beta_j}$ to $(1-\lambda_j^{-1})^{\beta_j}$.

(iv) The spectral-set condition changes from $\operatorname{Re}\lambda_j > 0$ to $|\lambda_j| > 1$.

None of these differences affect the coefficient-space algebra.

Why the mirror is exact

The explanation is simple: both theories are realizations of the *same abstract coefficient-space model*. The abstract model consists of:

(a) a lattice $\mathbb{N}_0^r$ with coordinate forward and backward shifts,

(b) a spectral set Λ with pointwise multipliers,

(c) a tensor-product index set $\mathbb{N}_0^r \times \Lambda$ carrying finitely supported coefficient families.

Every algebraic identity proved in this chapter, and every algebraic identity proved in Chapter 7, is a statement about this abstract model. The continuous and discrete theories merely provide two concrete function-space incarnations.

This perspective will be formalized in Chapter 14, where the abstract coefficient-space framework is developed explicitly and the continuous and discrete transform models, completion theories, and semigroup constructions are shown to be instances of a single unified theory.

Example 11.4.1. The simplest case $r = s = 1$, $\alpha \in (0,1)$, $\Lambda = \{\lambda\} \subset \mathbb{E}$, has basis vectors $e_{k,\lambda}(n,m) = h_{k\alpha}(n)\,\lambda^m$ for $k \in \mathbb{N}_0$. The forward shift is $J = \nabla_0^{-\alpha}$, the backward shift is $C = {}^{\mathrm{C}}\nabla_0^{\alpha}$, and every discrete Weyl difference W^β multiplies by $(1-\lambda^{-1})^\beta$. The algebra is a one-dimensional unilateral shift tensored with a one-dimensional multiplier—the absolute simplest case, identical in structure to Example 7.3.13 of the continuous theory.

11.5 Why boundary geometry is even more visible on lattices

On a continuous domain like $(0,\infty)^r$, the boundary at $x_i = 0$ is approached as a limit: a vector "reaches the boundary" in the sense that its support extends to the origin. On a lattice, the boundary is not a limit but a concrete location: the hyperplane $\{n_i = 0\}$ consists of actual lattice points. This makes the boundary mechanism more visible and more concrete.

The vacuum is a lattice point. In the continuous theory, the vacuum vector $e_{\mathbf{0},\lambda}$ is the function $e^{\langle\lambda,y\rangle}$, which is defined at $x = 0$ only as a boundary value. In the discrete theory, the vacuum vector $e_{\mathbf{0},\lambda}(n,m) = h_0(n_1)\cdots h_0(n_r)\,\lambda^m = \lambda^m$ is a well-defined lattice function at every point of $\mathbb{N}_0^r \times \mathbb{Z}^s$—including the corner $n = \mathbf{0}$, where $h_0(0) = 1$. The boundary is not a limiting concept but a genuine lattice stratum.

Boundary layers are finite sets of lattice positions. The boundary-layer sector $\Pi_{<\mathbf{m}}\mathcal{H}^{\mathrm{disc,alg}}_{\alpha,\Lambda}$ consists of all basis vectors $e_{\mathbf{k},\lambda}$ with $\mathbf{k} \not\geq \mathbf{m}$. For fixed $\mathbf{m}$ and fixed λ, the set of such $\mathbf{k}$ is finite and can be enumerated explicitly. For instance, with $r = 2$ and $\mathbf{m} = (1,1)$, the boundary layer consists of $\mathbf{k} \in \{(0,0), (0,1), (0,2), \ldots, (1,0), (2,0), \ldots\}$—the "L-shaped" region of the lattice where at least one coordinate is zero. This combinatorial explicitness makes the discrete theory an excellent setting for concrete calculations and examples.

Trace words correspond to lattice paths. In the boundary-augmented theory that will be developed in Chapter 13, ordered trace words record the sequence of boundary walls that are touched as the Caputo operators lower grades to zero. On the lattice, this corresponds to a concrete path along the boundary: the word $(1,2)$ means that the path first reached the wall $\{n_1 = 0\}$ and then the wall $\{n_2 = 0\}$, while $(2,1)$ means the reverse order. The lattice makes the ordering of boundary events geometrically vivid.

Remark 11.5.1. The discrete setting is also natural for numerical computation. Every element of the algebraic module is a finite linear combination of sequences, and every operator action can be computed by finite summation. This makes the discrete hybrid algebra directly implementable, and the concrete examples of the next two chapters can be verified by direct computation on a computer.

Looking ahead

This chapter has established the discrete hybrid shift-spectral algebra on $\mathbb{N}^r \times \mathbb{Z}^s$: the rising-factorial–character basis, the commuting shift tuple, the diagonal spectral multipliers, the defect localization, the uniqueness theorem, and the algebraic inversion theory for mixed operators. The abstract operator algebra is identical to its continuous counterpart.

The next chapter develops the discrete Banach completion theory and the fiberwise Z-transform model (the discrete counterpart of Chapter 8). There, the weighted coefficient spaces, shift-admissible weights, geometric weights, discrete Mittag–Leffler eigenvectors, and the Banach-space inversion theorem will all be constructed in exact parallel with the continuous theory. After that, Chapter 13 will develop the discrete boundary-augmented ambient space and classify the maximal commuting sectors (the discrete counterpart of Chapter 9).

Chapter 12

Discrete Completion and the Z-Transform

Chapter 11 constructed the discrete hybrid shift-spectral algebra on $\mathbb{N}^r \times \mathbb{Z}^s$: a canonical rising-factorial–character basis on which the one-sided operators act as commuting shifts and the whole-space operators act as diagonal multipliers. That entire chapter was algebraic—every element was a finite linear combination, and every proof used only finitely many terms.

This chapter carries out the same passage from algebra to analysis that Chapter 8 accomplished for the continuous theory. We replace the algebraic direct sum by a weighted Banach space of coefficient sequences, show that all algebraic identities survive on the completion, prove that the discrete Mittag–Leffler generating vectors become genuine Banach-space elements, establish a Banach-space inversion theorem for mixed operators, and construct a fiberwise Z-transform model on a polydisk.

The reader who has absorbed Chapter 8 will find the present chapter structurally familiar. Every definition, every theorem, and every proof has a continuous counterpart. The differences are confined to the function-space realization (rising factorials replacing monomials, lattice characters replacing exponentials, the discrete Weyl symbol $(1 - \lambda_j^{-1})^\beta$ replacing λ_j^β) and to the spectral-set condition ($|\lambda_j| > 1$ replacing $\operatorname{Re} \lambda_j > 0$). The coefficient-space framework is identical.

The material follows the research paper AD09. Throughout, we fix

$$r \geq 1, \quad s \geq 1, \quad \boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_r) \in (0,1)^r, \quad \Lambda \subset \mathbb{E}^s, \quad \mathbb{E} := \{z \in \mathbb{C} : |z| > 1\}.$$

12.1 Weighted discrete completions

The construction is identical in structure to the continuous case of Chapter 8. We recall the algebraic core and equip it with a weighted ℓ^p norm.

Definition 12.1.1. A *weight* on $I := \mathbb{N}_0^r \times \Lambda$ is a function $\omega : I \to (0, \infty)$. Fix $1 \leq p < \infty$. For a finite-support element $u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}$ in the discrete hybrid module $\mathcal{H}_{\boldsymbol{\alpha},\Lambda}^{\mathrm{disc,alg}}$, define

$$\|u\|_{p,\omega} := \left(\sum |a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k}, \lambda)^p \right)^{1/p}.$$

The *discrete weighted Banach completion* $X_\omega^{p,\mathrm{disc}}$ is the completion of the algebraic core under this norm.

As in the continuous theory, $X_\omega^{p,\mathrm{disc}}$ is canonically isometric to the weighted sequence space $\ell_\omega^p(I)$, via the map that sends each element to its coefficient family. The algebraic core corresponds to the finitely supported families, which are dense for $1 \le p < \infty$.

Definition 12.1.2. The weight ω is *shift-admissible* if, for every $i \in \{1, \dots, r\}$,

$$L_i := \sup_{(\mathbf{k},\lambda)} \frac{\omega(\mathbf{k}+\mathbf{e}_i,\lambda)}{\omega(\mathbf{k},\lambda)} < \infty \qquad \text{and} \qquad R_i := \sup_{\substack{(\mathbf{k},\lambda)\\ k_i \ge 1}} \frac{\omega(\mathbf{k}-\mathbf{e}_i,\lambda)}{\omega(\mathbf{k},\lambda)} < \infty.$$

Theorem 12.1.3 (Bounded discrete shift extensions)**.** *Let ω be shift-admissible. Then each J_i and each C_i extends uniquely to a bounded operator on $X_\omega^{p,\mathrm{disc}}$ with exact norms $\|J_i\| = L_i$ and $\|C_i\| = R_i$.*

Proof. The proof is identical to the continuous case (Theorem 8.3.3 of Chapter 8): the upper bounds follow from the weight-ratio estimates on the algebraic core, and the lower bounds from testing on single basis vectors. □

Each vacuum projection Π_i and each boundary-layer projection $\Pi_{<\mathbf{m}}$ extends to a contractive projection on the completion, because restricting a sum to a subset of its terms can only decrease the ℓ^p norm.

Corollary 12.1.4 (Discrete algebraic identities on the completion)**.** *For every $i \in \{1, \dots, r\}$,*

$$C_i J_i = I, \qquad J_i C_i = I - \Pi_i, \qquad [C_i, J_i] = \Pi_i$$

on $X_\omega^{p,\mathrm{disc}}$. For all $i \ne j$, $J_i J_j = J_j J_i$, $C_i C_j = C_j C_i$, and $C_i J_j = J_j C_i$. The higher-order identities $C^{\mathbf{m}} J^{\mathbf{m}} = I$ and $J^{\mathbf{m}} C^{\mathbf{m}} = I - \Pi_{<\mathbf{m}}$ hold for every $\mathbf{m} \in \mathbb{N}_0^r$.

Proof. All identities hold on the dense algebraic core by Chapter 11. Every operator is bounded, so the identities extend by continuity. □

12.2 Discrete spectral multipliers and closed operator calculus

The spectral block on Λ gives a diagonal multiplier calculus, exactly as in the continuous theory.

Theorem 12.2.1 (Discrete isometric multiplier calculus)**.** *For every bounded $\sigma : \Lambda \to \mathbb{C}$, the multiplier M_σ extends to a bounded operator on $X_\omega^{p,\mathrm{disc}}$ with $\|M_\sigma\| = \|\sigma\|_\infty$. The map $\sigma \mapsto M_\sigma$ is an isometric unital algebra homomorphism. Every bounded multiplier commutes with every J_i, C_i, Π_i, and $\Pi_{<\mathbf{m}}$.*

Proof. The diagonal action on coefficients gives $\|M_\sigma u\|_{p,\omega} \le \|\sigma\|_\infty \|u\|_{p,\omega}$, and testing on a single basis vector gives the reverse inequality. Commutativity follows as in the continuous case because the shift operators change only $\mathbf{k}$ while M_σ changes only the scalar attached to λ. □

For arbitrary (unbounded) symbols, the multiplier M_σ with maximal domain

$$\mathcal{D}(M_\sigma) := \left\{ u \in X_\omega^{p,\mathrm{disc}} : \sum |\sigma(\lambda)\, a_{\mathbf{k},\lambda}|^p\, \omega(\mathbf{k},\lambda)^p < \infty \right\}$$

is densely defined and closed, by the same continuity-of-coordinate-functionals argument as in Theorem 8.4.5 of Chapter 8.

Definition 12.2.2. The *standard discrete Weyl operator on* $X_\omega^{p,\mathrm{disc}}$ is the closed multiplier

$$W^{\boldsymbol{\beta}} := M_{\omega_{\boldsymbol{\beta}}}, \qquad \omega_{\boldsymbol{\beta}}(\lambda) := \prod_{j=1}^{s}(1-\lambda_j^{-1})^{\beta_j},$$

with maximal domain.

Corollary 12.2.3. *For every $\boldsymbol{\beta} \in [0,\infty)^s$, the operator $W^{\boldsymbol{\beta}}$ is densely defined and closed on $X_\omega^{p,\mathrm{disc}}$. It is bounded if and only if $\omega_{\boldsymbol{\beta}}$ is bounded on Λ. The shift operators J_i and C_i preserve $\mathcal{D}(W^{\boldsymbol{\beta}})$ and commute with $W^{\boldsymbol{\beta}}$ on that domain.*

Proof. Domain invariance uses the weight-ratio bounds L_i and R_i exactly as in Corollary 8.4.7 of Chapter 8, with the discrete Weyl symbol $\omega_{\boldsymbol{\beta}}$ in place of $\lambda^{\boldsymbol{\beta}}$. □

Remark 12.2.4. The only difference from the continuous spectral theory is the symbol: in the continuous case the Weyl eigenvalue is $\lambda^{\boldsymbol{\beta}}$, in the discrete case it is $\omega_{\boldsymbol{\beta}}(\lambda) = \prod_j (1-\lambda_j^{-1})^{\beta_j}$. The abstract framework (bounded multipliers are isometric, unbounded symbols give closed operators, shifts preserve domains and commute) is identical.

12.3 Geometric weights and discrete generating vectors

Definition 12.3.1. The *discrete geometric weight* is

$$\omega_{\rho,\eta}(\mathbf{k},\lambda) := \eta(\lambda)\,\boldsymbol{\rho}^{\mathbf{k}},$$

with $\boldsymbol{\rho} \in (0,\infty)^r$ and $\eta : \Lambda \to (0,\infty)$. Write $X_{\rho,\eta}^{p,\mathrm{disc}}$ for the corresponding completion.

Proposition 12.3.2. *The discrete geometric weight is shift-admissible with*

$$\|J_i\| = \rho_i, \qquad \|C_i\| = \rho_i^{-1}, \qquad \|J_i\| \cdot \|C_i\| = 1.$$

Proof. The weight ratios $\omega(\mathbf{k}+\mathbf{e}_i)/\omega(\mathbf{k}) = \rho_i$ and $\omega(\mathbf{k}-\mathbf{e}_i)/\omega(\mathbf{k}) = \rho_i^{-1}$ are constant, as in the continuous case (Proposition 8.5.2 of Chapter 8). □

Discrete Mittag–Leffler vectors

Definition 12.3.3. Fix $\lambda \in \Lambda$ and $\boldsymbol{\zeta} \in \mathbb{C}^r$ with $|\zeta_i| < \rho_i^{-1}$ for every i. The *discrete generating vector* is

$$E_{\zeta,\lambda}^{\mathrm{disc}} := \sum_{\mathbf{k}\in\mathbb{N}_0^r} \boldsymbol{\zeta}^{\mathbf{k}}\, e_{\mathbf{k},\lambda}.$$

Proposition 12.3.4 (Convergence)**.** *The vector* $E^{\mathrm{disc}}_{\zeta,\lambda}$ *belongs to* $X^{p,\mathrm{disc}}_{\rho,\eta}$ *with*

$$\|E^{\mathrm{disc}}_{\zeta,\lambda}\|^p_{p,\omega} = \eta(\lambda)^p \prod_{i=1}^{r} \frac{1}{1 - |\zeta_i|^p \rho_i^p}.$$

Proof. The norm is $\eta(\lambda)^p \sum_{\mathbf{k}} |\boldsymbol{\zeta}^{\mathbf{k}}|^p \boldsymbol{\rho}^{p\mathbf{k}} = \eta(\lambda)^p \prod_i (1-|\zeta_i|^p \rho_i^p)^{-1}$, a product of convergent geometric series. □

Theorem 12.3.5 (Discrete eigenvectors)**.** *Let* $|\zeta_i| < \rho_i^{-1}$ *for all i. Then:*

(i) $C_i\, E^{\mathrm{disc}}_{\zeta,\lambda} = \zeta_i\, E^{\mathrm{disc}}_{\zeta,\lambda}$ *for every i.*

(ii) $M_\sigma\, E^{\mathrm{disc}}_{\zeta,\lambda} = \sigma(\lambda)\, E^{\mathrm{disc}}_{\zeta,\lambda}$ *for every bounded* σ*.*

(iii) $W^{\boldsymbol{\beta}}\, E^{\mathrm{disc}}_{\zeta,\lambda} = \omega_{\boldsymbol{\beta}}(\lambda)\, E^{\mathrm{disc}}_{\zeta,\lambda}$ *for every* $\boldsymbol{\beta} \in [0,\infty)^s$*.*

Proof. Since C_i is bounded, we apply it term by term to the convergent series. Reindexing $\mathbf{m} = \mathbf{k} - \mathbf{e}_i$ after dropping the $k_i = 0$ terms gives $C_i\, E^{\mathrm{disc}} = \zeta_i\, E^{\mathrm{disc}}$. Parts (ii) and (iii) follow from the single-fiber support of the vector. □

Pointwise realization

Corollary 12.3.6. *If in addition* $|\zeta_i| < 1$ *for every i, then for every* $(n,m) \in \mathbb{N}_0^r \times \mathbb{Z}^s$,

$$E^{\mathrm{disc}}_{\zeta,\lambda}(n,m) \;=\; \left(\prod_{i=1}^{r} \mathcal{E}^{\nabla}_{\alpha_i}(n_i;\zeta_i)\right) \lambda^m,$$

where $\mathcal{E}^{\nabla}_{\alpha}(n;\zeta) := \sum_{q=0}^{\infty} \zeta^q\, h_{q\alpha}(n)$ *is the one-variable discrete generating kernel.*

Proof. The condition $|\zeta_i| < 1$ ensures pointwise absolute convergence of each one-variable series (since $h_{q\alpha}(n) = O(q^{n-1})$ for fixed n by gamma-ratio asymptotics). The multi-variable series then factors coordinatewise. □

Remark 12.3.7. The Banach-space vector $E^{\mathrm{disc}}_{\zeta,\lambda}$ exists as soon as $|\zeta_i| < \rho_i^{-1}$, while the pointwise discrete kernel series converges only when additionally $|\zeta_i| < 1$. When some $\rho_i < 1$, the completion contains genuine abstract eigenvectors whose pointwise realization on the physical lattice does not converge. This subtlety—abstract Banach-space elements without pointwise meaning—has no counterpart in the continuous theory, where the Mittag–Leffler function converges everywhere on $(0,\infty)$.

Example 12.3.8. In one variable ($r = 1$, $s = 0$), with $\alpha = 1/2$ and $\rho = 2$, the generating eigenvector $E_\zeta = \sum \zeta^k e_k$ exists for $|\zeta| < 1/2$ and satisfies $C_{1/2}\, E_\zeta = \zeta\, E_\zeta$. Its pointwise realization is $E_\zeta(n) = \mathcal{E}^{\nabla}_{1/2}(n;\zeta)$, which converges for $|\zeta| < 1$. Since $\rho^{-1} = 1/2 < 1$, every eigenvector in the completion also has a convergent pointwise realization. But if instead $\rho = 1/2$ (so $\rho^{-1} = 2$), then eigenvectors with $1 \le |\zeta| < 2$ exist in the Banach space but cannot be evaluated pointwise on the lattice.

12.4 The fiberwise Z-transform model

For geometric completions, the coefficient sequences assemble into holomorphic functions on a polydisk, exactly as in the continuous theory of Section 8.6 in Chapter 8. The resulting transform is called the *fiberwise Z-transform*.

Definition 12.4.1. The *reciprocal polydisk* is $\mathbb{D}_\rho := \{z \in \mathbb{C}^r : |z_i| < \rho_i \text{ for all } i\}$.

Definition 12.4.2 (Fiberwise Z-transform). For $u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda} \in X_{\rho,\eta}^{p,\mathrm{disc}}$ and fixed $\lambda \in \Lambda$, define

$$(\mathcal{Z}u)(z,\lambda) := \sum_{\mathbf{k}\in\mathbb{N}_0^r} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}}, \qquad z \in \mathbb{D}_\rho.$$

Proposition 12.4.3. *For every $u \in X_{\rho,\eta}^{p,\mathrm{disc}}$ and every $\lambda \in \Lambda$, the series converges absolutely on $\mathbb{D}_\rho$ and defines a holomorphic function. The transform $\mathcal{Z}$ is injective.*

Proof. The argument is identical to the continuous case (Proposition 8.6.3 of Chapter 8): Hölder's inequality controls the absolute convergence, Weierstrass's theorem gives holomorphy, and vanishing of Taylor coefficients gives injectivity. ◻

Theorem 12.4.4 (Discrete operator intertwining). *For $u \in X_{\rho,\eta}^{p,\mathrm{disc}}$:*

(i) Forward shift becomes multiplication: $(\mathcal{Z}(J_i u))(z,\lambda) = z_i\,(\mathcal{Z}u)(z,\lambda)$.

(ii) Spectral multiplier passes through: $(\mathcal{Z}(M_\sigma u))(z,\lambda) = \sigma(\lambda)\,(\mathcal{Z}u)(z,\lambda)$.

(iii) Backward shift becomes a quotient operator:

$$(\mathcal{Z}(C_i u))(z,\lambda) = \frac{(\mathcal{Z}u)(z,\lambda) - (\mathcal{Z}u)(z|_{z_i=0},\lambda)}{z_i}$$

for $z_i \neq 0$, extending holomorphically across $z_i = 0$.

(iv) Discrete Weyl passes through: *for $u \in \mathcal{D}(W^\beta)$, $(\mathcal{Z}(W^\beta u))(z,\lambda) = \omega_\beta(\lambda)\,(\mathcal{Z}u)(z,\lambda)$.*

Proof. Each identity follows from the basis-level actions by the same rearrangements as in Theorem 8.6.4 of Chapter 8. For part (i): J_i shifts the $\mathbf{k}$-index by $\mathbf{e}_i$, so $\mathcal{Z}(J_i u) = z_i\,\mathcal{Z}(u)$. For part (iii): C_i drops all terms with $k_i = 0$ and shifts the rest down by $\mathbf{e}_i$; dividing by z_i and subtracting the $k_i = 0$ component gives the quotient formula. For part (iv): W^β multiplies each coefficient by $\omega_\beta(\lambda)$. ◻

Remark 12.4.5. In the transform picture, $X_{\rho,\eta}^{p,\mathrm{disc}}$ becomes a mixed Hardy-type space of holomorphic functions on $\mathbb{D}_\rho$, parametrized fiberwise by $\lambda \in \Lambda$. The intertwining rules are formally identical to those of the continuous holomorphic model (Chapter 8): $J_i \leftrightarrow$ multiplication by z_i, $C_i \leftrightarrow$ backward quotient, $M_\sigma \leftrightarrow$ fiberwise multiplication by $\sigma(\lambda)$. The only discrete fingerprint is the Weyl eigenvalue $\omega_\beta(\lambda)$ in place of λ^β.

Example 12.4.6. Under $\mathscr{Z}$, the discrete generating vector $E^{\mathrm{disc}}_{\zeta,\lambda} = \sum \zeta^{\mathbf{k}} e_{\mathbf{k},\lambda}$ maps to

$$(\mathscr{Z} E^{\mathrm{disc}}_{\zeta,\lambda})(z,\lambda) = \prod_{i=1}^{r} \frac{1}{1-\zeta_i z_i},$$

the same product of Cauchy kernels as in the continuous transform (Example 8.6.6 of Chapter 8). This identity illustrates that the Cauchy kernel is the universal transform-side representation of the generating eigenvector, independent of the continuous or discrete realization.

12.5 The role of AD09

This chapter stands in exactly the same structural position as Chapter 8 in the continuous theory. The results organize into the same four themes.

The operators have norms. On the discrete completion, the forward and backward shifts J_i, C_i are bounded with exact norms determined by the weight ratios. For geometric weights, $\|J_i\| = \rho_i$ and $\|C_i\| = \rho_i^{-1}$, achieving the optimal product $\|J_i\| \cdot \|C_i\| = 1$. Spectral multipliers are isometric for bounded symbols and closed for arbitrary symbols. The discrete Weyl operators W^{β} are closed diagonal operators with symbol $\omega_{\beta}(\lambda)$.

The generating eigenvectors exist. The discrete Mittag–Leffler vectors $E^{\mathrm{disc}}_{\zeta,\lambda}$ are genuine Banach-space elements, not formal series. They are joint eigenvectors of the commuting Caputo tuple and of every spectral multiplier. When the parameters also satisfy $|\zeta_i| < 1$, the vectors can be evaluated pointwise on the lattice via the discrete generating kernel $\mathcal{E}^{\nabla}_{\alpha}$.

A Z-transform model has emerged. The fiberwise Z-transform converts the geometric completion into a Hardy-type space on a polydisk, with operators represented as multiplication (J_i), backward quotient (C_i), and fiberwise scalar multiplication (M_σ, W^{β}). The transform is injective and its intertwining rules are formally identical to those of the continuous holomorphic model.

Inversion is genuinely analytic. On the algebraic core, mixed lower-triangular operators are inverted by local nilpotence. On the completion, this fails, and one needs a Banach-space perturbation argument.

Theorem 12.5.1 (Discrete Banach-space invertibility). *Let ω be shift-admissible and let $T = M_{\sigma_{\mathbf{0}}} + \sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} C^{\mathbf{m}} M_{\sigma_{\mathbf{m}}}$ on $X^{p,\mathrm{disc}}_{\omega}$ with each $\sigma_{\mathbf{m}} \in \ell^{\infty}(\Lambda)$. If*

(a) $m_0 := \inf_{\lambda} |\sigma_{\mathbf{0}}(\lambda)| > 0$, and

(b) $\sum_{\mathbf{0}<\mathbf{m}\le\mathbf{M}} R^{\mathbf{m}} \|\sigma_{\mathbf{m}}\|_{\infty} < m_0$,

then T is boundedly invertible with $T^{-1} = \sum_{q=0}^{\infty} (-K)^q M_{\sigma_{\mathbf{0}}^{-1}}$, where the Neumann series converges in operator norm.

Proof. Factor $T = M_{\sigma_{\mathbf{0}}}(I + K)$ with $K = M_{\sigma_{\mathbf{0}}^{-1}} \sum C^{\mathbf{m}} M_{\sigma_{\mathbf{m}}}$. Then $\|K\| < 1$ by the same computation as in Theorem 8.7.1 of Chapter 8, and the Neumann lemma applies. □

12.6 Concrete examples

We close with two small worked examples that illustrate the discrete theory in its simplest instances.

Example 12.6.1 (One-variable discrete eigenvector)**.** Consider $r = 1$, $s = 0$, $\alpha = 1/2$, and the geometric weight with $\rho = 2$. The Banach space is $X_2^{p,\mathrm{disc}} = \ell^p_{2^k}(\mathbb{N}_0)$, with $\|J_{1/2}\| = 2$ and $\|C_{1/2}\| = 1/2$.

The eigenvector equation $C_{1/2}\,u = \zeta\,u$ for $|\zeta| < 1/2$ has the unique solution (up to scalar) $E_\zeta = \sum_{k=0}^{\infty} \zeta^k\, e_k$ in the completion. Under the Z-transform, this becomes $\mathscr{Z}(E_\zeta)(z) = 1/(1-\zeta z)$ on the disk $|z| < 2$, and the eigenvector equation becomes the elementary identity $(1/(1-\zeta z)-1)/z = \zeta/(1-\zeta z)$.

Now consider the operator $T = I + a\,C_{1/2}$ for $a \in \mathbb{C}$. On the algebraic core, $T^{-1} = \sum_{q=0}^{\infty} (-a)^q C_{1/2}^q$ terminates on every finite-support vector. On the completion, the Neumann series converges whenever $|a| \cdot \|C_{1/2}\| = |a|/2 < 1$, i.e., $|a| < 2$. Under $\mathscr{Z}$, the operator T acts by multiplication by $1 + a \cdot (\varphi(z) - \varphi(0))/z$ (where $\varphi = \mathscr{Z}(u)$), so solving $Tu = f$ reduces to dividing $\mathscr{Z}(f)$ by a known function.

Example 12.6.2 (Mixed discrete equation with Weyl symbol)**.** Take $r = s = 1$, $\alpha = 1/3$, $\Lambda = \{3\} \subset \mathbb{E}$, and a geometric weight with $\rho = 1$, $\eta(3) = 1$. The discrete Weyl symbol at $\lambda = 3$ is $(1 - 3^{-1})^\beta = (2/3)^\beta$.

Consider the operator $T = W^\beta + a\,C_{1/3}\,W^\gamma$ on the completion. Here $\sigma_0(3) = (2/3)^\beta$ and $\sigma_1(3) = a(2/3)^\gamma$. Since Λ is a single point, condition (a) of Theorem 12.5.1 is $(2/3)^\beta > 0$ (automatic), and condition (b) is

$$\|C_{1/3}\| \cdot |a| \cdot (2/3)^\gamma = 1 \cdot |a| \cdot (2/3)^\gamma < (2/3)^\beta,$$

so T is invertible whenever $|a| < (2/3)^{\beta-\gamma}$. The inverse is a convergent Neumann series whose terms involve successively higher powers of $C_{1/3}$.

Remark 12.6.3. The discrete setting lends itself to direct numerical verification. Every element of the algebraic core is a finite sequence, every operator acts by a finite sum, and the generating vectors can be truncated to any desired precision. The reader is encouraged to implement the examples above and check the eigenvector and inversion formulas by direct computation.

Looking ahead

This chapter has completed the discrete analytic theory: weighted Banach completions, bounded and closed operators, geometric weights, discrete generating eigenvectors, the fiberwise Z-transform, and Banach-space inversion. The structure is identical to the continuous theory of Chapter 8, with only the function-space realization and the spectral symbol changed.

The next chapter develops the discrete boundary-augmented theory and classifies the maximal commuting sectors for the discrete extended Caputo tuple—the discrete counterpart of Chapter 9. After that, Chapter 14 will unify the continuous and discrete completion theories in a single abstract coefficient-space framework, where the common structure becomes manifest and the transform models, optimal weights, and semigroup generation are treated once for both settings simultaneously.

Chapter 13

Discrete Boundary Structure and Maximal Commuting Sectors

Chapter 9 asked what happens to the continuous hybrid theory when one enlarges the canonical completion by adjoining boundary-trace layers that record the order of boundary extractions. The answer was a precise classification: the maximal graded invariant sector on which the extended Caputo tuple remains commuting consists of the canonical completion plus all trace blocks whose free rank is at most one. Noncommutativity is entirely boundary-generated, and the spectral block plays no role in it.

This chapter carries out the same analysis in the discrete setting. The canonical discrete Banach completion of Chapter 12 is embedded in a larger boundary-augmented space, the extended discrete Caputo operators are defined, their commutator is computed, and the maximal commuting sector is classified. The result is identical in structure to the continuous case: the obstruction to commutativity occurs only at simultaneous vacuum in two or more free one-sided lattice coordinates, and the maximal commuting sector is determined by the same free-rank condition $q(w) \leq 1$.

The reader who has absorbed Chapter 9 will find the present chapter structurally familiar. Every definition, every theorem, and every proof follows the same combinatorial logic. The reason is that the mechanism of noncommutativity—ordered boundary creation at simultaneous vacuum—is a property of the coefficient-space lattice $\mathbb{N}_0^r$, not of the specific function-space realization. Whether the underlying functions are Gamma-normalized monomials or normalized rising factorials, the boundary geometry is the same.

The material follows the research paper AD10. Throughout, we fix

$$r \geq 1, \quad s \geq 1, \quad \boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_r) \in (0,1)^r, \quad \Lambda \subset \mathbb{E}^s, \quad \mathbb{E} := \{z \in \mathbb{C} : |z| > 1\},$$

and geometric weights $\boldsymbol{\rho} = (\rho_1, \ldots, \rho_r) \in (0,\infty)^r$ and a positive spectral function $\eta : \Lambda \to (0,\infty)$.

13.1 The boundary-augmented discrete ambient space

On the canonical discrete Banach completion $X_{\boldsymbol{\rho},\eta}^{p,\mathrm{disc}}$ of Chapter 12, the partial Caputo nabla difference C_i^{disc} acts as a bounded backward shift: it lowers the i-th one-sided grade by one and annihilates any basis vector with $k_i = 0$. This vacuum annihilation is the discrete counterpart of the continuous phenomenon studied in Chapter 9—and, as in the continuous case, one can ask whether there is a natural larger space in which the boundary event at $n_i = 0$ is *recorded* rather than erased.

The answer follows exactly the same path as in the continuous theory. We adjoin formal trace layers indexed by ordered words.

Recollection of the canonical discrete completion

We recall from Chapter 12 that the canonical discrete hybrid basis vectors are

$$e_{\mathbf{k},\lambda}(n,m) = \left(\prod_{i=1}^{r} h_{k_i\alpha_i}(n_i)\right)\lambda^m, \qquad (\mathbf{k},\lambda) \in \mathbb{N}_0^r \times \Lambda,$$

where $h_\beta(n) = n^{\overline{\beta}}/\Gamma(\beta+1)$ is the normalized rising-factorial monomial and λ^m is the lattice character. The geometric norm on the finite-support span is

$$\left\|\sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda}\right\|_{p,\boldsymbol{\rho},\eta}^p = \sum_{\mathbf{k},\lambda} |a_{\mathbf{k},\lambda}|^p\, \eta(\lambda)^p\, \boldsymbol{\rho}^{p\mathbf{k}},$$

where $\boldsymbol{\rho}^{\mathbf{k}} = \rho_1^{k_1}\cdots\rho_r^{k_r}$. The Banach completion under this norm is denoted $X_{\boldsymbol{\rho},\eta}^{p,\mathrm{disc}}$.

On this space, the partial discrete Caputo operators act as a commuting family of weighted backward shifts. The question we now address is: what is the *maximal* graded sector, inside a natural ambient space, on which this commutativity survives?

Ordered trace words

The trace layers are parametrized by the same combinatorial data as in the continuous case (Definition 9.2.1 of Chapter 9). We recall the essential definitions in the present notation.

Definition 13.1.1. An *ordered trace word* in the alphabet $\{1,\dots,r\}$ is a finite sequence $w = (i_1,\dots,i_m)$ of distinct elements. The empty word is $\varnothing$. The set of all such words is

$$\mathcal{W}_r := \{\varnothing\} \cup \big\{(i_1,\dots,i_m) : 1 \le m \le r,\ i_\nu \in \{1,\dots,r\}\ \text{distinct}\big\}.$$

For each word $w = (i_1,\dots,i_m)$:

(i) the *length* is $\ell(w) := m$ (with $\ell(\varnothing) := 0$),

(ii) the *support* is $\operatorname{supp}(w) := \{i_1,\dots,i_m\}$,

(iii) the *free set* is $\mathrm{F}(w) := \{1,\dots,r\} \setminus \operatorname{supp}(w)$,

(iv) the *free rank* is $q(w) := |\mathrm{F}(w)| = r - \ell(w)$.

If $i \in \mathrm{F}(w)$, write wi for the word obtained by appending i to w.

Remark 13.1.2. The word set $\mathcal{W}_r$ and the notions of support, free set, and free rank are identical to those of the continuous theory (Definition 9.2.1 in Chapter 9). This is no coincidence: the combinatorics of ordered boundary extraction depend only on the number r of one-sided coordinates, not on whether those coordinates carry continuous or discrete grading.

Discrete trace blocks

Definition 13.1.3. For each word $w \in \mathcal{W}_r$, the *algebraic discrete trace block* is

$$\mathcal{T}_w^{\mathrm{alg,disc}} := \bigoplus_{\mathbf{k}\in\mathbb{N}_0^{\mathrm{F}(w)},\, \lambda\in\Lambda} \mathbb{C}\, t_{w,\mathbf{k},\lambda},$$

where $t_{w,\mathbf{k},\lambda}$ is a purely formal basis symbol. The index $\mathbf{k} \in \mathbb{N}_0^{\mathrm{F}(w)}$ records the residual grades in the free one-sided coordinates, and $\lambda \in \Lambda$ is the spectral label.

The symbol $t_{w,\mathbf{k},\lambda}$ should be read as an ordered lattice-boundary trace state. The word w records which one-sided lattice coordinates have already been extracted at $n_i = 0$ and in what order; the residual multi-index $\mathbf{k}$ records the remaining grades in the coordinates not yet extracted; and λ is the spectral fiber, unchanged from the canonical setting.

Definition 13.1.4. The *boundary-augmented algebraic discrete ambient space* is

$$\mathfrak{H}^{\mathrm{alg,disc}} := \mathcal{H}_{\alpha,\Lambda}^{\mathrm{alg,disc}} \oplus \bigoplus_{w\in\mathcal{W}_r} \mathcal{T}_w^{\mathrm{alg,disc}},$$

where $\mathcal{H}_{\alpha,\Lambda}^{\mathrm{alg,disc}}$ is the canonical discrete hybrid module of Chapter 11.

The boundary-augmented discrete Banach space

We complete the algebraic ambient space under a weighted norm that extends the canonical geometric norm by assigning controlled weight to the trace layers.

Definition 13.1.5. Fix $1 \le p < \infty$ and a *trace-weight parameter* $\tau > 0$. For a finite-support vector

$$u = \sum a_{\mathbf{k},\lambda}\, e_{\mathbf{k},\lambda} + \sum_{w\in\mathcal{W}_r} \sum b_{w,\mathbf{k},\lambda}\, t_{w,\mathbf{k},\lambda} \;\in\; \mathfrak{H}^{\mathrm{alg,disc}},$$

define

$$\|u\|_{p,\boldsymbol{\rho},\tau,\eta}^p := \sum_{\mathbf{k},\lambda} |a_{\mathbf{k},\lambda}|^p\, \eta(\lambda)^p\, \boldsymbol{\rho}^{p\mathbf{k}} + \sum_{w\in\mathcal{W}_r} \sum_{\mathbf{k},\lambda} |b_{w,\mathbf{k},\lambda}|^p\, \tau^{p(\ell(w)+1)}\, \eta(\lambda)^p\, \boldsymbol{\rho}_w^{p\mathbf{k}},$$

where $\boldsymbol{\rho}_w^{\mathbf{k}} := \prod_{i\in\mathrm{F}(w)} \rho_i^{k_i}$ (the empty product equals 1). The *boundary-augmented discrete completion* is the Banach space

$$\mathfrak{X}_{\boldsymbol{\rho},\tau,\eta}^{p,\mathrm{disc}} := \overline{\mathfrak{H}^{\mathrm{alg,disc}}} \quad \text{under } \|\cdot\|_{p,\boldsymbol{\rho},\tau,\eta}.$$

For each word $w \in \mathcal{W}_r$, let $T_w^{p,\mathrm{disc}}$ denote the closure of $\mathcal{T}_w^{\mathrm{alg,disc}}$ in the ambient norm.

The factor $\tau^{\ell(w)+1}$ plays the same role as in the continuous theory: it measures the cost of entering a trace layer. The exponent $\ell(w)+1$ ensures that even the empty-word trace block $T^{p,\mathrm{disc}}_{\varnothing}$ sits one weight level above the canonical completion.

Proposition 13.1.6. *There is an isometric direct-sum decomposition*

$$\mathfrak{x}^{p,\mathrm{disc}}_{\boldsymbol{\rho},\tau,\eta} = \mathrm{X}^{p,\mathrm{disc}}_{\boldsymbol{\rho},\eta} \oplus \bigoplus_{w\in\mathcal{W}_r} T^{p,\mathrm{disc}}_{w}.$$

In particular, the canonical discrete completion $\mathrm{X}^{p,\mathrm{disc}}_{\boldsymbol{\rho},\eta}$ *embeds isometrically as a complemented subspace.*

Proof. The basis families $\{e_{\mathbf{k},\lambda}\}$ and $\{t_{w,\mathbf{k},\lambda}\}$ are pairwise disjoint, and the ambient norm is the ℓ^p-sum of their individual block norms. Since the word set $\mathcal{W}_r$ is finite (it has at most $\sum_{m=0}^{r} r!/(r-m)!$ elements), the completion is a finite ℓ^p-direct sum of block completions. The first summand is $\mathrm{X}^{p,\mathrm{disc}}_{\boldsymbol{\rho},\eta}$ by construction. □

Remark 13.1.7. On the lattice $\mathbb{N}_0^r$, the boundary faces are especially concrete. The face $\{n_i = 0\}$ is a codimension-one sublattice, and the intersection $\{n_i = n_j = 0\}$ is a codimension-two sublattice. A trace word w of length m corresponds to a codimension-m boundary stratum in which exactly $q(w) = r - m$ lattice coordinates remain "active" (with positive grades). The lattice geometry makes the boundary strata more visible than in the continuous setting, but the algebraic structure of the trace blocks is identical.

13.2 Ordered words and extended discrete Caputo operators

We now define the extended discrete Caputo operators on the boundary-augmented space. On the canonical block, they reproduce the commuting backward shifts of Chapter 12. On the trace blocks, they act by grade-lowering whenever the relevant coordinate has positive grade, and by ordered boundary creation when the grade reaches zero.

Definition 13.2.1. For each $i \in \{1,\dots,r\}$, define the *extended discrete Caputo operator* $C_i^{\mathrm{disc}} : \mathfrak{H}^{\mathrm{alg,disc}} \longrightarrow \mathfrak{H}^{\mathrm{alg,disc}}$ on basis vectors as follows.

On the canonical block:

$$C_i^{\mathrm{disc}}\, e_{\mathbf{k},\lambda} := \begin{cases} 0, & k_i = 0, \\ e_{\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i \geq 1. \end{cases}$$

This is the canonical discrete backward shift from Chapter 12.

On a trace basis vector $t_{w,\mathbf{k},\lambda}$*:*

$$C_i^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} := \begin{cases} 0, & i \notin \mathrm{F}(w), \\ t_{w,\mathbf{k}-\mathbf{e}_i,\lambda}, & i \in \mathrm{F}(w),\ k_i \geq 1, \\ t_{wi,\mathbf{k}^{\widehat{i}},\lambda}, & i \in \mathrm{F}(w),\ k_i = 0, \end{cases}$$

where $\mathbf{k}^{\widehat{i}} \in \mathbb{N}_0^{\mathrm{F}(wi)}$ denotes the residual multi-index obtained by removing the i-th component.

The three cases have the same interpretations as in the continuous theory (Definition 9.3.1 of Chapter 9):

(i) If i already appears in the word w (i.e., $i \notin \mathrm{F}(w)$), then the i-th lattice boundary has already been recorded. The operator annihilates.

(ii) If i is free and $k_i \geq 1$, then C_i^{disc} lowers the residual grade by one, exactly as in the canonical shift model. The trace word is unchanged.

(iii) If i is free and $k_i = 0$, then the lattice coordinate n_i is at the face $\{n_i = 0\}$. The operator records this boundary event by appending i to the trace word, producing the longer word wi, and removing i from the set of free coordinates.

The third case—ordered boundary creation at vacuum—is again the sole source of noncommutativity.

Definition 13.2.2. For a bounded function $\sigma : \Lambda \to \mathbb{C}$, the *ambient discrete spectral multiplier* acts by

$$M_\sigma^{\mathrm{disc}}\, e_{\mathbf{k},\lambda} := \sigma(\lambda)\, e_{\mathbf{k},\lambda}, \qquad M_\sigma^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} := \sigma(\lambda)\, t_{w,\mathbf{k},\lambda}.$$

Theorem 13.2.3 (Boundedness of the extended tuple)**.** *Each operator C_i^{disc} extends to a bounded operator on $\mathfrak{X}_{\rho,\tau,\eta}^{p,\mathrm{disc}}$ with norm*

$$\|C_i^{\mathrm{disc}}\| = \max\{\rho_i^{-1},\ \tau\}.$$

Each bounded multiplier M_σ^{disc} extends to a bounded operator with $\|M_\sigma^{\mathrm{disc}}\| = \|\sigma\|_\infty$. Moreover,

$$C_i^{\mathrm{disc}}\, M_\sigma^{\mathrm{disc}} = M_\sigma^{\mathrm{disc}}\, C_i^{\mathrm{disc}}$$

on the ambient space.

Proof. On each basis vector, the ratio $\|C_i^{\mathrm{disc}}\,\xi\|/\|\xi\|$ takes one of three values: 0 (annihilation), ρ_i^{-1} (grade-lowering), or τ (trace creation). The last value arises because the norm factor changes from $\tau^{\ell(w)+1}$ to $\tau^{\ell(w)+2}$ while $\rho_i^{k_i}$ drops out with $k_i = 0$. Hence the operator norm is $\max\{\rho_i^{-1}, \tau\}$.

For the multiplier, $\|M_\sigma^{\mathrm{disc}}\,\xi\|/\|\xi\| = |\sigma(\lambda)|$ on every basis vector ξ with spectral label λ, giving $\|M_\sigma^{\mathrm{disc}}\| = \|\sigma\|_\infty$.

The identity $C_i^{\mathrm{disc}} M_\sigma^{\mathrm{disc}} = M_\sigma^{\mathrm{disc}} C_i^{\mathrm{disc}}$ holds because C_i^{disc} changes only the trace word and the one-sided grades, while M_σ^{disc} changes only the scalar attached to λ. Both operators are bounded, so the identity extends from the algebraic core. □

Remark 13.2.4. The operator-norm formula $\|C_i^{\mathrm{disc}}\| = \max\{\rho_i^{-1}, \tau\}$ is identical to the continuous formula $\|C_i\| = \max\{\rho_i^{-1}, \tau\}$ from Theorem 9.3.3 of Chapter 9. This is another instance of the general principle that the norm estimates depend on the coefficient-space geometry, not on the function-space realization.

13.3 The discrete commutator formula

We now compute the commutator $[C_i^{\mathrm{disc}}, C_j^{\mathrm{disc}}]$ on each basis vector. The result is, once again, that the commutator vanishes everywhere except at a single locus: simultaneous vacuum in two free one-sided coordinates inside a trace block.

Theorem 13.3.1 (Explicit discrete commutator)**.** *Let $i, j \in \{1, \ldots, r\}$ with $i \neq j$.*

On the canonical block:

$$[C_i^{\mathrm{disc}}, C_j^{\mathrm{disc}}]\, e_{\mathbf{k},\lambda} = 0 \qquad \textit{for all } (\mathbf{k},\lambda) \in \mathbb{N}_0^r \times \Lambda.$$

On a trace basis vector $t_{w,\mathbf{k},\lambda}$:

$$[C_i^{\mathrm{disc}}, C_j^{\mathrm{disc}}]\, t_{w,\mathbf{k},\lambda} = \begin{cases} t_{wji,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda} - t_{wij,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda}, & i,j \in \mathrm{F}(w) \textit{ and } k_i = k_j = 0, \\ 0, & \textit{otherwise}, \end{cases}$$

where $\mathbf{k}^{\widehat{\{i,j\}}} \in \mathbb{N}_0^{\mathrm{F}(w)\setminus\{i,j\}}$ denotes the residual multi-index with both the i-th and j-th components removed.

Proof. Canonical block. On $e_{\mathbf{k},\lambda}$, both operators are the canonical discrete backward shifts, which commute on the discrete hybrid module by the results of Chapter 11.

Trace block: setup. Fix a trace basis vector $t_{w,\mathbf{k},\lambda}$. If either i or j is not in $\mathrm{F}(w)$, then the corresponding operator annihilates, making both compositions zero.

Assume $i,j \in \mathrm{F}(w)$. We consider four cases.

Case $k_i \geq 1$, $k_j \geq 1$. Both operators lower residual grades in different coordinates:

$$C_i^{\mathrm{disc}} C_j^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} = C_j^{\mathrm{disc}} C_i^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} = t_{w,\,\mathbf{k}-\mathbf{e}_i-\mathbf{e}_j,\,\lambda}.$$

The commutator vanishes.

Case $k_i = 0$, $k_j \geq 1$. Applying C_j^{disc} first lowers the j-th grade:

$$C_j^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} = t_{w,\,\mathbf{k}-\mathbf{e}_j,\,\lambda}.$$

Since the i-th coordinate is still free with $k_i = 0$, the next application of C_i^{disc} appends the letter i:

$$C_i^{\mathrm{disc}} C_j^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} = t_{wi,\,(\mathbf{k}-\mathbf{e}_j)^{\hat{i}},\,\lambda}.$$

Applying in the other order: C_i^{disc} first appends i (since $k_i = 0$), producing $t_{wi,\,\mathbf{k}^{\hat{i}},\,\lambda}$. Then C_j^{disc} lowers the j-th coordinate (which is still free in $\mathrm{F}(wi)$ with grade $k_j \geq 1$), giving the same result. The commutator vanishes. The case $k_i \geq 1$, $k_j = 0$ is symmetric.

Case $k_i = 0$, $k_j = 0$. Applying C_j^{disc} first appends the letter j:

$$C_j^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} = t_{wj,\,\mathbf{k}^{\hat{j}},\,\lambda}.$$

Since $i \in \mathrm{F}(wj) = \mathrm{F}(w) \setminus \{j\}$ and the i-th residual grade is still zero, the next application of C_i^{disc} appends i:

$$C_i^{\mathrm{disc}} C_j^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} = t_{wji,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda}.$$

Reversing the order:

$$C_j^{\mathrm{disc}} C_i^{\mathrm{disc}}\, t_{w,\mathbf{k},\lambda} = t_{wij,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda}.$$

The words wji and wij are distinct (the last two letters are transposed), so the commutator is $t_{wji,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda} - t_{wij,\,\mathbf{k}^{\widehat{\{i,j\}}},\,\lambda} \neq 0$. □

Remark 13.3.2. The commutator formula of Theorem 13.3.1 is identical, symbol by symbol, to the continuous commutator formula of Theorem 9.4.1 in Chapter 9. The spectral label λ plays no role in either formula. The commutator depends only on the trace word, the free set, and the residual grades at the free coordinates—all of which are coefficient-space data. This is the strongest evidence that the boundary mechanism is truly function-space-independent.

Example 13.3.3 (The basic discrete commutator for $r = 2$)**.** Take $r = 2$ and the empty trace word $w = \varnothing$ with $\mathbf{k} = (0, 0)$. Then

$$\begin{aligned} C_1^{\mathrm{disc}} C_2^{\mathrm{disc}} \, t_{\varnothing,(0,0),\lambda} &= t_{(2,1),\varnothing,\lambda}, \\ C_2^{\mathrm{disc}} C_1^{\mathrm{disc}} \, t_{\varnothing,(0,0),\lambda} &= t_{(1,2),\varnothing,\lambda}. \end{aligned}$$

The two results are distinct: applying C_2^{disc} first records the word (2), then C_1^{disc} extends it to $(2, 1)$; the reverse order produces $(1, 2)$. The commutator is

$$[C_1^{\mathrm{disc}}, C_2^{\mathrm{disc}}] \, t_{\varnothing,(0,0),\lambda} \;=\; t_{(2,1),\varnothing,\lambda} \;-\; t_{(1,2),\varnothing,\lambda} \;\neq\; 0.$$

On the lattice $\mathbb{N}_0^2$, the point $(0, 0)$ is the corner where both faces $\{n_1 = 0\}$ and $\{n_2 = 0\}$ meet. At this corner, the two boundary extractions compete for priority, and their order produces the nonzero commutator.

13.4 The discrete maximal commuting sector

We now classify the largest graded invariant sector on which the extended discrete Caputo tuple remains commuting. As in the continuous case, the graded hypothesis is essential: without it, accidental finite-dimensional commuting subspaces destroy any meaningful maximality statement.

Definition 13.4.1. A closed subspace $U \subset \mathfrak{X}^{p,\mathrm{disc}}_{\rho,\tau,\eta}$ is called *closed graded* if there exists a subset S of the full basis $\{e_{\mathbf{k},\lambda}\} \cup \{t_{w,\mathbf{k},\lambda}\}$ such that $U = \overline{\mathrm{span}(S)}$.

Definition 13.4.2. The *thin discrete trace sector* is the algebraic span

$$\mathfrak{T}^{\mathrm{alg,disc}}_{\leq 1} \;:=\; \bigoplus_{\substack{w \in \mathcal{W}_r \\ q(w) \leq 1}} \mathcal{T}^{\mathrm{alg,disc}}_w,$$

consisting of all trace blocks whose free rank is at most 1. The *maximal commuting algebraic sector* is

$$\mathfrak{K}^{\mathrm{alg,disc}} \;:=\; \mathcal{H}^{\mathrm{alg,disc}}_{\boldsymbol{\alpha},\Lambda} \;\oplus\; \mathfrak{T}^{\mathrm{alg,disc}}_{\leq 1},$$

and its Banach closure is

$$\mathfrak{K}^{p,\mathrm{disc}}_{\rho,\tau,\eta} \;:=\; X^{p,\mathrm{disc}}_{\rho,\eta} \;\oplus\; \bigoplus_{\substack{w \in \mathcal{W}_r \\ q(w) \leq 1}} T^{p,\mathrm{disc}}_w.$$

The thin trace sector contains precisely those blocks in which at most one free one-sided lattice coordinate remains. When $q(w) = 0$, no free coordinate exists and every C_i^{disc} annihilates.

When $q(w) = 1$, only one free coordinate i_0 remains; every other C_j^{disc} with $j \neq i_0$ annihilates, so no pair of operators can simultaneously act nontrivially, and no commutator can form.

Proposition 13.4.3 (Invariance). *The sector $\mathfrak{K}^{\mathrm{alg,disc}}$ is invariant under every C_i^{disc} and every bounded multiplier M_σ^{disc}.*

Proof. The canonical block is invariant by construction. For a trace basis vector $t_{w,\mathbf{k},\lambda}$ with $q(w) \leq 1$: if $q(w) = 0$, every operator annihilates; if $q(w) = 1$ with unique free coordinate i_0, then $C_{i_0}^{\mathrm{disc}}$ either lowers the residual grade (staying in the same block) or appends i_0 to produce a word of free rank 0—still inside the thin sector. Every other C_j^{disc} with $j \neq i_0$ annihilates. Multipliers preserve every block by diagonal action. □

Proposition 13.4.4 (Commutativity). *The tuple $(C_1^{\mathrm{disc}}, \ldots, C_r^{\mathrm{disc}})$ is commuting on $\mathfrak{K}^{\mathrm{alg,disc}}$.*

Proof. On the canonical block, commutativity holds by the results of Chapters 11–12. On a trace basis vector $t_{w,\mathbf{k},\lambda}$ with $q(w) \leq 1$: for any pair of distinct indices $i \neq j$, at most one can belong to $\mathrm{F}(w)$, so at least one of C_i^{disc} or C_j^{disc} annihilates. Hence both compositions give zero, and the commutator vanishes. By linearity, the full tuple commutes on $\mathfrak{K}^{\mathrm{alg,disc}}$. □

The maximality theorem says this is the best one can do.

Theorem 13.4.5 (Maximal discrete commuting graded sector). *Let $U \subset \mathfrak{X}_{\rho,\tau,\eta}^{p,\mathrm{disc}}$ be a closed graded subspace such that:*

(i) $X_{\rho,\eta}^{p,\mathrm{disc}} \subset U$,

(ii) $C_i^{\mathrm{disc}} U \subset U$ *for all* $i \in \{1, \ldots, r\}$,

(iii) the tuple $(C_1^{\mathrm{disc}}, \ldots, C_r^{\mathrm{disc}})$ *is commuting on* U.

Then $U \subset \mathfrak{K}_{\rho,\tau,\eta}^{p,\mathrm{disc}}$.

Consequently, $\mathfrak{K}_{\rho,\tau,\eta}^{p,\mathrm{disc}}$ is the unique maximal closed graded invariant sector containing the canonical discrete completion and carrying a commuting discrete Caputo tuple.

Proof. Write $U = \overline{U^{\mathrm{alg}}}$ for a graded algebraic subspace U^{alg}. We must show that U^{alg} contains no trace basis vector $t_{w,\mathbf{k},\lambda}$ with $q(w) \geq 2$.

Suppose for contradiction that such a vector lies in U^{alg}. Since U^{alg} is graded and invariant under every C_i^{disc}, repeated application of the operators along the free coordinates lowers all positive residual grades to zero. (Each application of C_i^{disc} to a vector with $k_i \geq 1$ lowers that grade by one without appending a new trace letter.) Hence the vector $t_{w,\mathbf{0},\lambda}$—with all free coordinates at grade zero—also belongs to U^{alg}.

Since $q(w) \geq 2$, there exist distinct $i, j \in \mathrm{F}(w)$. By Theorem 13.3.1:

$$C_i^{\mathrm{disc}} C_j^{\mathrm{disc}} t_{w,\mathbf{0},\lambda} = t_{wji,\varnothing,\lambda}, \qquad C_j^{\mathrm{disc}} C_i^{\mathrm{disc}} t_{w,\mathbf{0},\lambda} = t_{wij,\varnothing,\lambda}.$$

These are distinct basis vectors (the words wji and wij differ), so

$$C_i^{\mathrm{disc}} C_j^{\mathrm{disc}} t_{w,\mathbf{0},\lambda} \neq C_j^{\mathrm{disc}} C_i^{\mathrm{disc}} t_{w,\mathbf{0},\lambda},$$

contradicting the assumption that the tuple commutes on U^{alg}.

Therefore U^{alg} contains no trace basis vector with $q(w) \geq 2$, and $U^{\mathrm{alg}} \subset \mathfrak{K}^{\mathrm{alg,disc}}$. Taking closures gives $U \subset \mathfrak{K}^{p,\mathrm{disc}}_{\rho,\tau,\eta}$.

Maximality follows: $\mathfrak{K}^{p,\mathrm{disc}}_{\rho,\tau,\eta}$ itself satisfies (i)–(iii) by Propositions 13.4.3 and 13.4.4, and every other sector satisfying (i)–(iii) is contained in it. □

Remark 13.4.6. The graded hypothesis is not a technicality but a structural requirement. Without it, one could take a two-dimensional subspace spanned by $t_{wji,\mathbf{0},\lambda} + t_{wij,\mathbf{0},\lambda}$ (the *symmetric* combination) and produce a subspace on which the commutator happens to vanish. Such accidental cancellations are invisible to the basis-level structure and do not extend to a meaningful invariant sector. The graded condition ensures that commutativity is tested basis line by basis line, which is the natural framework for the canonical models developed throughout this book.

13.5 Shared boundary logic in the continuous and discrete theories

We have now completed the boundary-augmented theory in both the continuous setting (Chapter 9) and the discrete setting (this chapter). The structural parallel is exact, and this section makes it explicit.

The following table displays the correspondence.

Feature	**Continuous (Ch. 9)**	**Discrete (Ch. 13)**
One-sided domain	$(0,\infty)^r$	$\mathbb{N}^r$
Whole-space domain	$\mathbb{R}^s$	$\mathbb{Z}^s$
Canonical basis	$\frac{x^{k\alpha}}{\Gamma(k\alpha+1)} \cdot e^{\langle\lambda,x\rangle}$	$\frac{n^{\overline{k\alpha}}}{\Gamma(k\alpha+1)} \cdot \lambda^m$
Spectral condition	$\operatorname{Re}\lambda_j > 0$	$\lvert\lambda_j\rvert > 1$
Trace word set	$\mathcal{W}_r$	$\mathcal{W}_r$
Extended Caputo action	Definition 9.3.1	Definition 13.2.1
Operator norm	$\max\{\rho_i^{-1},\tau\}$	$\max\{\rho_i^{-1},\tau\}$
Commutator formula	Theorem 9.4.1	Theorem 13.3.1
Maximal sector condition	$q(w) \leq 1$	$q(w) \leq 1$

Every entry in the "Discrete" column is either identical to the continuous entry or differs only in the function-space realization (rising factorials replacing monomials, lattice characters replacing exponentials, $|\lambda_j| > 1$ replacing $\operatorname{Re}\lambda_j > 0$). The coefficient-space structure—the word combinatorics, the norm formula, the commutator, and the maximality criterion—is the same.

Remark 13.5.1 (Why the parallelism is exact). The boundary mechanism depends on three facts: (a) the one-sided grades take values in $\mathbb{N}_0$, (b) the backward shift annihilates at grade zero, and (c) the extended operator records the order of boundary extraction when annihilation would occur. All three facts are common to both settings. The specific formula for the basis vectors and the spectral condition on Λ enter only through the weight structure of the completion, not through the combinatorial core that drives noncommutativity. This is why the commutator formula and the maximal-sector classification are identical.

This observation will be elevated to a theorem in Chapter 14, where the continuous and discrete completions are recognized as two instances of a single abstract coefficient-space model. In that framework, the boundary-augmented theory becomes literally the same construction, parametrized by different choices of spectral data.

13.6 Low-dimensional examples

The maximality theorem is most transparent in the smallest cases. We examine $r = 1$ and $r = 2$ in detail.

Example 13.6.1 (The case $r = 1$: no obstruction)**.** When there is a single one-sided coordinate ($r = 1$), the trace words are $\varnothing$ (free rank 1) and (1) (free rank 0). Both have free rank at most 1, so the thin sector exhausts the full ambient space:

$$\mathfrak{K}^{p,\mathrm{disc}}_{\rho,\tau,\eta} = \mathfrak{X}^{p,\mathrm{disc}}_{\rho,\tau,\eta}.$$

This is consistent with the fact that a single operator C_1^{disc} trivially commutes with itself. In one one-sided lattice dimension, there is only one face $\{n_1 = 0\}$, only one boundary event can occur, and no ordering ambiguity ever arises.

Example 13.6.2 (The case $r = 2$: the first obstruction)**.** With $r = 2$, the words and their free ranks are:

w	$\varnothing$	(1)	(2)	$(1,2)$	$(2,1)$
$q(w)$	2	1	1	0	0

The maximal commuting sector is

$$\mathfrak{K}^{p,\mathrm{disc}}_{\rho,\tau,\eta} = X^{p,\mathrm{disc}}_{\rho,\eta} \oplus T^{p,\mathrm{disc}}_{(1)} \oplus T^{p,\mathrm{disc}}_{(2)} \oplus T^{p,\mathrm{disc}}_{(1,2)} \oplus T^{p,\mathrm{disc}}_{(2,1)}.$$

The unique excluded block is $T^{p,\mathrm{disc}}_{\varnothing}$, the empty-word trace block with free rank 2. This is the block in which the basic commutator of Example 13.3.3 lives.

On the lattice $\mathbb{N}_0^2$, the geometry is especially vivid. The canonical completion lives on the interior of the first quadrant (all grades positive) together with the axes and the origin. The empty-word trace block represents states at the corner $(0,0)$ that have not yet recorded any boundary extraction. At the corner, both faces $\{n_1 = 0\}$ and $\{n_2 = 0\}$ meet, and the two boundary extractions compete for priority. The maximal commuting sector resolves this by excluding the unresolved corner states and retaining only those states that have already committed to a boundary ordering.

Example 13.6.3 (The case $r = 3$: counting excluded blocks)**.** With $r = 3$, the words have the following distribution by free rank:

$q(w)$	3	2	1	0
number of words	1	3	6	6

The single word with $q(w) = 3$ is $\varnothing$; the three words with $q(w) = 2$ are (1), (2), (3). These four trace blocks are excluded from the maximal commuting sector. The remaining twelve blocks (six with $q(w) = 1$ and six with $q(w) = 0$) are included. Thus the ambient space has $1 + 16 = 17$ blocks (one canonical plus sixteen trace), and the maximal commuting sector retains $1 + 12 = 13$ of them.

The excluded blocks are precisely those corresponding to boundary strata of codimension 0 or 1 (in the word sense), where two or more lattice faces remain simultaneously active. As r grows, the proportion of excluded blocks increases: the "noncommutative part" of the ambient theory becomes larger, while the canonical commuting core remains the fixed anchor.

Corollary 13.6.4. *The classification of maximal graded commuting sectors is independent of the spectral data. Every bounded multiplier* M_σ^{disc} *preserves both* $\mathfrak{R}^{p,\mathrm{disc}}_{\rho,\tau,\eta}$ *and each trace block* $T_w^{p,\mathrm{disc}}$*, and the obstruction to commutativity is determined entirely by the one-sided free-rank condition* $q(w) \geq 2$.

Proof. Every multiplier acts diagonally on the spectral label λ and leaves the trace word and the one-sided grades unchanged. The commutator formula (Theorem 13.3.1) never involves λ. Hence the maximality criterion depends only on the one-sided structure. □

Remark 13.6.5. The lattice setting offers a useful geometric picture for the maximality theorem. A word w of length m corresponds to a boundary stratum of codimension m in $\mathbb{N}_0^r$: the intersection of the m faces $\{n_{i_\nu} = 0\}$ for $\nu = 1, \ldots, m$. The free rank $q(w) = r - m$ counts the number of remaining "bulk" directions. The maximal commuting sector retains a trace block if and only if the corresponding boundary stratum has at most one remaining bulk direction.

In other words: the obstruction to commutativity requires a lattice corner where at least two faces meet and both are still "live" (i.e., neither has been recorded yet in the trace word). At such a corner, the two boundary extractions can occur in two different orders, producing the nonzero commutator. Once all but one face has been recorded, no further ordering ambiguity exists, and commutativity is automatic.

Looking ahead

This chapter completes the discrete counterpart of Chapters 11–9: the discrete hybrid algebra (Chapter 11), its Banach completion and Z-transform model (Chapter 12), and its boundary-augmented maximal commuting structure (this chapter). The parallel with the continuous sequence (Chapters 7–9) is exact in every respect: the same word combinatorics, the same commutator formula, the same maximality criterion, and the same operator-norm identities.

The next two chapters draw the consequences of this parallelism. Chapter 14 unifies the continuous and discrete completion theories in a single abstract coefficient-space framework, where the common structure is made manifest and the transform models, optimal weights, and semigroup generation are treated once for both settings. Chapter 15 then returns to the boundary-augmented theory and asks a different question: instead of *restricting* to a maximal commuting sector, can one *quotient* the ambient space to recover commutativity? The answer leads to the theory of commutator ideals, ordered boundary traces, and the universal commuting quotient—the culminating algebraic structure of the entire program.

Remark 13.6.6 (Restriction and quotient). The maximal commuting sector of this chapter is obtained by restriction: one discards the trace blocks whose free rank is too large. Chapter 15 will develop the complementary strategy of quotient reduction, in which one keeps all trace blocks but identifies ordered histories that differ only in the ordering of the same letters. The two reductions—restriction and quotient—are genuinely different. Restriction preserves the ambient-space structure on a smaller domain; the quotient preserves the full domain but coarsens the space. Both are natural, and comparing them will reveal the full algebraic picture of boundary-generated noncommutativity.

Chapter 14

Unified Coefficient-Space Theory, Optimal Weights, and Semigroups

The preceding chapters developed two parallel completion theories. In Chapter 8, the algebraic hybrid module on $(0,\infty)^r \times \mathbb{R}^s$ was completed to a weighted Banach space carrying all of the canonical shift and spectral relations, and a fiberwise holomorphic model on a polydisk was constructed. In Chapter 12, the discrete algebraic hybrid module on $\mathbb{N}^r \times \mathbb{Z}^s$ was completed in the same manner, producing a fiberwise Z-transform model on a polydisk. The reader will have noticed a striking pattern: every definition, every theorem, and every proof in the two chapters followed the same logic, differing only in the function-space realization (Gamma-normalized monomials versus normalized rising factorials, exponential characters versus lattice characters, $\operatorname{Re}\lambda_j > 0$ versus $|\lambda_j| > 1$).

This chapter makes the common structure explicit. Instead of working with continuous functions on mixed regions or discrete sequences on mixed lattices, we work with the underlying *coefficient space*: a weighted ℓ^p space over the grade lattice $\mathbb{N}_0^r$ and the spectral index set Λ. From this vantage point, three questions that were left open by the earlier completion chapters receive sharp answers.

Which weights are the right ones? In Chapters 8 and 12, geometric weights provided especially transparent examples, but no intrinsic characterization was given. We prove here that geometric weights are exactly the *balanced* weights—the unique weights for which each forward and backward shift pair achieves the minimal possible product norm $\|J_i\|\,\|C_i\| = 1$.

To what extent are the two transform models different? We construct a single coefficient transform on a polydisk that simultaneously underlies both the continuous holomorphic realization and the discrete Z-transform realization. At the coefficient level, the two theories are literally the same.

Which operators generate semigroups? We prove an exact semigroup criterion for the diagonal spectral block, establish semigroup generation for mixed shift–spectral generators by bounded perturbation, and optimize the exponential growth bound over the geometric weight parameters.

The material follows the research paper AD11. Throughout, we fix $r \geq 1$ and a nonempty spectral index set Λ. The exponent $1 \leq p < \infty$ is fixed unless otherwise stated.

14.1 Why a unified abstract theory is now necessary

Consider the key ingredients of both completion theories. In each case one starts with a canonical basis indexed by pairs $(\mathbf{k}, \lambda) \in \mathbb{N}_0^r \times \Lambda$. In each case the one-sided operators act as forward and backward shifts on the grade index $\mathbf{k}$, the spectral operators act diagonally on λ, and the two blocks commute. In each case one equips the finite-support span with a weighted norm and completes. The resulting Banach spaces are weighted ℓ^p spaces over the same index set, carrying the same operator relations.

The only difference is the concrete interpretation of the basis vectors. In the continuous theory, $e_{\mathbf{k},\lambda}$ is a product of Gamma-normalized monomials and exponentials; in the discrete theory, it is a product of normalized rising factorials and lattice characters. This interpretation matters for pointwise function-space properties (integral representations, asymptotic estimates, kernel formulae), but it plays no role at all in the coefficient-space theory of norms, operator bounds, transform models, or semigroup generation.

The purpose of this chapter is to develop the coefficient-space theory once, abstractly, and then transport the results to both realizations by a single isometric identification. This eliminates the redundancy of proving the same theorems twice and reveals the structure that governs both branches.

Remark 14.1.1. The coefficient-space viewpoint is standard in harmonic analysis and operator theory. The Fourier transform identifies function-space properties with sequence-space properties, and many classical theorems (Plancherel, Hausdorff–Young, multiplier theorems) are most naturally proved on the coefficient side. What is distinctive here is that the coefficient lattice $\mathbb{N}_0^r$ carries a nontrivial boundary at the coordinate hyperplanes $\{k_i = 0\}$, and this boundary is the source of the vacuum projections, defect phenomena, and—ultimately—the noncommutativity studied in Chapters 9 and 13.

14.2 The abstract coefficient-space model

We now set up the abstract framework common to both realizations.

Definition 14.2.1. The *abstract hybrid algebraic core* is the vector space

$$\mathfrak{H}_\Lambda^{\mathrm{alg}} := \bigoplus_{(\mathbf{k},\lambda)\in\mathbb{N}_0^r\times\Lambda} \mathbb{C}\,\mathbf{e}_{\mathbf{k},\lambda},$$

where $\{\mathbf{e}_{\mathbf{k},\lambda}\}$ is a formal basis indexed by the pairs $(\mathbf{k},\lambda)\in\mathbb{N}_0^r\times\Lambda$. Every element of $\mathfrak{H}_\Lambda^{\mathrm{alg}}$ is a finite linear combination of these basis vectors.

The three families of operators that have appeared throughout the book—forward shifts, backward shifts, and spectral multipliers—are now defined purely in terms of index manipulation, without any reference to monomials, rising factorials, or characters.

Definition 14.2.2. For each $i\in\{1,\dots,r\}$, define the *abstract forward shift* and *abstract backward shift* on $\mathfrak{H}_\Lambda^{\mathrm{alg}}$ by

$$J_i\,\mathbf{e}_{\mathbf{k},\lambda} := \mathbf{e}_{\mathbf{k}+\mathbf{e}_i,\lambda},$$

and

$$C_i\,\mathbf{e}_{\mathbf{k},\lambda} := \begin{cases} 0, & k_i = 0,\\ \mathbf{e}_{\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i\ge 1.\end{cases}$$

For any function $\sigma : \Lambda \to \mathbb{C}$, define the *diagonal multiplier*

$$M_\sigma \, \mathbf{e}_{\mathbf{k},\lambda} \ := \ \sigma(\lambda) \, \mathbf{e}_{\mathbf{k},\lambda}.$$

Definition 14.2.3. For each $i \in \{1, \ldots, r\}$, the *coordinate-vacuum projection* is

$$\Pi_i \, \mathbf{e}_{\mathbf{k},\lambda} \ := \ \begin{cases} \mathbf{e}_{\mathbf{k},\lambda}, & k_i = 0, \\ 0, & k_i \geq 1. \end{cases}$$

More generally, for $\mathbf{m} = (m_1, \ldots, m_r) \in \mathbb{N}_0^r$, the *boundary-layer projection* $\Pi_{<\mathbf{m}}$ projects onto the span of all $\mathbf{e}_{\mathbf{k},\lambda}$ with $\mathbf{k} \not\geq \mathbf{m}$ (i.e., with $k_i < m_i$ for at least one coordinate i).

The following proposition collects the algebraic relations that have appeared in concrete form throughout the book. In the abstract setting they become transparent.

Proposition 14.2.4 (Algebraic relations)**.** *On $\mathfrak{H}_\Lambda^{\mathrm{alg}}$, the following hold.*

(i) Commutativity. *For all $i, j \in \{1, \ldots, r\}$:*

$$J_i J_j = J_j J_i, \qquad C_i C_j = C_j C_i, \qquad C_i J_j = J_j C_i \quad (i \neq j).$$

(ii) Shift–vacuum identities. *For each i:*

$$C_i J_i = I, \qquad J_i C_i = I - \Pi_i, \qquad [C_i, J_i] = \Pi_i.$$

(iii) Higher-order identities. *For any $\mathbf{m} \in \mathbb{N}_0^r$, if $J^{\mathbf{m}} := J_1^{m_1} \cdots J_r^{m_r}$ and $C^{\mathbf{m}} := C_1^{m_1} \cdots C_r^{m_r}$, then*

$$C^{\mathbf{m}} J^{\mathbf{m}} = I, \qquad J^{\mathbf{m}} C^{\mathbf{m}} = I - \Pi_{<\mathbf{m}}.$$

(iv) Spectral commutation. *Every multiplier M_σ commutes with every J_i, every C_i, and every projection Π_i, $\Pi_{<\mathbf{m}}$.*

Proof. Each identity is verified directly on a basis vector $\mathbf{e}_{\mathbf{k},\lambda}$. For the central example,

$$C_i J_i \, \mathbf{e}_{\mathbf{k},\lambda} = C_i \, \mathbf{e}_{\mathbf{k}+\mathbf{e}_i,\lambda} = \mathbf{e}_{\mathbf{k},\lambda},$$

since $k_i + 1 \geq 1$. On the other hand, $J_i C_i \, \mathbf{e}_{\mathbf{k},\lambda}$ equals $\mathbf{e}_{\mathbf{k},\lambda}$ when $k_i \geq 1$ and 0 when $k_i = 0$; this is precisely $(I - \Pi_i) \, \mathbf{e}_{\mathbf{k},\lambda}$. The commutator identity $[C_i, J_i] = C_i J_i - J_i C_i = I - (I - \Pi_i) = \Pi_i$ follows immediately.

The higher-order identities follow by induction on the components of $\mathbf{m}$, and the multiplier commutation holds because M_σ changes only the spectral scalar while J_i, C_i, and Π_i change only the grade index. □

Remark 14.2.5. The relations in Proposition 14.2.4 are precisely the shift–vacuum identities proved for the continuous model in Chapters 4–7 and for the discrete model in Chapter 11. In those chapters, the proofs required explicit computation with Gamma-normalized monomials or rising factorials. In the abstract setting, the same identities become tautological—they hold because they are built into the definition of the shift operators, independently of any function-space interpretation. This is the payoff of the coefficient-space viewpoint.

Weighted completions

Definition 14.2.6. A *weight* on $\mathbb{N}_0^r \times \Lambda$ is a function $\omega : \mathbb{N}_0^r \times \Lambda \to (0, \infty)$.

Definition 14.2.7. For $1 \le p < \infty$ and a finite-support element $u = \sum a_{\mathbf{k},\lambda}\, \mathbf{e}_{\mathbf{k},\lambda} \in \mathfrak{H}_\Lambda^{\mathrm{alg}}$, define

$$\|u\|_{p,\omega} \;:=\; \left(\sum_{\mathbf{k},\lambda} |a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k},\lambda)^p \right)^{1/p}.$$

The Banach completion of $\mathfrak{H}_\Lambda^{\mathrm{alg}}$ under this norm is denoted X_ω^p.

Proposition 14.2.8. *The space X_ω^p is canonically isometric to the weighted sequence space*

$$\ell_\omega^p(\mathbb{N}_0^r \times \Lambda) \;:=\; \left\{ a = (a_{\mathbf{k},\lambda}) : \sum_{\mathbf{k},\lambda} |a_{\mathbf{k},\lambda}|^p \, \omega(\mathbf{k},\lambda)^p < \infty \right\},$$

with the algebraic core $\mathfrak{H}_\Lambda^{\mathrm{alg}}$ corresponding to the finite-support sequences.

Proof. The coefficient map $\sum a_{\mathbf{k},\lambda}\, \mathbf{e}_{\mathbf{k},\lambda} \mapsto (a_{\mathbf{k},\lambda})$ is linear and isometric on the core. Its image—the finite-support families—is dense in ℓ_ω^p because $1 \le p < \infty$. The result follows by completion. □

This simple observation is the key to the chapter. Every Banach-space result we prove will be a statement about weighted ℓ^p spaces, and therefore it will apply simultaneously to the continuous and discrete realizations once the appropriate isometric identification is made.

14.3 Admissibility and exact norm formulae

We now characterize the weights for which the shift operators extend boundedly to the completion and give exact formulae for their operator norms.

Definition 14.3.1. A weight ω is called *shift-admissible* if, for every $i \in \{1, \ldots, r\}$, both of the following quantities are finite:

$$L_i \;:=\; \sup_{\mathbf{k},\lambda} \frac{\omega(\mathbf{k}+\mathbf{e}_i,\lambda)}{\omega(\mathbf{k},\lambda)}, \qquad R_i \;:=\; \sup_{\substack{\mathbf{k},\lambda \\ k_i \ge 1}} \frac{\omega(\mathbf{k}-\mathbf{e}_i,\lambda)}{\omega(\mathbf{k},\lambda)}.$$

The ratio L_i measures the maximal cost of moving one grade step *forward* in the i-th coordinate: it compares the weight of a basis vector with the weight of its forward-shifted image. The ratio R_i measures the maximal cost of moving one step *backward*. If the weight grows rapidly in the i-th direction, then L_i is large and R_i is small, and vice versa.

Theorem 14.3.2 (Exact norm formulae)**.** *For each $i \in \{1, \dots, r\}$, the operators J_i and C_i extend uniquely to bounded operators on X^p_ω if and only if ω is shift-admissible. In that case,*

$$\|J_i\| = L_i, \qquad \|C_i\| = R_i.$$

Moreover, the projections Π_i and $\Pi_{<\mathbf{m}}$ extend to bounded contractive projections on X^p_ω.

Proof. Upper bounds. For a finite-support vector $u = \sum a_{\mathbf{k},\lambda}\, \mathbf{e}_{\mathbf{k},\lambda}$,

$$\|J_i u\|^p_{p,\omega} = \sum_{\mathbf{k},\lambda} |a_{\mathbf{k},\lambda}|^p\, \omega(\mathbf{k} + \mathbf{e}_i, \lambda)^p \le L_i^p\, \|u\|^p_{p,\omega},$$

since on each basis vector the weight ratio is at most L_i. Similarly,

$$\|C_i u\|^p_{p,\omega} = \sum_{\substack{\mathbf{k},\lambda \\ k_i \ge 1}} |a_{\mathbf{k},\lambda}|^p\, \omega(\mathbf{k} - \mathbf{e}_i, \lambda)^p \le R_i^p\, \|u\|^p_{p,\omega}.$$

Hence J_i and C_i extend boundedly to the completion with $\|J_i\| \le L_i$ and $\|C_i\| \le R_i$.

Lower bounds. Testing on unit basis vectors:

$$\frac{\|J_i\, \mathbf{e}_{\mathbf{k},\lambda}\|}{\|\mathbf{e}_{\mathbf{k},\lambda}\|} = \frac{\omega(\mathbf{k} + \mathbf{e}_i, \lambda)}{\omega(\mathbf{k}, \lambda)},$$

and when $k_i \ge 1$,

$$\frac{\|C_i\, \mathbf{e}_{\mathbf{k},\lambda}\|}{\|\mathbf{e}_{\mathbf{k},\lambda}\|} = \frac{\omega(\mathbf{k} - \mathbf{e}_i, \lambda)}{\omega(\mathbf{k}, \lambda)}.$$

Taking suprema gives $\|J_i\| \ge L_i$ and $\|C_i\| \ge R_i$.

Necessity. If J_i is bounded on the completion, then the weight ratio $\omega(\mathbf{k} + \mathbf{e}_i, \lambda)/\omega(\mathbf{k}, \lambda)$ is bounded by $\|J_i\|$ for every $(\mathbf{k}, \lambda)$, so $L_i < \infty$. The argument for C_i and R_i is analogous.

Projections. The estimate $\|\Pi_i u\|^p_{p,\omega} \le \|u\|^p_{p,\omega}$ holds because Π_i selects a subset of the terms in the norm sum. The same argument applies to $\Pi_{<\mathbf{m}}$. □

Corollary 14.3.3. *If ω is shift-admissible, then all algebraic identities of Proposition 14.2.4 extend from the core to the full Banach space X^p_ω.*

Proof. All operators appearing in the identities are bounded by Theorem 14.3.2, and the core is dense. Hence the identities extend by continuity. □

Example 14.3.4. If $\omega(\mathbf{k}, \lambda) = \eta(\lambda)\, \rho_1^{k_1} \cdots \rho_r^{k_r}$ for some $\boldsymbol{\rho} \in (0, \infty)^r$ and $\eta : \Lambda \to (0, \infty)$, then $L_i = \rho_i$ and $R_i = \rho_i^{-1}$ for every i, and the weight is shift-admissible. We will see in the next section that these geometric weights are singled out by a sharp optimality property.

Example 14.3.5. Consider the one-variable weight $\omega(k) = (k + 1)^\beta$ for some $\beta > 0$. Then $L = \sup_k((k+2)/(k+1))^\beta = 2^\beta$ and $R = \sup_{k \ge 1}(k/(k+1))^\beta = 1$, so the weight is shift-admissible. But the product $LR = 2^\beta > 1$: this weight is admissible but not balanced.

14.4 The optimal-weight theorem

Among all shift-admissible weights, the geometric weights occupy a distinguished position: they are the unique weights for which each forward–backward shift pair is perfectly balanced.

Theorem 14.4.1 (Sharp product inequality and balanced weights)**.** *Let ω be shift-admissible. Then for each $i \in \{1, \dots, r\}$,*

$$\|J_i\| \, \|C_i\| \;\geq\; 1.$$

Moreover, the following are equivalent:

(i) $\|J_i\| \, \|C_i\| = 1$ *for every* $i \in \{1, \dots, r\}$.

(ii) There exist constants $\rho_1, \dots, \rho_r \in (0, \infty)$ *and a function* $\eta : \Lambda \to (0, \infty)$ *such that*

$$\omega(\mathbf{k}, \lambda) = \eta(\lambda) \, \rho_1^{k_1} \cdots \rho_r^{k_r} \qquad \textit{for all } (\mathbf{k}, \lambda) \in \mathbb{N}_0^r \times \Lambda.$$

Proof. Fix a coordinate i and define the weight ratio

$$r_i(\mathbf{k}, \lambda) \;:=\; \frac{\omega(\mathbf{k} + \mathbf{e}_i, \lambda)}{\omega(\mathbf{k}, \lambda)}.$$

By Theorem 14.3.2, $\|J_i\| = \sup r_i$. For the backward shift, reindex by setting $\mathbf{m} = \mathbf{k} - \mathbf{e}_i$:

$$\|C_i\| = \sup_{\substack{\mathbf{k}, \lambda \\ k_i \geq 1}} \frac{\omega(\mathbf{k} - \mathbf{e}_i, \lambda)}{\omega(\mathbf{k}, \lambda)} = \sup_{\mathbf{m}, \lambda} \frac{\omega(\mathbf{m}, \lambda)}{\omega(\mathbf{m} + \mathbf{e}_i, \lambda)} = \sup_{\mathbf{m}, \lambda} \frac{1}{r_i(\mathbf{m}, \lambda)} = \frac{1}{\inf r_i}.$$

Therefore

$$\|J_i\| \, \|C_i\| = \frac{\sup r_i}{\inf r_i} \;\geq\; 1,$$

with equality if and only if r_i is constant.

Now assume $\|J_i\| \, \|C_i\| = 1$ for every i. Then each ratio r_i is constant, say $r_i \equiv \rho_i > 0$. This gives the recurrence $\omega(\mathbf{k} + \mathbf{e}_i, \lambda) = \rho_i \, \omega(\mathbf{k}, \lambda)$ for all $(\mathbf{k}, \lambda)$ and all i. Iterating in each coordinate:

$$\omega(\mathbf{k}, \lambda) = \omega(\mathbf{0}, \lambda) \, \rho_1^{k_1} \cdots \rho_r^{k_r}.$$

Setting $\eta(\lambda) := \omega(\mathbf{0}, \lambda)$ gives (ii).

Conversely, if (ii) holds, then $r_i(\mathbf{k}, \lambda) = \rho_i$ for all $(\mathbf{k}, \lambda)$, so $\|J_i\| = \rho_i$ and $\|C_i\| = \rho_i^{-1}$, giving $\|J_i\| \, \|C_i\| = 1$ for every i. □

Definition 14.4.2. A weight of the form

$$\omega_{\rho, \eta}(\mathbf{k}, \lambda) \;:=\; \eta(\lambda) \, \rho_1^{k_1} \cdots \rho_r^{k_r}, \qquad \boldsymbol{\rho} = (\rho_1, \dots, \rho_r) \in (0, \infty)^r,$$

is called *geometric*. The corresponding completion is denoted $X^p_{\rho, \eta}$.

Corollary 14.4.3. *For a geometric weight* $\omega_{\rho, \eta}$*,*

$$\|J_i\| = \rho_i, \qquad \|C_i\| = \rho_i^{-1} \qquad (1 \leq i \leq r).$$

In particular, J_i *is contractive if and only if* $\rho_i \leq 1$*, and* C_i *is contractive if and only if* $\rho_i \geq 1$.

Remark 14.4.4. The product inequality $\|J_i\| \, \|C_i\| \geq 1$ is a kind of uncertainty principle for the graded setting: the forward and backward shifts in any given coordinate cannot simultaneously be made small. Making the forward shift contractive forces the backward shift to be expansive, and vice versa. The geometric weights achieve the sharpest possible balance, with the identity $\|J_i\| = \rho_i$, $\|C_i\| = \rho_i^{-1}$ making the trade-off explicit. In the language of Chapter 8, this explains why geometric weights were the natural choice for the completion theory: they are the only weights that treat the shift pair optimally.

Corollary 14.4.5 (Simultaneously contractive weights)**.** *The following are equivalent:*

(i) $\|J_i\| \le 1$ *and* $\|C_i\| \le 1$ *for every* i*.*

(ii) The weight is independent of $\mathbf{k}$*:* $\omega(\mathbf{k}, \lambda) = \eta(\lambda)$ *for some positive function* η *on* Λ*.*

Proof. If (i) holds, then $1 \le \|J_i\|\,\|C_i\| \le 1$ for every i, so the weight is geometric with $\rho_i \le 1$ and $\rho_i^{-1} \le 1$. Hence $\rho_i = 1$ for all i, giving (ii). The converse is immediate. □

Example 14.4.6 ($r = 1$: the balance picture)**.** For a single one-sided coordinate, the weight is geometric if and only if $\omega(k, \lambda) = \eta(\lambda)\,\rho^k$ for some $\rho > 0$. The forward shift J has norm ρ and the backward shift C has norm ρ^{-1}. The parameter ρ controls the asymmetry: $\rho > 1$ favors the backward shift (making C contractive at the cost of expanding J), while $\rho < 1$ favors the forward shift. The choice $\rho = 1$ is the unique weight for which both J and C are isometries.

In the continuous realization, this means the Riemann–Liouville integral has norm ρ and the Caputo derivative has norm ρ^{-1}. In the discrete realization, the same norms hold for the nabla fractional sum and the Caputo nabla difference. The coincidence is structural: it follows from the common coefficient algebra.

14.5 The unified transform model

For geometric weights, the coefficient transform on a polydisk provides a common analytic realization that underlies both the continuous holomorphic model of Chapter 8 and the discrete Z-transform model of Chapter 12.

Definition 14.5.1. For $\boldsymbol{\rho} \in (0, \infty)^r$, define the open polydisk

$$\mathbb{D}_{\rho} := \{z = (z_1, \ldots, z_r) \in \mathbb{C}^r : |z_i| < \rho_i \text{ for all } i\}.$$

Definition 14.5.2. For $u = \sum a_{\mathbf{k},\lambda}\, \mathbf{e}_{\mathbf{k},\lambda} \in X^p_{\rho,\eta}$, define the *coefficient transform*

$$(\mathscr{T} u)(z, \lambda) := \sum_{\mathbf{k}\in\mathbb{N}_0^r} a_{\mathbf{k},\lambda}\, z^{\mathbf{k}}, \qquad z^{\mathbf{k}} := z_1^{k_1} \cdots z_r^{k_r}.$$

Proposition 14.5.3. *For every* $u \in X^p_{\rho,\eta}$ *and every* $\lambda \in \Lambda$*, the series defining* $(\mathscr{T} u)(\cdot, \lambda)$ *converges absolutely on* $\mathbb{D}_\rho$ *and defines a holomorphic function there. The map* $\mathscr{T}$ *is injective.*

Proof. Fix λ and $z \in \mathbb{D}_\rho$. For $1 < p < \infty$ with conjugate exponent q, Hölder's inequality gives

$$\sum_{\mathbf{k}} |a_{\mathbf{k},\lambda}\, z^{\mathbf{k}}| \le \frac{1}{\eta(\lambda)} \Big(\sum_{\mathbf{k}} |a_{\mathbf{k},\lambda}|^p\, \eta(\lambda)^p\, \boldsymbol{\rho}^{p\mathbf{k}} \Big)^{1/p} \prod_{i=1}^{r} \Big(\sum_{n=0}^{\infty} (|z_i|/\rho_i)^{qn} \Big)^{1/q}.$$

Each geometric series converges because $|z_i| < \rho_i$. For $p = 1$, the simpler estimate $|z^{\mathbf{k}}| \le \boldsymbol{\rho}^{\mathbf{k}}$ suffices. Weierstrass theory gives holomorphy.

If $\mathscr{T} u = 0$, then for each λ the holomorphic function vanishes on $\mathbb{D}_\rho$, so all Taylor coefficients at the origin are zero. Hence $a_{\mathbf{k},\lambda} = 0$ for all $(\mathbf{k}, \lambda)$ and $u = 0$. □

The key result is that the transform intertwines the abstract operators with familiar operations on holomorphic functions.

Definition 14.5.4. For a holomorphic function F on $\mathbb{D}_\rho$ and $i \in \{1, \dots, r\}$, define the *backward quotient operator*

$$(\mathcal{B}_i F)(z) := \frac{F(z) - F(z_1, \dots, z_{i-1}, 0, z_{i+1}, \dots, z_r)}{z_i},$$

with the removable singularity at $z_i = 0$ filled in holomorphically.

Theorem 14.5.5 (Unified transform identities)**.** *Let $u \in X^p_{\rho,\eta}$, let $\sigma : \Lambda \to \mathbb{C}$ be bounded, and let $i \in \{1, \dots, r\}$. Then:*

(i) $(\mathscr{T}(J_i u))(z, \lambda) = z_i \, (\mathscr{T} u)(z, \lambda)$.

(ii) $(\mathscr{T}(C_i u))(z, \lambda) = (\mathcal{B}_i \, (\mathscr{T} u)(\cdot, \lambda))(z)$.

(iii) $(\mathscr{T}(M_\sigma u))(z, \lambda) = \sigma(\lambda) \, (\mathscr{T} u)(z, \lambda)$.

In words: the forward shift becomes multiplication by z_i, the backward shift becomes the backward quotient $\mathcal{B}_i$, and every spectral multiplier remains pointwise diagonal in λ.

Proof. Write $u = \sum a_{\mathbf{k},\mu} \, \mathbf{e}_{\mathbf{k},\mu}$ and fix λ.

(i) Since $J_i \, \mathbf{e}_{\mathbf{k},\lambda} = \mathbf{e}_{\mathbf{k}+\mathbf{e}_i,\lambda}$,

$$(\mathscr{T}(J_i u))(z, \lambda) = \sum_{\mathbf{k}} a_{\mathbf{k},\lambda} \, z^{\mathbf{k}+\mathbf{e}_i} = z_i \sum_{\mathbf{k}} a_{\mathbf{k},\lambda} \, z^{\mathbf{k}} = z_i \, (\mathscr{T} u)(z, \lambda).$$

(ii) Since C_i drops the $k_i = 0$ terms and shifts the remaining ones,

$$\begin{aligned}(\mathscr{T}(C_i u))(z, \lambda) &= \sum_{k_i \ge 1} a_{\mathbf{k},\lambda} \, z^{\mathbf{k}-\mathbf{e}_i} = \frac{1}{z_i} \sum_{k_i \ge 1} a_{\mathbf{k},\lambda} \, z^{\mathbf{k}} \\ &= \frac{(\mathscr{T} u)(z, \lambda) - (\mathscr{T} u)(z_1, \dots, 0, \dots, z_r, \lambda)}{z_i},\end{aligned}$$

which is $\mathcal{B}_i \, (\mathscr{T} u)(\cdot, \lambda)$.

(iii) Since M_σ multiplies each coefficient by $\sigma(\lambda)$, the result is immediate. □

Remark 14.5.6. In the continuous realization, the transform $\mathscr{T}$ becomes the fiberwise holomorphic coefficient model of Chapter 8. In the discrete realization, it becomes the fiberwise Z-transform of Chapter 12. At the abstract level, there is no distinction: the same power series $\sum a_{\mathbf{k},\lambda} \, z^{\mathbf{k}}$ underlies both. The continuous–discrete difference enters only when one maps the abstract basis to a concrete function-space basis.

Generating eigenvectors revisited

The transform model also gives a unified treatment of the generating eigenvectors from Chapters 8 and 12.

Corollary 14.5.7 (Abstract generating eigenvectors)**.** *Fix $\lambda \in \Lambda$ and $\boldsymbol{\zeta} = (\zeta_1, \dots, \zeta_r) \in \mathbb{C}^r$ with $|\zeta_i| < \rho_i^{-1}$ for every i. Then*

$$E_{\zeta,\lambda} \;:=\; \sum_{\mathbf{k}\in\mathbb{N}_0^r} \boldsymbol{\zeta}^{\mathbf{k}}\, \mathbf{e}_{\mathbf{k},\lambda}$$

converges in $X^p_{\rho,\eta}$ and satisfies

$$C_i\, E_{\zeta,\lambda} = \zeta_i\, E_{\zeta,\lambda} \qquad (1 \le i \le r), \qquad M_\sigma\, E_{\zeta,\lambda} = \sigma(\lambda)\, E_{\zeta,\lambda}.$$

Under the transform, $(\mathscr{T}\, E_{\zeta,\lambda})(z,\lambda) = \prod_{i=1}^r (1-\zeta_i z_i)^{-1}$.

Proof. The norm is $\|E_{\zeta,\lambda}\|^p_{p,\rho,\eta} = \eta(\lambda)^p \prod_{i=1}^r (1-|\zeta_i\rho_i|^p)^{-1} < \infty$ because $|\zeta_i\rho_i| < 1$. Since C_i is bounded, applying it termwise (after reindexing $\mathbf{m} = \mathbf{k} - \mathbf{e}_i$) gives $C_i E_{\zeta,\lambda} = \zeta_i E_{\zeta,\lambda}$. The multiplier identity and the transform formula follow from the diagonal action and the geometric series. □

Remark 14.5.8. Under the continuous transport (Section 14.7), $E_{\zeta,\lambda}$ becomes the hybrid Mittag–Leffler vector of Chapter 8. Under the discrete transport, it becomes the discrete generating eigenvector of Chapter 12. The eigenvector property depends only on the shift algebra, not on the function-space realization.

14.6 Diagonal multipliers and C_0-semigroup generation

We now study the dynamics generated by the spectral block. This section introduces the minimum amount of semigroup theory needed for the application.

A brief orientation to C_0-semigroups

A *C_0-semigroup* on a Banach space X is a family of bounded operators $(T(t))_{t\ge 0}$ satisfying: (a) $T(0) = I$; (b) $T(t+s) = T(t)\,T(s)$ for all $t, s \ge 0$; (c) $\|T(t)u - u\| \to 0$ as $t \downarrow 0$ for every $u \in X$. The *generator* A of the semigroup is defined by $Au := \lim_{t\downarrow 0}(T(t)u - u)/t$ on the domain where this limit exists. The reader unfamiliar with semigroup theory may consult the standard references by Engel–Nagel or Pazy; we develop only what is needed here.

A bounded operator B always generates the *uniformly continuous* semigroup $T(t) = e^{tB} := \sum_{n=0}^{\infty}(tB)^n/n!$; the series converges in operator norm. For unbounded generators, semigroup generation is subtler. The Hille–Yosida theorem provides necessary and sufficient conditions, but in our setting the transparent structure of diagonal multipliers makes generation verifiable by hand.

Maximal multipliers

When the symbol σ is unbounded, the multiplier M_σ is no longer everywhere defined on X^p_ω, but it can be defined on its natural maximal domain.

Definition 14.6.1. For any $\sigma : \Lambda \to \mathbb{C}$, define the maximal domain

$$\operatorname{dom}(M_\sigma) \;:=\; \Big\{ u = \sum a_{\mathbf{k},\lambda}\, \mathbf{e}_{\mathbf{k},\lambda} \in X^p_\omega : \sum |\sigma(\lambda)|^p\, |a_{\mathbf{k},\lambda}|^p\, \omega(\mathbf{k},\lambda)^p < \infty \Big\},$$

and for $u \in \operatorname{dom}(M_\sigma)$ set $M_\sigma u := \sum \sigma(\lambda)\, a_{\mathbf{k},\lambda}\, \mathbf{e}_{\mathbf{k},\lambda}$.

Proposition 14.6.2. *For every $\sigma : \Lambda \to \mathbb{C}$, the operator M_σ is densely defined and closed on X^p_ω.*

Proof. The algebraic core is contained in $\operatorname{dom}(M_\sigma)$—every finite-support element involves only finitely many spectral labels—so the domain is dense.

For closedness: suppose $u_n \to u$ and $M_\sigma u_n \to v$ in X^p_ω. The coordinate functional $u \mapsto a_{\mathbf{k},\lambda}$ is continuous (with $|a_{\mathbf{k},\lambda}| \le \omega(\mathbf{k},\lambda)^{-1}\, \|u\|_{p,\omega}$), so $a^{(n)}_{\mathbf{k},\lambda} \to a_{\mathbf{k},\lambda}$ and $\sigma(\lambda)\, a^{(n)}_{\mathbf{k},\lambda} \to b_{\mathbf{k},\lambda}$. Hence $b_{\mathbf{k},\lambda} = \sigma(\lambda)\, a_{\mathbf{k},\lambda}$. Since $v \in X^p_\omega$, it follows that $u \in \operatorname{dom}(M_\sigma)$ and $M_\sigma u = v$. ☐

The exact semigroup criterion

Theorem 14.6.3 (Exact semigroup criterion)**.** *Let $\sigma : \Lambda \to \mathbb{C}$. Then M_σ generates a C_0-semigroup on X^p_ω if and only if*

$$\omega_0 \;:=\; \sup_{\lambda\in\Lambda} \operatorname{Re} \sigma(\lambda) \;<\; \infty.$$

In that case, the semigroup is

$$T_\sigma(t) \;:=\; M_{e^{t\sigma}} \qquad (t \ge 0),$$

with exact growth $\|T_\sigma(t)\| = e^{t\omega_0}$, and the generator is the maximal multiplier M_σ.

Proof. Sufficiency. Assume $\omega_0 < \infty$. For each $t \ge 0$, $|e^{t\sigma(\lambda)}| = e^{t \operatorname{Re} \sigma(\lambda)} \le e^{t\omega_0}$, so $T_\sigma(t) = M_{e^{t\sigma}}$ is bounded with $\|T_\sigma(t)\| = \sup_\lambda |e^{t\sigma(\lambda)}| = e^{t\omega_0}$.

The semigroup law follows from $M_{e^{t\sigma}} M_{e^{s\sigma}} = M_{e^{(t+s)\sigma}}$.

For strong continuity, it suffices to check on the dense algebraic core. A finite-support element u involves only finitely many spectral labels $\lambda_1, \dots, \lambda_N$, and $\|T_\sigma(t)u - u\| \to 0$ as $t \downarrow 0$ because each factor $e^{t\sigma(\lambda_\nu)} - 1 \to 0$.

Identification of the generator. For $u \in \operatorname{dom}(M_\sigma)$, the difference quotient $(T_\sigma(t)u - u)/t$ has coefficients $(e^{t\sigma(\lambda)} - 1)/t \cdot a_{\mathbf{k},\lambda}$. The estimate $|(e^{t\sigma(\lambda)} - 1)/t| \le e^{\omega_0}\, |\sigma(\lambda)|$ for $0 < t \le 1$, combined with $u \in \operatorname{dom}(M_\sigma)$, provides the domination needed for convergence in ℓ^p_ω norm to $M_\sigma u$.

Conversely, if $(T_\sigma(t)u - u)/t \to v$, then coordinatewise limits give $b_{\mathbf{k},\lambda} = \sigma(\lambda)\, a_{\mathbf{k},\lambda}$, so $u \in \operatorname{dom}(M_\sigma)$ and $v = M_\sigma u$. The generator is exactly M_σ.

Necessity. If M_σ generates a C_0-semigroup, then $\|T_\sigma(t)\| < \infty$ for every $t > 0$, forcing $\sup_\lambda e^{t \operatorname{Re} \sigma(\lambda)} < \infty$ and hence $\omega_0 < \infty$. ☐

Remark 14.6.4. The semigroup $T_\sigma(t) = M_{e^{t\sigma}}$ is "exactly diagonal": it multiplies each spectral fiber by $e^{t\sigma(\lambda)}$. This is the simplest possible semigroup structure. In the continuous realization, $\sigma(\lambda) = \lambda^\beta$ is a Weyl symbol and $T_\sigma(t)$ describes time evolution under a fractional Weyl operator. In the discrete realization, $\sigma(\lambda) = (1 - \lambda^{-1})^\beta$, and the semigroup has the same abstract form.

Corollary 14.6.5 (Resolvent and functional calculus)**.** *If $\inf_\lambda |\zeta - \sigma(\lambda)| > 0$, then $(\zeta I - M_\sigma)^{-1} = M_{(\zeta-\sigma)^{-1}}$ with $\|(\zeta I - M_\sigma)^{-1}\| = \sup_\lambda |\zeta - \sigma(\lambda)|^{-1}$.*

More generally, if $U \subset \mathbb{C}$ is open, $\sigma(\Lambda) \subset U$, and f is holomorphic on U with $f \circ \sigma$ bounded on Λ, then $f(M_\sigma) := M_{f\circ\sigma}$ is bounded with $\|f(M_\sigma)\| = \sup_\lambda |f(\sigma(\lambda))|$.

Proof. The resolvent formula is a direct computation on the maximal domain. The functional calculus statement follows from the definition of diagonal multipliers and their norm formula. □

14.7 Mixed generators and broader semigroup background

The spectral block generates diagonal semigroups; the shift block is bounded on any geometric completion. Combining them by the classical bounded perturbation theorem produces semigroups for mixed generators.

Bounded shift blocks

Definition 14.7.1. On a geometric completion $X^p_{\rho,\eta}$ and for $\mathbf{a}, \mathbf{b} \in \mathbb{C}^r$, define the bounded operators

$$B^+_{\mathbf{a}} := \sum_{i=1}^{r} a_i J_i, \qquad B^-_{\mathbf{b}} := \sum_{i=1}^{r} b_i C_i, \qquad B_{\mathbf{a},\mathbf{b}} := B^+_{\mathbf{a}} + B^-_{\mathbf{b}}.$$

Proposition 14.7.2. *On $X^p_{\rho,\eta}$,*

$$\|B_{\mathbf{a},\mathbf{b}}\| \;\le\; \sum_{i=1}^{r} \big(|a_i|\,\rho_i + |b_i|\,\rho_i^{-1}\big).$$

In particular, $B_{\mathbf{a},\mathbf{b}}$ generates the uniformly continuous semigroup $e^{tB_{\mathbf{a},\mathbf{b}}}$.

Proof. Corollary 14.4.3 gives $\|J_i\| = \rho_i$ and $\|C_i\| = \rho_i^{-1}$; the triangle inequality does the rest. □

The bounded perturbation theorem

We recall the classical tool for combining unbounded generators with bounded perturbations.

Theorem 14.7.3 (Bounded perturbation; Engel–Nagel, Pazy)**.** *Let A generate a C_0-semigroup $(T(t))_{t\ge 0}$ on a Banach space X with $\|T(t)\| \le Me^{\omega t}$. Let B be a bounded operator on X. Then $A + B$, with domain* $\operatorname{dom}(A)$*, generates a C_0-semigroup $(S(t))_{t\ge 0}$ satisfying $\|S(t)\| \le Me^{(\omega + M\|B\|)t}$.*

In our setting, $A = M_\sigma$ generates the diagonal semigroup $M_{e^{t\sigma}}$ with $M = 1$ and $\omega = \omega_0$, and $B = B_{\mathbf{a},\mathbf{b}}$ is bounded.

Mixed generators

Theorem 14.7.4 (Mixed shift–spectral generators)**.** *Let $\sigma : \Lambda \to \mathbb{C}$ satisfy $\omega_0 := \sup_\lambda \operatorname{Re}\, \sigma(\lambda) < \infty$, and let $\mathbf{a}, \mathbf{b} \in \mathbb{C}^r$. Define*

$$A_{\sigma,\mathbf{a},\mathbf{b}} \;:=\; M_\sigma + \sum_{i=1}^{r} a_i J_i + \sum_{i=1}^{r} b_i C_i,$$

with domain $\operatorname{dom}(M_\sigma)$*. Then $A_{\sigma,\mathbf{a},\mathbf{b}}$ generates a C_0-semigroup on $X^p_{\rho,\eta}$ with*

$$\big\|e^{tA_{\sigma,\mathbf{a},\mathbf{b}}}\big\| \;\le\; \exp\!\left(t\omega_0 + t\sum_{i=1}^{r}\big(|a_i|\,\rho_i + |b_i|\,\rho_i^{-1}\big)\right) \qquad (t \ge 0).$$

Proof. By Theorem 14.6.3, M_σ generates the C_0-semigroup $M_{e^{t\sigma}}$ with $\|M_{e^{t\sigma}}\| = e^{t\omega_0}$. By Proposition 14.7.2, $B_{\mathbf{a},\mathbf{b}}$ is bounded. Theorem 14.7.3 gives the result. □

Remark 14.7.5 (Transform picture for mixed generators). When $\mathbf{b} = \mathbf{0}$, the semigroup factors as $e^{tA_{\sigma,\mathbf{a},\mathbf{0}}} = M_{e^{t\sigma}}\, e^{t\sum a_i J_i}$, and the transform gives

$$\big(\mathscr{T}\big(e^{tA_{\sigma,\mathbf{a},\mathbf{0}}}u\big)\big)(z,\lambda) \;=\; e^{t(\sigma(\lambda)+a_1z_1+\cdots+a_rz_r)}\,(\mathscr{T}u)(z,\lambda).$$

For general $\mathbf{b}$, the backward quotient operators $\mathcal{B}_i$ enter, and the transformed orbit $F(t,z,\lambda) := (\mathscr{T}(e^{tA_{\sigma,\mathbf{a},\mathbf{b}}}u))(z,\lambda)$ satisfies the evolution equation

$$\partial_t F = \big(\sigma(\lambda) + a_1z_1 + \cdots + a_rz_r\big)F + \sum_{i=1}^{r} b_i\,\mathcal{B}_iF$$

on the polydisk, reflecting the non-local character of the backward shifts.

Optimization of the geometric parameters

The growth bound in Theorem 14.7.4 depends on $\boldsymbol{\rho}$ through $\Phi_{\mathbf{a},\mathbf{b}}(\boldsymbol{\rho}) := \sum_{i=1}^{r}(|a_i|\,\rho_i + |b_i|\,\rho_i^{-1})$. For a given mixed generator, which geometric weight minimizes the growth bound?

Theorem 14.7.6 (Optimal geometric parameters)**.** *Fix* $\mathbf{a},\mathbf{b} \in \mathbb{C}^r$.

(i) If $a_ib_i \neq 0$ *for every* i, *the functional* $\Phi_{\mathbf{a},\mathbf{b}}$ *has a unique minimizer*

$$\rho_i^* = \sqrt{|b_i|/|a_i|} \qquad (1 \le i \le r),$$

with minimal value $\Phi_{\mathbf{a},\mathbf{b}}(\boldsymbol{\rho}^*) = 2\sum_{i=1}^{r}\sqrt{|a_ib_i|}$.

(ii) If $a_i = 0$ *and* $b_i \neq 0$, *the infimum of* $|b_i|\,\rho_i^{-1}$ *over* $\rho_i > 0$ *is* 0, *approached as* $\rho_i \to \infty$. *If* $b_i = 0$ *and* $a_i \neq 0$, *the infimum of* $|a_i|\,\rho_i$ *is* 0, *approached as* $\rho_i \to 0$.

Proof. The functional separates coordinatewise. For $a_ib_i \neq 0$, set $\phi_i(\rho) := |a_i|\,\rho + |b_i|\,\rho^{-1}$. Then $\phi_i'(\rho) = |a_i| - |b_i|\,\rho^{-2}$ vanishes at $\rho = \sqrt{|b_i|/|a_i|}$, and $\phi_i''(\rho) = 2|b_i|\,\rho^{-3} > 0$ confirms the minimum. The value is $2\sqrt{|a_ib_i|}$. □

Example 14.7.7. With $r = 1$, $a = b = 1$, and $\omega_0 = 0$, the growth bound is $\exp(t(\rho + \rho^{-1}))$. The optimum is $\rho^* = 1$ with bound e^{2t}. At $\rho = 1$ both J and C are isometries: the balanced weight gives the tightest growth control.

Example 14.7.8. With $r = 1$, $a = 1$, $b = 4$, the optimum is $\rho^* = 2$ with minimal contribution $1 \cdot 2 + 4 \cdot \frac{1}{2} = 4 = 2\sqrt{4}$. The optimal weight tilts toward larger ρ, making C more contractive ($\|C\| = \frac{1}{2}$) to compensate for the stronger backward-shift coefficient.

Remark 14.7.9. Theorem 14.7.6 and Theorem 14.4.1 together give two complementary characterizations of the geometric weights. Algebraically, they are the unique balanced weights for the shift algebra. Analytically, they are the weights that optimize the basic semigroup-growth bound for mixed shift–spectral generators. The geometric weight parameters ρ_i are therefore not merely a convenient convention; they are control parameters that the analyst can tune to the problem at hand, with the product $\rho_i = \sqrt{|b_i|/|a_i|}$ as the natural balance point.

Transport to the continuous and discrete theories

We conclude by explaining how every result of this chapter specializes simultaneously to both completion theories.

Definition 14.7.10. Let $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r) \in (0,1)^r$.

(i) In the *continuous case*, let $\Lambda_{\mathrm{c}} \subset (\mathbb{C}_+)^s$ with $\mathbb{C}_+ := \{z \in \mathbb{C} : \operatorname{Re} z > 0\}$, and define the *continuous transport map* $U_{\mathrm{c}} : \mathbf{e}_{\mathbf{k},\lambda} \longmapsto e^{\mathrm{cont}}_{\mathbf{k},\lambda}$, where

$$e^{\mathrm{cont}}_{\mathbf{k},\lambda}(x,y) \;:=\; \left(\prod_{i=1}^{r} \frac{x_i^{k_i\alpha_i}}{\Gamma(k_i\alpha_i+1)}\right) e^{\langle \lambda, y\rangle}.$$

(ii) In the *discrete case*, let $\Lambda_{\mathrm{d}} \subset \mathbb{E}^s$ with $\mathbb{E} := \{z \in \mathbb{C} : |z| > 1\}$, and define $U_{\mathrm{d}} : \mathbf{e}_{\mathbf{k},\lambda} \longmapsto e^{\mathrm{disc}}_{\mathbf{k},\lambda}$, where

$$e^{\mathrm{disc}}_{\mathbf{k},\lambda}(n,m) \;:=\; \left(\prod_{i=1}^{r} \frac{n_i^{\overline{k_i\alpha_i}}}{\Gamma(k_i\alpha_i+1)}\right) \lambda^{m}.$$

Theorem 14.7.11 (Transport theorem)**.** *Each transport map extends to an isometric isomorphism from the abstract completion X^p_ω onto the corresponding Banach completion of Chapter 8 (for U_{c}) or Chapter 12 (for U_{d}). Under these isomorphisms:*

(i) the abstract forward shifts J_i become the one-sided Riemann–Liouville integrals (resp. nabla fractional sums);

(ii) the abstract backward shifts C_i become the Caputo derivatives (resp. Caputo nabla differences);

(iii) the abstract multipliers M_σ become the Weyl spectral multipliers;

(iv) the abstract transform $\mathscr{T}$ becomes the fiberwise holomorphic model (resp. the fiberwise Z-transform model).

Proof. In both settings, the Banach completions are defined coefficientwise relative to the canonical bases. The transport map sends the abstract basis to the concrete basis, preserving the coefficient norm by construction, and therefore extends isometrically to the completion. The transport of operators follows from the basis-level identities proved in the earlier chapters: the concrete operators act on the concrete basis exactly as J_i, C_i, and M_σ act on $\mathbf{e}_{\mathbf{k},\lambda}$. The transform transport is immediate because $\mathscr{T}$ depends only on the coefficient sequence. □

Corollary 14.7.12. *All results of this chapter—exact norm formulae, the optimal-weight theorem, the unified transform model, the semigroup criterion, the mixed generator theorem, and the optimal geometric parameters—apply verbatim to both the continuous and discrete completion theories through the transport maps. In particular:*

(i) the geometric weights are the balanced weights in both realizations;

(ii) every diagonal Weyl-symbol multiplier with bounded real part generates a C_0-semigroup in both realizations;

(iii) every mixed shift–spectral generator produces a C_0-semigroup in both realizations;

(iv) the geometric parameters minimizing the growth bound are the same in both cases.

Proof. Every proof in this chapter is purely coefficient-theoretic and therefore invariant under the isometric transport. □

Remark 14.7.13 (The central message). Corollary 14.7.12 is the main conceptual conclusion of this chapter. The continuous and discrete hybrid calculi have different underlying domains, different function-space bases, different pointwise kernels, and different spectral conditions. But their completion theories, transform models, semigroup generation, and optimal-weight structures are governed by the same coefficient algebra. The distinction between the two branches lies entirely in the concrete realization of the basis, not in the analytic mechanism.

This chapter therefore completes the analytic unification of the continuous and discrete programs. Together with the boundary-augmentation chapters (Chapters 9 and 13), which demonstrated the identical combinatorial structure of noncommutativity in both settings, it establishes that the parallel between the two branches is a structural identity grounded in the common coefficient lattice $\mathbb{N}_0^r$.

Looking ahead

This chapter has placed the continuous and discrete completion theories in a common abstract framework. The three conclusions are: (a) the geometric weights are intrinsically balanced, (b) the transform model is the same in both realizations, and (c) semigroup generation for mixed shift–spectral operators follows from the diagonal–bounded decomposition.

The final chapter of the book returns to the boundary-augmented theory and asks a question complementary to that of Chapters 9 and 13. Those chapters recovered commutativity by *restriction*: they identified the maximal graded sector on which the extended Caputo tuple commutes and discarded the trace blocks whose free rank was too large. Chapter 15 asks instead: what happens if one *quotients* the ambient space by the commutator ideal? The answer leads to the theory of ordered boundary traces, commutator-generated defect spaces, and the universal commuting quotient—the culminating algebraic structure of the entire AD01–AD12 research program.

Chapter 15

Ordered Boundary Trace Calculus and Commutator Ideals

Chapters 9 and 13 posed the question: on what part of the boundary-augmented ambient space does the extended Caputo tuple remain commuting? The answer was obtained by *restriction*—discarding the trace blocks whose free rank is too large. The resulting maximal commuting sectors are natural and cleanly characterized, but the restriction strategy discards information: the trace blocks with $q(w) \geq 2$ are simply removed.

This chapter asks a different, and complementary, question. Instead of asking "where does the tuple commute?" it asks "what exactly *is* the obstruction to commutativity, and can one quotient it out?" The answer requires a deeper analysis of the trace mechanism. One must understand not merely that the commutator is nonzero on certain strata, but that the commutator generates a precise invariant subspace—the *ordering-defect space*—and that quotienting by this subspace produces a *universal commuting reduction* with a clean algebraic structure.

The conceptual shift is from classification to structure theory. The maximal-sector theorems classify subspaces; this chapter constructs an ideal and a quotient. In the language of ring theory, this is the difference between finding a maximal subring on which a relation holds and factoring out the ideal generated by the relation's failure.

Because the entire analysis is coefficient-theoretic, it applies simultaneously to the continuous ambient model of Chapter 9 and the discrete ambient model of Chapter 13. Following the unification philosophy of Chapter 14, we work throughout in the abstract boundary-augmented coefficient space.

The material follows the research paper AD12. Throughout, we fix $r \geq 1$, a spectral index set Λ, geometric weight parameters $\boldsymbol{\rho} \in (0, \infty)^r$, a trace-weight parameter $\tau > 0$, a spectral weight $\eta : \Lambda \to (0, \infty)$, and $1 \leq p < \infty$.

15.1 From classification to structure theory

The boundary-augmented theory of Chapters 9 and 13 established three central facts:

(i) the extended Caputo operators are bounded on the ambient completion and commute with every spectral multiplier;

(ii) the commutator $[C_i, C_j]$ is nonzero only when two free one-sided coordinates are simultaneously at vacuum inside a trace block, and the commutator is the difference of two ordered trace states;

(iii) the maximal graded invariant commuting sector consists of the canonical completion plus all trace blocks with free rank at most one.

These results tell us *where* commutativity holds and where it fails. They do not tell us what the failure *means* in structural terms.

Consider the commutator formula from Chapter 9:

$$[C_i, C_j]\, t_{w,\mathbf{0},\lambda} = t_{wji,\varnothing,\lambda} - t_{wij,\varnothing,\lambda}.$$

The right-hand side is the difference of two trace states that carry the *same boundary support* $\mathrm{supp}(w) \cup \{i, j\}$ but differ in the *ordering* of the last two letters. This suggests a precise algebraic interpretation: the commutator does not create genuinely new boundary information; it merely detects the ambiguity in the order of boundary extraction.

If one could identify all trace states that differ only in their ordering—collapsing the ordered word w to its unordered support $\mathrm{supp}(w)$—then the resulting quotient space should carry a commuting tuple. This is the idea that the present chapter makes rigorous.

Remark 15.1.1. For a reader who has studied ring theory, the analogy is helpful. In a noncommutative ring R, one often studies the two-sided ideal I generated by all commutators $[a, b] = ab - ba$. The quotient R/I is the largest commutative quotient of R, and it has a universal property: every ring homomorphism from R to a commutative ring factors through R/I. The construction of this chapter is the Banach-space operator-algebra analogue of that classical algebraic construction.

15.2 Ordered boundary trace operators

We work in the abstract boundary-augmented coefficient space of Chapters 9, 13, and 14. We recall the notation briefly: the ambient Banach space is

$$\mathfrak{X}^p_{\rho,\tau,\eta} \;=\; X^p_{\rho,\eta} \;\oplus\; \bigoplus_{w \in \mathcal{W}_r} T^p_w,$$

where $X^p_{\rho,\eta}$ is the canonical completion and T^p_w is the closure of the algebraic trace block for the word w. The extended Caputo operators C_i act as canonical backward shifts on the core, as residual backward shifts on the trace blocks when the relevant coordinate has positive grade, and as ordered trace-creation operators when the grade reaches zero.

The trace-creation mechanism was introduced in Chapters 9 and 13 as part of the definition of the extended operators. In this section we isolate it as an independent family of operators and study its algebraic properties.

Vacuum projections

Definition 15.2.1. For each $i \in \{1, \ldots, r\}$, define the *core vacuum projection* $\Pi^{\mathrm{core}}_i : X^p_{\rho,\eta} \to X^p_{\rho,\eta}$ by

$$\Pi^{\mathrm{core}}_i\, e_{\mathbf{k},\lambda} := \begin{cases} e_{\mathbf{k},\lambda}, & k_i = 0, \\ 0, & k_i \geq 1. \end{cases}$$

For each word $w \in \mathcal{W}_r$ and each $i \in \mathrm{F}(w)$, define the *residual vacuum projection* $\Pi^{(w)}_i : T^p_w \to$

T^p_w by

$$\Pi^{(w)}_i \, t_{w,\mathbf{k},\lambda} := \begin{cases} t_{w,\mathbf{k},\lambda}, & k_i = 0, \\ 0, & k_i \geq 1. \end{cases}$$

If $i \notin \mathrm{F}(w)$, set $\Pi^{(w)}_i := 0$.

The vacuum projections identify the part of each block that sits at grade zero in a given coordinate. It is precisely on the simultaneous vacuum subspace $\Pi^{(w)}_i \Pi^{(w)}_j T^p_w$ that the commutator acts nontrivially.

Trace-creation operators

Definition 15.2.2. For each $i \in \{1, \dots, r\}$, define the *initial trace operator* $\mathrm{Tr}^{\mathrm{core}}_i : X^p_{\rho,\eta} \to T^p_{(i)}$ by

$$\mathrm{Tr}^{\mathrm{core}}_i \, e_{\mathbf{k},\lambda} := \begin{cases} t_{(i),\mathbf{k}^{\hat{i}},\lambda}, & k_i = 0, \\ 0, & k_i \geq 1. \end{cases}$$

For each word $w \in \mathcal{W}_r$ and each $i \in \mathrm{F}(w)$, define the *residual trace operator* $\mathrm{Tr}^{(w)}_i : T^p_w \to T^p_{wi}$ by

$$\mathrm{Tr}^{(w)}_i \, t_{w,\mathbf{k},\lambda} := \begin{cases} t_{wi,\mathbf{k}^{\hat{i}},\lambda}, & k_i = 0, \\ 0, & k_i \geq 1. \end{cases}$$

If $i \notin \mathrm{F}(w)$, set $\mathrm{Tr}^{(w)}_i := 0$.

The initial trace operator extracts the boundary event at coordinate i from a core vector at vacuum; the residual trace operator does the same from a trace-block vector whose i-th residual grade is zero. In both cases, the operator appends the letter i to the trace word and removes i from the set of free coordinates.

Definition 15.2.3 (Ordered trace operators)**.** For a word $w = (i_1, \dots, i_m) \in \mathcal{W}_r$ with $m \geq 1$, define the *ordered trace operator associated with* w by the composition

$$\mathrm{Tr}_w \; := \; \mathrm{Tr}^{(i_1 \cdots i_{m-1})}_{i_m} \circ \cdots \circ \mathrm{Tr}^{(i_1)}_{i_2} \circ \mathrm{Tr}^{\mathrm{core}}_{i_1} \; : \; X^p_{\rho,\eta} \longrightarrow T^p_w.$$

For the empty word, set $\mathrm{Tr}_\varnothing := I$ on $X^p_{\rho,\eta}$.

Proposition 15.2.4 (Trace norms)**.** *For every $w \in \mathcal{W}_r$, the ordered trace operator Tr_w is bounded with $\| \mathrm{Tr}_w \| = \tau^{\ell(w)+1}$. Each residual trace operator satisfies $\| \mathrm{Tr}^{(w)}_i \| = \tau$. Moreover, every bounded spectral multiplier commutes with every trace operator.*

Proof. On a basis vector $t_{w,\mathbf{k},\lambda}$ with $k_i = 0$, the residual trace operator $\mathrm{Tr}^{(w)}_i$ maps to $t_{wi,\mathbf{k}^{\hat{i}},\lambda}$, and the norm ratio is

$$\frac{\tau^{\ell(w)+2} \, \eta(\lambda) \, \rho^{\mathbf{k}^{\hat{i}}}}{\tau^{\ell(w)+1} \, \eta(\lambda) \, \rho^{\mathbf{k}}} = \tau,$$

because $k_i = 0$ implies $\boldsymbol{\rho}^{\mathbf{k}^{\widehat{i}}} = \boldsymbol{\rho}^{\mathbf{k}}$. The initial trace operator satisfies $\|\,\mathrm{Tr}_i^{\mathrm{core}}\,\| = \tau^2$ by the same computation with the canonical-block norm factor. Composing gives $\|\,\mathrm{Tr}_w\,\| \le \tau^{\ell(w)+1}$, and equality is achieved on any core basis vector at simultaneous vacuum in all coordinates of $\mathrm{supp}(w)$.

Multiplier commutation holds because the trace operators act only on the one-sided indices and the word, while M_σ acts only on the spectral scalar. □

The following theorem gives the explicit action of the ordered trace operator on a core basis vector.

Theorem 15.2.5 (Explicit ordered trace formula)**.** *Let $w \in \mathcal{W}_r$ with support $S := \mathrm{supp}(w)$. Then for every core basis vector $e_{\mathbf{k},\lambda}$,*

$$\mathrm{Tr}_w\; e_{\mathbf{k},\lambda} \;=\; \begin{cases} t_{w,\mathbf{k}^{\widehat{S}},\lambda}, & k_i = 0 \textit{ for all } i \in S, \\ 0, & \textit{otherwise.} \end{cases}$$

In particular, the domain condition (all coordinates in S at vacuum) depends only on the support *of w, while the target basis vector remembers the* order *encoded by w.*

Proof. By induction on $\ell(w)$. The case $\ell(w) = 0$ is trivial, and $\ell(w) = 1$ is the definition of the initial trace operator. For the inductive step, $\mathrm{Tr}_{wi} = \mathrm{Tr}_i^{(w)} \circ \mathrm{Tr}_w$. If any coordinate in $\mathrm{supp}(w)$ is not at vacuum, then Tr_w annihilates; if they are all at vacuum but $k_i \ge 1$, then $\mathrm{Tr}_i^{(w)}$ annihilates the result; if all coordinates in $\mathrm{supp}(w) \cup \{i\}$ are at vacuum, the composition produces $t_{wi,\mathbf{k}^{\widehat{S\cup\{i\}}},\lambda}$. □

Remark 15.2.6. Theorem 15.2.5 encapsulates the central asymmetry of the trace calculus. Two words w and w' with the same support S act on the same set of core vectors (those at simultaneous vacuum in S), but they produce different trace states (in different trace blocks T_w^p and $T_{w'}^p$). It is precisely this asymmetry that generates the commutator and the ordering-defect space.

15.3 Defect decomposition

We now decompose each extended Caputo operator on the trace blocks into two structurally transparent pieces: a residual backward shift and a trace-creation term.

Definition 15.3.1. For each $w \in \mathcal{W}_r$ and each $i \in \mathrm{F}(w)$, define the *residual backward shift* $S_i^{(w)} : T_w^p \to T_w^p$ by

$$S_i^{(w)}\, t_{w,\mathbf{k},\lambda} := \begin{cases} 0, & k_i = 0, \\ t_{w,\mathbf{k}-\mathbf{e}_i,\lambda}, & k_i \ge 1. \end{cases}$$

If $i \notin \mathrm{F}(w)$, set $S_i^{(w)} := 0$.

The residual backward shift is simply the "ordinary" part of the extended Caputo operator: it lowers a positive residual grade without touching the trace word. It is the trace-creation part that is new.

Proposition 15.3.2 (Defect decomposition)**.** *On each trace block T_w^p, one has*

$$C_i\big|_{T_w^p} \;=\; S_i^{(w)} \;+\; \mathrm{Tr}_i^{(w)} \qquad (1 \le i \le r).$$

Proof. This is a restatement of the basis action from the definition of the extended Caputo operator (Definition 9.3.1 of Chapter 9). If $i \notin \mathrm{F}(w)$, all three operators vanish. If $i \in \mathrm{F}(w)$ and $k_i \geq 1$, then C_i lowers the grade (the residual shift) and the trace operator vanishes. If $k_i = 0$, the residual shift vanishes and the trace operator appends the letter i. The sum reproduces C_i on every basis vector. □

The defect decomposition immediately clarifies the commutator.

Theorem 15.3.3 (Commutator factorization via traces)**.** *Let $w \in \mathcal{W}_r$ and let $i, j \in \mathrm{F}(w)$ be distinct. On the simultaneous vacuum subspace $V^{(w)}_{i,j} := \Pi^{(w)}_i \Pi^{(w)}_j T^p_w$,*

$$[C_i, C_j]\big|_{V^{(w)}_{i,j}} = \mathrm{Tr}^{(wj)}_i \circ \mathrm{Tr}^{(w)}_j - \mathrm{Tr}^{(wi)}_j \circ \mathrm{Tr}^{(w)}_i .$$

On the complementary subspace $(I - \Pi^{(w)}_i \Pi^{(w)}_j) T^p_w$, the commutator vanishes.

Proof. Off the simultaneous vacuum, at least one of the residual grades k_i, k_j is positive. In that case both compositions $C_i C_j$ and $C_j C_i$ act by lowering residual coordinates, and these operations commute.

On $V^{(w)}_{i,j}$, both residual shifts $S^{(w)}_i$ and $S^{(w)}_j$ vanish. By Proposition 15.3.2, $C_i|_{V^{(w)}_{i,j}} = \mathrm{Tr}^{(w)}_i$ and $C_j|_{V^{(w)}_{i,j}} = \mathrm{Tr}^{(w)}_j$. Composing in the two orders gives the formula. □

Remark 15.3.4. The commutator factorization reveals the mechanism transparently. On the simultaneous vacuum, $C_i C_j$ applies $\mathrm{Tr}^{(w)}_j$ first (recording the letter j), then $\mathrm{Tr}^{(wj)}_i$ (recording i), producing a trace word ending in ji. Reversing the order produces a trace word ending in ij. The commutator is the difference of these two recordings. It does not create new boundary data; it detects the ambiguity in ordering the same boundary data.

15.4 The ordering-defect space and the commutator ideal

The commutator formula shows that the obstruction to commutativity lives in a very specific subspace of the ambient completion: the subspace spanned by differences of trace states that share the same boundary support but differ in their ordering.

Definition 15.4.1. For a subset $S \subseteq \{1, \ldots, r\}$, let $\mathcal{W}(S) := \{w \in \mathcal{W}_r : \mathrm{supp}(w) = S\}$ be the set of all words with support S. Define the *algebraic order-difference space for S* by

$$\mathfrak{N}^{\mathrm{alg}}_S := \mathrm{span}\big\{ t_{w,\mathbf{k},\lambda} - t_{w',\mathbf{k},\lambda} : w, w' \in \mathcal{W}(S),\ \mathbf{k} \in \mathbb{N}_0^{\{1,\ldots,r\}\setminus S},\ \lambda \in \Lambda \big\}.$$

The *ordering-defect space* is

$$\mathfrak{N} := \overline{\sum_{S \subseteq \{1,\ldots,r\}} \mathfrak{N}^{\mathrm{alg}}_S}^{\mathfrak{X}^p_{\rho,\tau,\eta}} .$$

In words: $\mathfrak{N}$ is the closed span of all differences between trace states that have the same boundary support, the same residual grades, and the same spectral label, but are indexed by different orderings of the same set of boundary coordinates.

Example 15.4.2. For $r = 2$, the only nontrivial order-difference space is $\mathfrak{N}^{\mathrm{alg}}_{\{1,2\}}$, spanned by the differences $t_{(1,2),\mathbf{k},\lambda} - t_{(2,1),\mathbf{k},\lambda}$ for all $(\mathbf{k}, \lambda)$ with $\mathbf{k} \in \mathbb{N}_0^{\varnothing} = \{0\}$. (Since $\{1, \dots, r\} \setminus S = \varnothing$, the residual index is trivial.) Thus $\mathfrak{N}$ identifies the two corner traces for each spectral label λ.

Proposition 15.4.3 (Invariance of the ordering-defect space)**.** *The space $\mathfrak{N}$ is a closed invariant subspace of $\mathfrak{X}^p_{\rho,\tau,\eta}$ for every C_i and every bounded spectral multiplier M_σ.*

Proof. Closedness is built into the definition. It suffices to verify invariance of the algebraic generators. Let $d = t_{w,\mathbf{k},\lambda} - t_{w',\mathbf{k},\lambda}$ with $\operatorname{supp}(w) = \operatorname{supp}(w') = S$.

For a multiplier: $M_\sigma\, d = \sigma(\lambda)(t_{w,\mathbf{k},\lambda} - t_{w',\mathbf{k},\lambda}) \in \mathfrak{N}^{\mathrm{alg}}_S$.

For C_i: if $i \notin \mathrm{F}(w)$, both terms are annihilated. If $i \in \mathrm{F}(w)$ and $k_i \geq 1$, then $C_i\, d = t_{w,\mathbf{k}-\mathbf{e}_i,\lambda} - t_{w',\mathbf{k}-\mathbf{e}_i,\lambda} \in \mathfrak{N}^{\mathrm{alg}}_S$. If $i \in \mathrm{F}(w)$ and $k_i = 0$, then $C_i\, d = t_{wi,\mathbf{k}^{\hat{i}},\lambda} - t_{w'i,\mathbf{k}^{\hat{i}},\lambda}$, and $\operatorname{supp}(wi) = \operatorname{supp}(w'i) = S \cup \{i\}$, so the result lies in $\mathfrak{N}^{\mathrm{alg}}_{S\cup\{i\}}$. ☐

Theorem 15.4.4 (Commutator ranges lie in $\mathfrak{N}$)**.** *For all distinct $i, j \in \{1, \dots, r\}$,*

$$[C_i, C_j]\big(\mathfrak{X}^p_{\boldsymbol{\rho},\tau,\eta}\big) \;\subseteq\; \mathfrak{N}.$$

Proof. On the canonical block the commutator vanishes. On a trace basis vector $t_{w,\mathbf{k},\lambda}$, the commutator is either zero or equal to $t_{wji,\mathbf{k}^{\widehat{\{i,j\}}},\lambda} - t_{wij,\mathbf{k}^{\widehat{\{i,j\}}},\lambda}$ (Theorem 9.4.1 of Chapter 9). The two words wji and wij have the same support, so the difference lies in $\mathfrak{N}$. ☐

The key theorem of this section says that the ordering-defect space is *exactly* the invariant subspace generated by the commutator ranges—no larger and no smaller.

Definition 15.4.5. Let $\mathcal{G}$ denote the smallest closed subspace of $\mathfrak{X}^p_{\rho,\tau,\eta}$ that is

(i) invariant under every C_i and every bounded multiplier M_σ, and

(ii) contains the ranges of all commutators $[C_i, C_j]$ for $i \neq j$.

Lemma 15.4.6 (Adjacent transposition generation)**.** *Let $w = u\, i\, j\, v$ and $w' = u\, j\, i\, v$ be two words that differ by a single adjacent transposition (where u, v are subwords and all letters are distinct). Then for every $\mathbf{k} \in \mathbb{N}_0^{\mathrm{F}(w)}$ and $\lambda \in \Lambda$,*

$$t_{w,\mathbf{k},\lambda} - t_{w',\mathbf{k},\lambda} \;\in\; \mathcal{G}.$$

Proof. Construct a basis vector $t_{u,\mathbf{m},\lambda}$ in T^p_u by setting the residual grades at coordinates i, j, and all letters of v equal to zero, and setting the remaining residual grades equal to the corresponding components of $\mathbf{k}$.

By the commutator formula (Theorem 15.3.3),

$$[C_j, C_i]\, t_{u,\mathbf{m},\lambda} = t_{uij,\,\mathbf{m}^{\widehat{\{i,j\}}},\,\lambda} - t_{uji,\,\mathbf{m}^{\widehat{\{i,j\}}},\,\lambda},$$

which belongs to $\mathcal{G}$ by definition. Since $\mathcal{G}$ is invariant under each C_{v_ν}, applying the letters of v in order yields

$$C_{v_{\ell(v)}} \cdots C_{v_1}\, [C_j, C_i]\, t_{u,\mathbf{m},\lambda} = t_{uijv,\,\mathbf{k},\,\lambda} - t_{ujiv,\,\mathbf{k},\,\lambda},$$

which is the claimed adjacent-transposition difference. ☐

Theorem 15.4.7 (The ordering-defect space equals $\mathcal{G}$)**.** *One has $\mathcal{G} = \mathfrak{N}$. That is, $\mathfrak{N}$ is the smallest closed invariant subspace containing all commutator ranges.*

Proof. *$\mathcal{G} \subseteq \mathfrak{N}$.* By Proposition 15.4.3, $\mathfrak{N}$ is closed and invariant. By Theorem 15.4.4, it contains all commutator ranges. Hence $\mathcal{G} \subseteq \mathfrak{N}$.

$\mathfrak{N} \subseteq \mathcal{G}$. It suffices to show that every algebraic generator $t_{w,\mathbf{k},\lambda} - t_{w',\mathbf{k},\lambda}$ with $\operatorname{supp}(w) = \operatorname{supp}(w')$ belongs to $\mathcal{G}$. Since any two permutations of a finite set are related by a finite sequence of adjacent transpositions, there exist words $w = w_0, w_1, \dots, w_m = w'$, all with the same support, such that each pair $(w_{\nu-1}, w_\nu)$ differs by one adjacent transposition. Then

$$t_{w,\mathbf{k},\lambda} - t_{w',\mathbf{k},\lambda} = \sum_{\nu=1}^{m} \bigl(t_{w_{\nu-1},\mathbf{k},\lambda} - t_{w_\nu,\mathbf{k},\lambda}\bigr).$$

Each summand belongs to $\mathcal{G}$ by Lemma 15.4.6. Hence the full difference lies in $\mathcal{G}$. □

Remark 15.4.8. The proof uses the classical fact that the symmetric group $\mathfrak{S}_m$ is generated by adjacent transpositions. This is the link between the combinatorics of permutation groups and the operator theory of boundary-augmented fractional calculus. The adjacent-transposition generation lemma shows that every ordering difference is reachable from the commutator ranges by the ambient operator algebra; the symmetric-group decomposition assembles these elementary differences into arbitrary permutations.

We now formalize the connection with the commutator ideal.

Definition 15.4.9. Let

$$\mathfrak{A} := \overline{\operatorname{Alg}}\bigl\{C_1, \dots, C_r,\ M_\sigma : \sigma \in \ell^\infty(\Lambda)\bigr\} \ \subseteq\ \mathcal{B}(\mathfrak{X}^p_{\rho,\tau,\eta})$$

be the norm-closed operator algebra generated by the extended Caputo tuple and the bounded spectral multipliers. The *commutator ideal* is the closed two-sided ideal

$$\mathfrak{J}_{\mathrm{com}} := \overline{\mathfrak{A}\left\{[C_i, C_j]\right\}_{i \neq j} \mathfrak{A}}^{\|\cdot\|}.$$

For any subset $\mathfrak{J} \subseteq \mathcal{B}(\mathfrak{X}^p_{\rho,\tau,\eta})$, define its *range space* by

$$\mathcal{R}(\mathfrak{J}) := \overline{\operatorname{span}}\bigl\{Tx : T \in \mathfrak{J},\ x \in \mathfrak{X}^p_{\rho,\tau,\eta}\bigr\}.$$

Theorem 15.4.10 (Commutator ideal theorem)**.** *The range space of the commutator ideal is exactly the ordering-defect space:*

$$\mathcal{R}(\mathfrak{J}_{\mathrm{com}}) = \mathfrak{N}.$$

Proof. By Theorem 15.4.7, $\mathfrak{N}$ is the smallest closed invariant subspace containing all commutator ranges. The range space $\mathcal{R}(\mathfrak{J}_{\mathrm{com}})$ is closed, invariant under $\mathfrak{A}$, and contains the commutator ranges (since $\mathfrak{J}_{\mathrm{com}}$ contains the commutators themselves). Hence $\mathfrak{N} \subseteq \mathcal{R}(\mathfrak{J}_{\mathrm{com}})$.

Conversely, each commutator has range in $\mathfrak{N}$ by Theorem 15.4.4. Since $\mathfrak{N}$ is invariant under every generator of $\mathfrak{A}$, every operator of the form $A[C_i, C_j]B$ with $A, B \in \mathfrak{A}$ also has range in $\mathfrak{N}$. Passing to the norm closure gives $\mathcal{R}(\mathfrak{J}_{\mathrm{com}}) \subseteq \mathfrak{N}$. □

Remark 15.4.11. Theorem 15.4.10 is the structural heart of this chapter. It says that the commutator ideal does not merely detect *some* portion of the defect sector; it detects *exactly* the part generated by ordered boundary ambiguity. The ordering-defect space $\mathfrak{N}$ is the operator-theoretic shadow of the algebraic commutator ideal, and the two objects determine each other completely.

15.5 The universal commuting quotient

Since the ordering-defect space is a closed invariant subspace, we may form the quotient.

Definition 15.5.1. The *unordered boundary quotient* is the Banach space

$$\widetilde{\mathfrak{X}} := \mathfrak{X}^{p}_{\rho,\tau,\eta}/\mathfrak{N},$$

with quotient map $\mathfrak{q} : \mathfrak{X}^{p}_{\rho,\tau,\eta} \to \widetilde{\mathfrak{X}}$.

For a subset $S \subseteq \{1, \dots, r\}$, a residual multi-index $\mathbf{k} \in \mathbb{N}_0^{\{1,\dots,r\}\setminus S}$, and $\lambda \in \Lambda$, define the *unordered trace class*

$$u_{S,\mathbf{k},\lambda} := \mathfrak{q}(t_{w,\mathbf{k},\lambda}),$$

where $w \in \mathcal{W}(S)$ is any word with support S.

Proposition 15.5.2. *The class $u_{S,\mathbf{k},\lambda}$ is well defined: it does not depend on the choice of the word $w \in \mathcal{W}(S)$.*

Proof. If $w, w' \in \mathcal{W}(S)$, then $t_{w,\mathbf{k},\lambda} - t_{w',\mathbf{k},\lambda} \in \mathfrak{N}$ by definition, so $\mathfrak{q}(t_{w,\mathbf{k},\lambda}) = \mathfrak{q}(t_{w',\mathbf{k},\lambda})$. □

In the quotient, the ordered trace words are replaced by unordered boundary supports. The trace state $t_{(1,2),\mathbf{k},\lambda}$ and the trace state $t_{(2,1),\mathbf{k},\lambda}$ become the same element $u_{\{1,2\},\mathbf{k},\lambda}$. The ordering information—which was the sole source of noncommutativity—has been identified.

Theorem 15.5.3 (Induced commuting quotient calculus)**.** *The operators C_i and all bounded multipliers M_σ descend to bounded operators $\widetilde{C}_i, \widetilde{M}_\sigma \in \mathcal{B}(\widetilde{\mathfrak{X}})$ satisfying*

$$\mathfrak{q} \circ C_i = \widetilde{C}_i \circ \mathfrak{q}, \qquad \mathfrak{q} \circ M_\sigma = \widetilde{M}_\sigma \circ \mathfrak{q}.$$

On the quotient basis classes:

$$\widetilde{C}_i\, u_{S,\mathbf{k},\lambda} = \begin{cases} 0, & i \in S, \\ u_{S,\mathbf{k}-\mathbf{e}_i,\lambda}, & i \notin S,\ k_i \geq 1, \\ u_{S\cup\{i\},\mathbf{k}^{\widehat{i}},\lambda}, & i \notin S,\ k_i = 0. \end{cases}$$

Moreover, the quotient Caputo tuple is commuting:

$$[\widetilde{C}_i, \widetilde{C}_j] = 0 \qquad (1 \leq i, j \leq r).$$

Proof. Since $\mathfrak{N}$ is invariant under every C_i and every M_σ (Proposition 15.4.3), the induced operators on the quotient are well defined and bounded.

The explicit action on quotient classes follows from the basis action of C_i. In the vacuum case: if $i \notin S$ and $k_i = 0$, choose any $w \in \mathcal{W}(S)$. Then $\widetilde{C}_i\, u_{S,\mathbf{k},\lambda} = \mathfrak{q}(t_{wi,\mathbf{k}^{\hat{i}},\lambda})$. Any other word $w' \in \mathcal{W}(S)$ gives the same class because $\operatorname{supp}(wi) = \operatorname{supp}(w'i) = S \cup \{i\}$.

The quotient tuple commutes because the commutator operators have range in $\mathfrak{N}$ (Theorem 15.4.4), so they vanish after applying the quotient map: $\widetilde{C}_i\, \widetilde{C}_j\, \mathfrak{q}(x) = \mathfrak{q}(C_i C_j x) = \mathfrak{q}(C_j C_i x) = \widetilde{C}_j\, \widetilde{C}_i\, \mathfrak{q}(x)$. ◻

The quotient has a universal property: it is the "freest" way to impose commutativity on the ambient calculus.

Theorem 15.5.4 (Universal property of the commuting quotient)**.** *Let Y be a Banach space carrying a commuting tuple $D_1, \dots, D_r \in \mathcal{B}(Y)$ and multipliers $N_\sigma \in \mathcal{B}(Y)$ commuting with every D_i. If*

$$\Phi : \mathfrak{X}^p_{\rho,\tau,\eta} \to Y$$

is a bounded linear map satisfying $\Phi \circ C_i = D_i \circ \Phi$ and $\Phi \circ M_\sigma = N_\sigma \circ \Phi$ for every i and every bounded σ, then $\mathfrak{N} \subseteq \ker \Phi$. Consequently, there exists a unique bounded map $\widetilde{\Phi} : \widetilde{\mathfrak{X}} \to Y$ with $\Phi = \widetilde{\Phi} \circ \mathfrak{q}$.

Proof. The kernel of Φ is a closed subspace invariant under every C_i and every M_σ. For any distinct i, j,

$$\Phi \circ [C_i, C_j] = [D_i, D_j] \circ \Phi = 0,$$

because the target tuple $(D_1, \dots, D_r)$ is commuting. Hence $\ker \Phi$ contains all commutator ranges. By Theorem 15.4.7, $\mathfrak{N}$ is the smallest closed invariant subspace with this property, so $\mathfrak{N} \subseteq \ker \Phi$. The factorization through the quotient is the standard universal property of Banach-space quotients. ◻

15.6 Restriction versus quotient

The book has now produced two canonical routes from the noncommutative ambient calculus to a commuting one.

Theorem 15.6.1 (Two commuting reductions)**.** *The following are two genuinely different commuting reductions of the boundary-augmented ambient calculus.*

(i) Restriction. *The maximal commuting sector*

$$\mathfrak{R}^p_{\rho,\tau,\eta} = X^p_{\rho,\eta} \oplus \bigoplus_{\substack{w \in \mathcal{W}_r \\ q(w) \le 1}} T^p_w$$

is the largest closed graded invariant subspace on which the extended Caputo tuple commutes. It is obtained by discarding the trace blocks whose free rank is ≥ 2.

(ii) Quotient. *The space $\widetilde{\mathfrak{X}} = \mathfrak{X}^p_{\rho,\tau,\eta}/\mathfrak{N}$ is the universal commuting quotient. It retains all boundary supports but identifies orderings with the same support.*

The restriction preserves the ordered trace data and removes the noncommuting strata. The quotient preserves all boundary supports and collapses the ordering information that causes the commutator.

Proof. Part (i) is the maximality theorem of Chapters 9 and 13. Part (ii) is Theorem 15.5.4. ◻

Example 15.6.2 ($r = 2$: the two reductions compared)**.** With $r = 2$, the words and their free ranks are: $\varnothing$ ($q = 2$), (1) ($q = 1$), (2) ($q = 1$), (1, 2) ($q = 0$), (2, 1) ($q = 0$).

Restriction excludes the empty-word trace block $T^p_\varnothing$ (the only block with $q(w) \geq 2$) and retains everything else:

$$\mathfrak{R}^p = X^p \oplus T^p_{(1)} \oplus T^p_{(2)} \oplus T^p_{(1,2)} \oplus T^p_{(2,1)}.$$

Quotient identifies $T^p_{(1,2)}$ and $T^p_{(2,1)}$ (the two blocks with support $\{1, 2\}$) into a single unordered class:

$$u_{\{1,2\},0,\lambda} = \mathfrak{q}(t_{(1,2),0,\lambda}) = \mathfrak{q}(t_{(2,1),0,\lambda}).$$

The quotient retains the empty-word block (which the restriction excluded) and all other blocks, but collapses the corner ordering.

The two reductions are visibly different. The restriction is "smaller" in the sense that it lives inside the ambient space; the quotient is "larger" in the sense that it surjects from the full ambient space. Neither contains the other.

Example 15.6.3 ($r = 1$: both reductions are trivial)**.** When $r = 1$, there are no mixed commutators. The ordering-defect space is $\mathfrak{N} = \{0\}$, and the quotient is the entire ambient space. The maximal commuting sector is also the entire ambient space. Both reductions are trivial, as expected.

Example 15.6.4 ($r = 3$: counting quotient classes)**.** With $r = 3$, there are 16 words in $\mathcal{W}_3$. The words are grouped by support: one support $\varnothing$, three supports of size 1, three of size 2, and one of size 3. The quotient collapses each group to a single unordered class, producing 8 classes (one for each subset of $\{1, 2, 3\}$). In the restriction, the 4 blocks with $q(w) \geq 2$ are removed, leaving 13 out of 17 total blocks. The quotient keeps all 17 blocks but identifies them into 9 classes (one canonical, plus 8 unordered trace classes).

15.7 Further directions

The ordered boundary trace calculus and the commutator ideal theorem bring the AD01–AD12 research program to a natural resting point. The program began with the search for the correct canonical spaces (Chapters 3–7), passed through completion and transform theory (Chapters 8, 12, 14), analyzed the boundary-generated noncommutativity (Chapters 9, 13), and arrived in this chapter at a structural understanding of the commutator as an ideal.

Several natural continuations suggest themselves.

Concrete realizations of trace layers. Throughout the boundary-augmented theory, the trace states $t_{w,\mathbf{k},\lambda}$ are formal symbols. A natural question is whether they can be realized as actual face traces, corner jets, or anisotropic distributions on the underlying continuous or discrete domains. Such a realization would connect the abstract coefficient-space theory to the analytic geometry of mixed-domain boundary value problems.

Cohomological interpretation. The commutator ideal is generated by differences of the form $t_{wji,\ldots} - t_{wij,\ldots}$, which are "boundary cocycles" measuring the failure of path-independence in the ordered trace calculus. It is natural to ask whether this structure admits a systematic cohomological interpretation, perhaps in terms of group cohomology of $\mathfrak{S}_r$ acting on the trace layers.

Representation theory of the quotient algebra. The quotient $\widetilde{\mathfrak{X}}$ carries a commuting tuple and a diagonal spectral block. Understanding its representations— especially the irreducible ones—may shed light on the spectral theory of mixed-domain fractional operators.

Generalized Weyl operators and mixed systems. The broader framework of Kostić for multidimensional fractional calculus suggests that the boundary-trace mechanism may extend to more general classes of mixed differential-difference operators. The coefficient-space viewpoint of this book is well suited to such extensions.

Remark 15.7.1. The central insight of this chapter—and of the boundary-augmented theory as a whole—is that noncommutativity in the partial Caputo calculus is not a deficiency or an obstacle. It is a structured phenomenon with a precise algebraic source (ordered boundary creation), a precise invariant space (the ordering-defect space $\mathfrak{N}$), and two precise resolutions (restriction to the maximal commuting sector, or quotient by the commutator ideal). The noncommutativity is, in a precise sense, organized by the geometry of the boundary.

Chapter 16

Looking Back at the Whole Program

This book began with a simple observation from ordinary calculus: when one chooses the right basis, differentiation and integration become transparent operations—a backward shift on normalized monomials, a diagonal multiplication on exponentials. It then asked whether the same kind of structural clarity could be achieved for operators of fractional order, and followed that question through fifteen chapters of increasingly rich mathematics.

This final chapter steps back from the technical development and surveys the landscape as a whole. It does not introduce new theorems. Instead, it distills the main themes that run through the book, indicates how they connect to the broader mathematical literature, and suggests how a reader might navigate the material on a first or second pass.

16.1 The complementarity of the shift and spectral pictures

The most basic structural dichotomy of the book is the distinction between the *shift picture* and the *spectral picture*. It appeared first in the integer-order setting of Chapter 1 and persisted, in increasingly sophisticated forms, through every subsequent stage.

In the shift picture, the canonical basis consists of Gamma-normalized fractional monomials (continuous) or normalized rising factorials (discrete). The Riemann–Liouville integral (or nabla fractional sum) advances the grade index by one step, and the Caputo derivative (or Caputo nabla difference) lowers it by one step. The grade-zero subspace is the vacuum, annihilated by the backward shift. This picture is natural on one-sided domains—domains bounded by a wall at the origin—because the wall creates a distinguished boundary and a natural grading.

In the spectral picture, the canonical basis consists of exponential characters (continuous) or lattice characters (discrete). Each character is an eigenvector of the Weyl-type fractional operators, and the operator algebra becomes a diagonal multiplier algebra. There is no vacuum, no boundary layer, and no defect. This picture is natural on whole-space domains—domains without a boundary—because the relevant symmetry is translation invariance rather than grading.

The hybrid theory of Chapters 7 and 11 unites the two pictures on mixed domains of the form $(0,\infty)^r \times \mathbb{R}^s$ (or $\mathbb{N}^r \times \mathbb{Z}^s$ in the discrete case). The canonical basis is a tensor product of one-sided monomials and whole-space characters, and the operator algebra has two commuting blocks: a shift block in the one-sided coordinates and a diagonal block in the whole-space coordinates. All defect phenomena are localized in the shift block; the spectral block contributes multiplicity but no new structure.

The complementarity is not a matter of taste or convention. It reflects the geometry of the underlying domain. Wherever there is a boundary, the shift picture governs; wherever the domain is homogeneous, the spectral picture governs; wherever both features coexist, the hybrid algebra is the natural framework.

Remark 16.1.1. The reader familiar with classical analysis will recognize a parallel with the Fourier and Laplace transforms. The Fourier transform diagonalizes constant-coefficient operators on $\mathbb{R}^d$ (the spectral picture), while the Laplace transform converts initial-value problems on $(0, \infty)$ into algebraic equations (the shift picture with boundary data). The hybrid algebra can be viewed as a fractional-order unification of these two classical viewpoints.

16.2 Why completion changes the theme

A recurring motif of the book is the passage from algebra to analysis: from the algebraic core (finite linear combinations of basis vectors) to a completed Banach space (infinite coefficient sequences with controlled growth).

The algebraic core is where the exact operator relations are discovered and proved. The shift identities $C_iJ_i = I$, $J_iC_i = I - \Pi_i$, and the commutativity of the hybrid algebra are all algebraic statements, verified on finite-support vectors by one-line computations. The algebraic theory is clean and self-contained.

But the algebraic core is too small for analysis. The generating eigenvectors (the hybrid Mittag–Leffler vectors of Chapters 8 and 12) are infinite series and do not belong to the algebraic module. There are no operator norms, no convergence questions, no spectral theory, and no semigroups. The passage to a Banach completion is not a technicality; it is the transition from algebraic model-building to genuine operator theory.

Once the completion is in place, new phenomena appear. The operators acquire norms, and the question of which weights make these norms optimal becomes meaningful (Chapter 14). The generating eigenvectors become genuine Banach-space elements, and the transform model becomes an injective map into a space of holomorphic functions. Closed unbounded multipliers and their semigroups enter the picture, and the bounded perturbation theorem produces mixed generators with explicit growth bounds.

The abstract coefficient-space framework of Chapter 14 reveals that the completion theory is the same in the continuous and discrete settings: the two branches are different function-space realizations of a single weighted ℓ^p space over the common grade lattice $\mathbb{N}_0^r$. The geometric weights are singled out as the unique balanced weights (satisfying the sharp product inequality $\|J_i\| \|C_i\| = 1$), and they also optimize the semigroup-growth bounds for mixed generators.

In summary, the algebraic core provides the skeleton; the completion provides the flesh. The skeleton determines what is possible; the completion determines what converges, what is bounded, and what generates dynamics.

16.3 Why boundary matters

The boundary is the most consequential geometric feature in the entire program. Its influence is felt at every level of the theory.

Vacuum and defect. At the algebraic level, the boundary creates the vacuum: the grade-zero subspace annihilated by the backward shift. The identity $J_iC_i = I - \Pi_i$ says that the backward-then-forward composition recovers everything except the vacuum component. This defect is the one-variable analogue of initial data in an ordinary differential equation; in the fractional setting, it is built into the algebraic structure of the shift model from the very beginning (Chapters 4–5).

Boundary layers. In several one-sided variables, the single vacuum becomes a family of boundary hyperplanes $\{k_i = 0\}$ in the grade lattice $\mathbb{N}_0^r$. The intersection pattern of these hyperplanes produces a hierarchy of boundary strata: faces, edges, corners. The defect structure of the multi-variable shift algebra reflects this lattice geometry (Chapter 5).

Ordered traces and noncommutativity. The most striking consequence of the boundary appears in the boundary-augmented theory (Chapters 9, 13, 15). When one enlarges the canonical completion by adjoining trace states that record the order of boundary extractions, the extended Caputo tuple is no longer commuting. The commutator is nonzero precisely when two free one-sided coordinates are simultaneously at vacuum: the two boundary extractions compete for priority, and the two possible orderings produce genuinely different trace states.

The commutator formula

$$[C_i, C_j]\, t_{w,\mathbf{0},\lambda} = t_{wji,\varnothing,\lambda} - t_{wij,\varnothing,\lambda}$$

shows that the obstruction is the difference of two ordered histories with the same boundary support. The ordering-defect space $\mathfrak{N}$—the closed span of all such differences—is the range of the commutator ideal, and quotienting by $\mathfrak{N}$ collapses ordered traces to unordered supports, producing the universal commuting quotient (Chapter 15).

Two resolutions. The book offers two canonical ways of recovering commutativity from the noncommutative ambient space. The first is *restriction*: discard the trace blocks whose free rank is ≥ 2, retaining the maximal graded invariant commuting sector. The second is *quotient*: identify trace states that differ only in their ordering, producing the universal commuting quotient. Restriction preserves the ordered structure on a smaller domain; the quotient preserves the full domain but coarsens the space. Neither subsumes the other, and together they give a complete structural picture of boundary-generated noncommutativity.

The spectral block is innocent. Throughout the boundary-augmented theory, the spectral multipliers remain diagonal on every block—canonical and trace alike. The commutator formula never involves the spectral label λ. The whole-space directions contribute multiplicity but create no obstruction. Noncommutativity is purely a one-sided boundary phenomenon.

16.4 Broader literature and next reading

The results of this book make contact with several established areas of mathematics. We indicate the main connections here, together with suggestions for further reading.

Fractional calculus. The standard references for the Riemann–Liouville and Caputo operators are Podlubny [9] and Kilbas–Srivastava–Trujillo. The Mittag–Leffler function and its role in fractional equations are treated in Mainardi [7]. The multidimensional fractional calculus framework of Kostić provides a broader context for the Weyl-type operators and mixed-domain settings used in this book.

Discrete fractional calculus. The discrete analogues of fractional sums and differences, rising factorials, and the nabla calculus are developed systematically in Goodrich–Peterson [5]. The discrete chapters of this book (Chapters 10–13) follow the notation and conventions of that reference, adapted to the canonical-basis viewpoint.

Operator theory and model theory. The shift-model philosophy of this book is rooted in the classical model theory of contractions (Sz.-Nagy–Foiaş) and the theory of weighted shift operators (Shields). The notion of a canonical model space on which an operator takes a simple form is a central theme of operator theory, and the constructions of this book can be viewed as fractional-order analogues of that program. Standard references include Conway [2] and Kreyszig for general functional analysis, and Engel–Nagel [4] or Pazy [8] for C_0-semigroups.

Complex analysis. The minimum complex analysis needed for the spectral picture (principal branches, complex powers, half-plane geometry) is covered in any standard text such as Conway's *Functions of One Complex Variable*. The polydisk transform models of Chapters 8, 12, and 14 use standard Weierstrass convergence and power-series theory in several complex variables; Krantz provides a comprehensive reference.

Algebra and ideal theory. The commutator ideal and the universal commuting quotient of Chapter 15 are operator-algebra analogues of classical constructions in ring theory. The reader who has studied commutative algebra (e.g., Atiyah–Macdonald) will recognize the quotient R/I construction; the Banach-space version requires additional care with closures and invariance, but the algebraic intuition is the same.

Matrix functions. The nonexistence theorem of Chapter 3 uses the theory of matrix roots and nilpotent Jordan blocks. A thorough treatment of matrix functions, including the question of when a matrix has a p-th root, can be found in Higham [6].

16.5 Reading paths through the book

The book is designed to be read linearly, but not every reader will want to read every chapter with the same intensity on a first pass. Here are two suggested paths.

A first pass: the core story. A reader encountering this material for the first time should focus on the following chapters:

Chapter 1	The global map and the central question.
Chapter 2	Fractional integrals, Caputo derivatives, and first computations.
Chapter 3	The impossibility on P_n: why the naive approach fails.
Chapter 4	The one-variable shift model: the heart of the program.
Chapter 7	The hybrid algebra: how shift and spectral pictures combine.
Chapter 8	Banach completions: the passage from algebra to analysis.
Chapter 9	Boundary augmentation: where noncommutativity comes from.

These seven chapters tell the essential story: impossibility, the canonical basis, hybridization, completion, and boundary structure. Each subsequent chapter adds depth—multi-variable extensions, discrete analogues, the abstract coefficient framework, the commutator ideal—but the conceptual spine is visible in the seven chapters listed above.

A second pass: the full theory. On a second reading, the reader should fill in the remaining chapters:

Chapter 5	Multi-variable one-sided calculus: from a chain to a lattice.
Chapter 6	Whole-space spectral models: exponentials and Weyl operators.
Chapter 10	Discrete preparation: the continuous/discrete dictionary.
Chapters 11–13	The discrete trilogy: hybrid algebra, completion, boundary.
Chapter 14	The unified coefficient-space theory and semigroups.
Chapter 15	The commutator ideal and the universal commuting quotient.

The second pass reveals the full scope of the program: the multi-variable and discrete theories, the abstract unification, and the culminating algebraic structure of ordered traces and commutator ideals.

A thematic reading. A reader interested in a particular theme may also read the book selectively:

- *Canonical bases and shift models*: Chapters 1, 2, 3, 4, 5.
- *Spectral models and complex analysis*: Chapters 6, 7 (the spectral block), 8 (the transform model).
- *The continuous/discrete parallel*: Chapters 10, 11, 12, 13, compared with Chapters 7, 8, 9.
- *Completion and semigroup theory*: Chapters 8, 12, 14.
- *Boundary, noncommutativity, and ideals*: Chapters 9, 13, 15.

Remark 16.5.1. The book does not contain formal exercise sets, but the reader is strongly encouraged to work through the low-dimensional examples ($r = 1$, $r = 2$) that appear throughout. The two-variable case $r = 2$ is the smallest setting in which every phenomenon of the theory—vacuum layers, commutator formulas, maximal commuting sectors, ordering-defect spaces, the quotient reduction—is nontrivially visible. Reproducing the $r = 2$ examples by hand is the single most effective way to internalize the ideas of the book.

Final perspective

The twelve research papers AD01–AD12 form a single arc. The program began with a negative result (the impossibility of internal fractional models on finite-dimensional polynomial spaces) and arrived at a positive structural theory in which fractional operators on mixed domains are understood through canonical bases, exact operator algebras, weighted completions, transform models, semigroup generation, and boundary trace calculus.

The central message can be stated simply. Fractional calculus is not, at its core, a collection of complicated integral formulas. It is an *operator algebra* that becomes transparent when viewed on the right spaces and in the right bases. Where there is a boundary, the algebra is a shift algebra. Where there is no boundary, it is a spectral algebra. Where both features coexist, it is a hybrid of the two. The passage to Banach completions turns the algebra into genuine operator theory. And the ordered boundary trace calculus reveals that even the noncommutativity of the multi-variable theory is not a pathology but a structured phenomenon, governed by the geometry of the boundary and resolvable by classical algebraic operations—restriction and quotient.

This perspective transforms fractional calculus from a technical specialty into a chapter of operator algebra and Banach-space geometry. The reader who has followed the book to this point is equipped not only with the specific results of the twelve papers, but with a way of thinking about fractional operators—as inhabitants of canonical spaces, as components of exact algebraic models, and as generators of structured dynamics—that extends naturally to problems beyond the scope of the present work.

Bibliography

[1] F. M. Atici and P. W. Eloe, *Initial value problems in discrete fractional calculus*, Proc. Amer. Math. Soc. **137** (2009), no. 3, 981–989.

[2] J. B. Conway, *A Course in Functional Analysis*, 2nd ed., Graduate Texts in Mathematics, Vol. 96, Springer, New York, 1990.

[3] K. Diethelm, *The Analysis of Fractional Differential Equations: An Application-Oriented Exposition Using Differential Operators of Caputo Type*, Lecture Notes in Mathematics, Vol. 2004, Springer, Berlin, 2010.

[4] K.-J. Engel and R. Nagel, *One-Parameter Semigroups for Linear Evolution Equations*, Graduate Texts in Mathematics, Vol. 194, Springer, New York, 2000.

[5] C. Goodrich and A. C. Peterson, *Discrete Fractional Calculus*, Springer, Cham, 2015.

[6] N. J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, 2008.

[7] F. Mainardi, *Fractional Calculus and Waves in Linear Viscoelasticity*, Imperial College Press, London, 2010.

[8] A. Pazy, *Semigroups of Linear Operators and Applications to Partial Differential Equations*, Applied Mathematical Sciences, Vol. 44, Springer, New York, 1983.

[9] I. Podlubny, *Fractional Differential Equations*, Mathematics in Science and Engineering, Vol. 198, Academic Press, San Diego, 1999.

Index