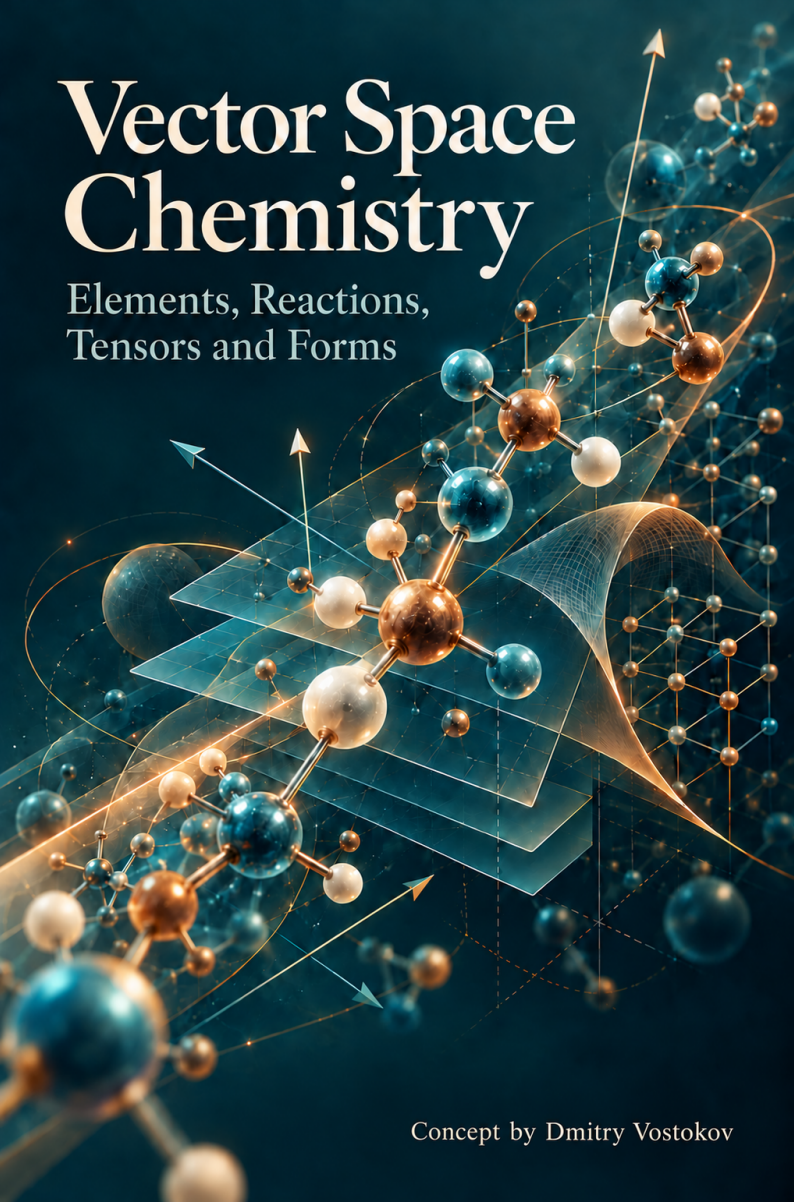
Vector Space Chemistry

Elements, Reactions, Tensors and Forms



Concept by Dmitry Vostokov

Vector Space Chemistry

Elements, Reactions, Tensors and Forms

Concept and editorial direction
Dmitry Vostokov

Version 14 • November 2026

Vector Space Chemistry

Elements, Reactions, Tensors and Forms

Copyright © 2026 Dmitry Vostokov. All rights reserved.

Copyright © 2026 OpenTask. All rights reserved.

ISBN-13: 978-1-906717-55-1

No part of this work may be reproduced, stored, or transmitted in any form or by any means without prior written permission from the copyright holders, except for brief quotations used in review, scholarship, or teaching as permitted by law.

Manuscript Version 14, November 2026.

Developed from the author's vector space chemistry idea of 1986–1987.

Drafting and revision assistance: GPT-6 Astra
Independent editorial review and additional suggestions: Claude Fable 5.1

Contents

Preface **8**

Representations and Conventions **12**

Part I Elements and Chemical Representations **14**

Chapter 1 The Original Idea **15**

1.1 Elements as directions 15

1.2 Multiplication and addition 16

1.3 Balancing as equality 17

1.4 What the first representation omits 18

Chapter 2 Vectors Scalars and Physical Inventories **19**

2.1 The free vector space 19

2.2 Four domains of coefficients 20

2.3 Order and feasibility 21

2.4 Independence is chosen at this level 21

Chapter 3 Formulas Multisets and Molecular Identity **23**

3.1 Multiplicity before geometry 23

3.2 Formula units and empirical formulas 24

3.3 Structure as additional data 25

3.4 Isotopes and coarse representations 25

Chapter 4 Species Space and the Composition Map **27**

4.1 Two bases for two questions 27

4.2 Kernel and image 28

4.3 Fibers and indistinguishability 29

4.4 The map does not supply reactions 30

Chapter 5 Balancing as a Nullspace Problem 31
5.1 Signed and unsigned formulations 31
5.2 A complete calculation 32
5.3 Rational solutions and integer normalization 32
5.4 Underdetermined equations 34
5.5 Integer kernels and images 35

Part II Reactions and Linear Algebra 37

Chapter 6 Reaction Spaces and Conservation Laws 38
6.1 The stoichiometric map 38
6.2 A network with one cycle 39
6.3 Left nullspaces 39
6.4 Three kernels with different meanings 40

Chapter 7 Amounts Extents and Feasible Regions 43
7.1 States move in affine spaces 43
7.2 Limiting reagents as interval endpoints 43
7.3 Normalization changes the extent coordinate 44
7.4 Several extents and nonuniqueness 45
7.5 Discrete counts and Petri nets 46

Chapter 8 Dual Spaces Charge and Mass 48
8.1 Observables as covectors 48
8.2 Pulling a measurement back to species 49
8.3 Charge is additional information 50
8.4 Conservation and electroneutrality 52

Chapter 9 Coordinates Projections and Lost Information 53
9.1 A new basis is not new chemistry 53
9.2 What counts as an inner product 54
9.3 Projection as an observational choice 54
9.4 Reduced models need compatible dynamics 55

Chapter 10 From Stoichiometry to Dynamics 57
10.1 The rate map 57
10.2 A reversible conversion 58
10.3 Linearization of a nonlinear system 59
10.4 Counts and stochastic events 59

Part III Tensor Algebra and Forms 61

Chapter 11 Bilinear and Quadratic Forms 62
11.1 Two linear arguments 62
11.2 Quadratic measurements 63
11.3 Local energy curvature 64
11.4 Distances and uncertainty weights 66

Chapter 12 Tensor Products and Chemical Context 68
12.1 Two kinds of labels 68
12.2 Pure tensors and correlations 69
12.3 Composition and aggregation by contraction 70
12.4 Transport and local reactions 70
12.5 Tensor language and physical tensors 71

Chapter 13 Polynomial Kinetics as Tensor Contraction 73
13.1 Multilinear maps and diagonal evaluation 73
13.2 A bimolecular example 74
13.3 Repeated reactants and polynomial degree 75
13.4 Conservation as a tensor constraint 75
13.5 Tensor powers do not close dynamics automatically 76

Chapter 14 Exterior Algebra and Alternating Forms 78
14.1 Alternation as a mathematical choice 78
14.2 A two-dimensional calculation 79
14.3 Reaction directions and rank 79

14.4 A special three-coordinate shortcut 80
14.5 Forms versus vectors 81

Chapter 15 Differential Forms and Thermodynamic Change 82
15.1 Forms at a state 82
15.2 Exact forms and path dependence 83
15.3 Chemical potentials as a one-form 84
15.4 A reaction-level example 85
15.5 Dissipation and its assumptions 87

Chapter 16 Measurements Inverse Problems and Uncertainty 88
16.1 An observation model 88
16.2 A worked reconstruction 89
16.3 Constraints and residuals 89
16.4 Covariance transforms with the map 91
16.5 Structural and dynamical inference 92

Part IV Applications Limits and Development 94

Chapter 17 Integrated Chemical Examples 95
17.1 Neutralization and the charge row 95
17.2 A redox balance 96
17.3 A catalytic mechanism 97
17.4 Reaction energies and linear dependence 98

Chapter 18 Reaction Graphs and Composed Models 100
18.1 Complexes are multisets of species 100
18.2 Factorizing the catalytic example 101
18.3 A structural count and its limits 101
18.4 Opening the system 102
18.5 Composing compartments 103

Chapter 19 Model Meaning and Validation 105
19.1 Labels and interpretation 105

19.2 Exact arithmetic 106
19.3 More than one solution 107
19.4 Validation by independent checks 108
19.5 The organization of a model 109

Chapter 20 Scope Originality and Further Directions 110
20.1 What has been developed 110
20.2 Exact correspondences and useful analogies 111
20.3 Enriching composition systematically 112
20.4 Constrained learning and observation design 114
20.5 Further investigation 115

Appendix A Notation and Terminology 116
A.1 Spaces and structural maps 116
A.2 States rates and measurements 118
A.3 Algebraic operations 119
A.4 Units and coefficient domains 120
A.5 A compact glossary 121

Appendix B Mathematical Background 126
B.1 Linear maps before matrices 126
B.2 Row reduction with an explicit kernel 127
B.3 Why rank and nullity add up 128
B.4 Fibers quotients and feasible slices 129
B.5 Dual maps and change of basis 130
B.6 Tensor products in coordinates 130
B.7 Differentiation and pullback 131
B.8 A practical order of calculation 132

Bibliography 134
Index 139

Preface

In 1986–1987, when I was in the ninth and tenth classes at school, I developed what I called vector space chemistry. I considered chemical elements as basis vectors and stoichiometric coefficients as scalars. We can then view a chemical expression as a linear combination and apply vector operations to it. In this book, I return to that idea and follow its connections with linear algebra, tensor algebra, and forms.

I described the idea publicly in my blog post *Vector Space Chemistry* on 23 January 2009 (Vostokov, 2009). It announced an OpenTask book, with ISBN 978-1-906717-55-1, and presented mathematics as a tool for chemical modeling in a popular science account. This is that book, developed further and retaining the announced ISBN.

The post provides a dated public account. Its reference to a preface written “after 25 years” differs from the 1986–1987 dates I recall here. The precise school-year chronology remains a matter of personal recollection.

The tensor constructions, differential forms, and extended chemical examples presented here are the current development of the starting idea. Linear methods in chemistry and reaction-network theory have an established history, including work before

1986–1987. I use that mathematics throughout the book; my personal starting point should not be taken as a claim of priority for those methods.

Let us first consider an ordinary chemical formula. We can read it as a string of symbols, or we can view each elemental count as a coordinate. Water and hydrogen peroxide then differ in their oxygen coordinate. Multiplying the water vector by two gives the composition of two formula units. If we apply the same counting to both sides of a reaction equation, balancing becomes equality of the resulting vectors. This simple correspondence leads us to linear maps and their kernels.

Sometimes different molecules give us the same vector. Ethanol and dimethyl ether have the same elemental counts, although their atoms are connected differently. We may also obtain the same total composition from different mixtures. The representation retains an inventory and loses the organization of that inventory. To preserve species identity, we introduce another space, with species as basis vectors, and connect it to the elemental space. We use this construction repeatedly in the following chapters.

The book retains the compact format and reference structure of my *Applied Set Theory*. There is also a direct mathematical connection: a chemical formula can be considered as a multiset of elemental labels. We can pass from such multisets to signed

combinations and then to real vector spaces. At each step, we examine which operations become available and what their results mean. Numbered definitions and worked examples make these connections explicit.

We begin with elements, formulas, and species in Part I. In Part II, we consider reactions as changes in species space and examine conservation, feasible amounts, observations, and dynamics. Part III develops the tensor and form constructions, including bilinear and quadratic forms, exterior algebra, and thermodynamic differentials. Part IV brings these approaches together in chemical examples and questions of model interpretation, and then considers further directions.

The examples use ordinary chemical formulas and specified model systems. A balanced expression tells us that elemental totals agree. To discuss whether a reaction occurs, we also need its chemical conditions. Similarly, a rate constant, an ideal-mixture assumption, or a measurement operator has to come from the model being studied. We keep these additions visible so that an algebraic result can be interpreted at the appropriate level.

We develop the argument through explicit representations and their consequences. Appendix A collects the notation and terminology, and Appendix B brings together the mathematical constructions used throughout the book. Integer and rational

examples make the relationships visible without introducing numerical approximation.

The school-period idea gives us a starting representation. We can now change that representation, compare it with others, and ask which chemical distinctions remain observable. The following chapters develop these questions through examples.

Representations and Conventions

Spaces and coordinates. We use E for elemental composition, W for species, and R for reaction coordinates. The maps between them are $A: W \to E$ and $S: R \to W$. Once a basis has been stated, we use the same letter for a map and its matrix. Vectors are columns. A superscript T denotes transpose, and an asterisk denotes a dual space.

Chemical notation. We distinguish elemental symbols such as H and O from the vectors they label. For example, e_{H} is an elemental basis vector and $u_{\mathrm{H_2O}}$ is a species basis vector. This keeps hydrogen atoms, hydrogen gas, and coordinate directions separate. Subscripts count atoms within a formula; coefficients before formulas specify numbers of formula units or reaction proportions. An amount of substance, usually measured in moles, is a further quantity.

Scalars and units. Before interpreting a coefficient, we determine what it measures. It may be an integer count, a rational balancing coefficient, a real amount, or a signed change. Negative coordinates are useful for changes, although they cannot describe a physical inventory. We attach units to chemical amounts and state when an example is dimensionless. For the conventions on physical quantities and symbols, see IUPAC (2023).

Kinds of assertion. We introduce mathematical objects in definitions and justify propositions by the arguments that follow them. Worked examples show how to use the constructions. When we add a physical model, we state its assumptions. When we use an analogy, we identify the correspondence and its limits. Numbering restarts for each chapter and each type of item, so Definition 4.1 and Example 4.1 refer to different passages.

Connections between chapters. Chapters 1–8 establish the representations used in the chemical examples of Chapter 17. Chapters 9–15 develop their algebraic and geometric extensions. Measurements and model interpretation connect Chapters 16 and 19, while Chapter 18 relates the constructions to reaction graphs. Chapter 20 considers their scope and further development. The index and Appendix A provide references to particular concepts.

Sources and examples. The bibliography points to the mathematical and chemical foundations used here. We develop the examples for the exposition; they should be read as calculations in stated models, rather than reports of experiments.

Part I Elements and Chemical Representations

A formula already contains a small inventory. We begin by writing that inventory as a vector and examining the operations we can perform on it. We then compare formulas with the species they describe. This comparison shows why one representation is sufficient for balancing and another is needed to preserve molecular identity.

Chapter 1 The Original Idea

1.1 Elements as directions

Consider formulas containing hydrogen, carbon, and oxygen. We can associate an independent formal direction with each of these elements. Here a direction is identified by an elemental label; it does not point towards a place in physical space. For our examples, let us choose the following order:

$$\mathcal{B}_E = (e_{\mathrm{H}}, e_{\mathrm{C}}, e_{\mathrm{O}}).$$

Definition 1.1. The elemental composition vector of a formula is the sum of its elemental basis vectors, each multiplied by the number of atoms of that element in one formula unit.

For example, the two hydrogen atoms and one oxygen atom in water give

$$a(\mathrm{H_2O}) = 2e_{\mathrm{H}} + e_{\mathrm{O}}.$$

The coordinate column is $(2,0,1)^T$. For carbon dioxide, we obtain $(0,1,2)^T$ in the same basis. Notice the zero coordinates: they record the absence of an element from a particular formula while retaining its place in our vocabulary. We can therefore compare different formulas without changing their coordinate positions.

We now have the basic correspondence of vector space chemistry. Elemental labels give us directions,

and their counts give us coefficients. At this stage the coordinates are nonnegative integers. Later, we allow a larger range of coefficients so that the same representation can also describe differences and macroscopic amounts.

1.2 Multiplication and addition

Example 1.1. Let us take three formula units of water. Multiplying its composition vector by three gives

$$3a(\mathrm{H_2O}) = 6e_{\mathrm{H}} + 3e_{\mathrm{O}}.$$

We have scaled the inventory. The expression $3\mathrm{H_2O}$ still refers to three formula units of water. It does not denote a single molecule containing six hydrogen atoms and three oxygen atoms. The distinction between a coefficient and a subscript in chemical notation must therefore survive our translation into algebra.

We can also add inventories. For example, one formula unit of carbon dioxide and two formula units of water give

$$a(\mathrm{CO_2}) + 2a(\mathrm{H_2O})$$
$$= 4e_{\mathrm{H}} + e_{\mathrm{C}} + 4e_{\mathrm{O}}.$$

From the sum we can recover the total count of each element. We cannot recover whether the original objects were in the same container, in different containers, or arranged in a particular structure.

Addition combines the selected counts. To interpret it physically, we have to retain the meaning of the objects that were added.

1.3 Balancing as equality

Example 1.2. Consider the familiar reaction equation

$$2\mathrm{H}_2 + \mathrm{O}_2 \rightarrow 2\mathrm{H}_2\mathrm{O}$$

and count the elements on each side. We obtain $4e_{\mathrm{H}} + 2e_{\mathrm{O}}$ on the left and the same vector on the right. Their difference is zero. We can therefore recognize elemental balance by comparing two vectors.

The reaction arrow adds another interpretation: it specifies an intended direction of change. The equality of the two inventories says nothing about the mechanism or the conditions needed for that change. A sample may react rapidly, slowly, or imperceptibly while the displayed equation remains balanced. We will keep this distinction when we introduce reaction spaces.

This is the usual counting procedure from introductory chemistry; see Flowers and colleagues (2019). By expressing it through vectors, we can connect it with linear maps. That connection also helps us separate balancing an equation from choosing which reactions belong to a model.

1.4 What the first representation omits

Now consider ethanol and dimethyl ether. Both have the formula C_2H_6O, so our construction gives the same vector for both. Their atom connectivity differs. We have therefore represented one shared attribute of the molecules, while leaving out the information that distinguishes them.

The usefulness of the representation depends on the question. For elemental balance, the omitted connectivity does not affect the calculation. For a structural question, we need to restore it. We can do this by introducing species labels or molecular graphs. Tensor indices will later allow us to retain further distinctions, such as the compartment in which a species occurs.

This also indicates how I develop the 1986–1987 idea here. We begin with its original correspondence and add mathematical structure when an example requires a distinction that the earlier representation cannot make. The later constructions belong to this development.

Chapter 2 Vectors Scalars and Physical Inventories

2.1 The free vector space

For a particular model, we usually need only some elements. Let $\mathcal{E}$ denote this finite selection. The basis can therefore be chosen from the vocabulary of the problem, without including the entire periodic table.

Definition 2.1. The real elemental space E consists of all real linear combinations

$$v = \sum_{X \in \mathcal{E}} v_X \, e_X .$$

We take the vectors e_X as a distinguished basis. Calling the space *free* means that we impose no linear relation between these elemental directions. If their linear combination is zero, all its coefficients must be zero. Each elemental count consequently has its own unambiguous coordinate.

Once we use a real vector space, subtraction and multiplication by any real number become available. Every vector has an additive inverse, and there is a unique zero vector. We use these mathematical properties to calculate with representations. Their chemical interpretation still has to be checked. Axler (2024) and Strang (2010) give the general linear-algebra background.

2.2 Four domains of coefficients

A formula, a difference of formulas, and an amount in moles need different kinds of coefficients. We can distinguish the corresponding domains explicitly.

Definition 2.2. The counting monoid is $\mathbb{N}^{\mathcal{E}}$, where $\mathbb{N}$ includes zero. The integer lattice is $\mathbb{Z}^{\mathcal{E}}$. The nonnegative cone is $\mathbb{R}^{\mathcal{E}}_{\geq 0}$, and the ambient vector space is $\mathbb{R}^{\mathcal{E}}$.

We can add counts and include a zero count within the monoid. Subtraction may take us outside it, so we use the integer lattice for signed differences. For continuously varying amounts, we use the nonnegative cone. The ambient vector space then contains all these states together with formal combinations that are useful as changes but cannot be physical inventories.

For example, one-half mole of water is an ordinary amount of substance, whereas one-half of a water molecule is not an intact molecule. Integer coordinates belong to molecular counting. Rational coefficients help us balance equations and can be rescaled by clearing a common denominator. Real coordinates allow the continuum description of macroscopic amounts. The choice of coefficients follows the level at which we are working.

2.3 Order and feasibility

Suppose we want to compare two inventories. We write $v \leq w$ when the comparison holds in every coordinate. This means that each elemental count or amount in v is no larger than its counterpart in w.

Proposition 2.1. A proposed change d can be applied to an inventory v without making any elemental coordinate negative exactly when $v + d \geq 0$.

We check this by examining the coordinates after the change. Passing the test establishes that none is negative. It leaves further chemical questions open: the species needed for the change may be absent, or the transformation may be unavailable under the stated conditions.

Example 2.1. Take the inventory $v = (5,2)^T$ and the proposed change $d = (-3,1)^T$. We obtain $(2,3)^T$, with both coordinates nonnegative. If we instead apply $(-6,1)^T$, the first coordinate becomes -1. The negative entry in a change records removal; a negative entry in the resulting inventory shows that the proposed removal exceeds what is available.

2.4 Independence is chosen at this level

Sometimes we may be tempted to add water as another elemental direction. Suppose we introduce w and require $w = 2e_{\mathrm{H}} + e_{\mathrm{O}}$. We have then made the enlarged family dependent. Hydrogen and oxygen

remain independent labels, but water has been defined as their combination and cannot be a third basis direction at this level.

We can resolve this by changing the objects that label the basis. In a space of species, water and hydrogen gas can have independent directions. Their images in the elemental space may still be dependent. Chapter 4 introduces the map between these two descriptions.

We also choose the scalar field according to the calculation. Real numbers support limits, rates, and measurement uncertainty. Complex numbers are useful in quantum mechanics and spectral methods, but a complex coefficient does not by itself describe an atom count. Extending the algebra requires us to explain the chemical meaning of any new outputs.

Chapter 3 Formulas Multisets and Molecular Identity

3.1 Multiplicity before geometry

A formula tells us both which elements occur and how often they occur. This suggests the multiset representation already used in *Applied Set Theory*.

Definition 3.1. An elemental multiset is a function from the selected elemental labels to the nonnegative integers. Its value at an element is that element's multiplicity.

For water, the values are two for hydrogen, one for oxygen, and zero for every other selected element. Listing these values as coordinates gives our composition vector. Combining two such multisets adds their multiplicities, just as vector addition adds coordinates. We can therefore carry the multiset discussion of Vostokov (2026) directly into formula arithmetic.

If we keep only a set of elemental names, water and hydrogen peroxide become indistinguishable: both contain hydrogen and oxygen. This set is called the support, and the map to it discards multiplicity. Our composition representation retains the counts but still omits connectivity. Comparing the two maps tells us exactly which information has been lost at each step.

3.2 Formula units and empirical formulas

Before counting formula units, we should identify what the formula describes. A molecular formula gives counts within a molecule. In an ionic solid, a formula unit commonly gives a proportion without asserting that isolated molecules of that formula exist. We can use the same vector operations in both cases, provided we retain this difference in interpretation.

Definition 3.2. For a nonzero integer composition vector, its primitive composition vector is obtained by dividing all entries by their greatest common divisor.

Example 3.1. For $C_6H_{12}O_6$, take the coordinate order carbon, hydrogen, oxygen. The vector is $(6,12,6)^T$. Dividing by the greatest common divisor gives $(1,2,1)^T$, corresponding to CH_2O. We have recovered the empirical proportion. The calculation alone cannot recover a molecular formula or identify the substance.

We can view this reduction geometrically. Positive multiples of a nonzero vector form a ray, so choosing the ray preserves proportions and discards scale. Reducing a molecular formula to its empirical formula performs the corresponding integer reduction. This representation is useful when proportions are the subject of the calculation.

3.3 Structure as additional data

Suppose we now need to distinguish different arrangements of the same atoms. We can introduce a molecular graph, with elemental labels on its atom vertices and specified bonds as edges. Summing the vertex labels recovers the composition vector. In this way, the earlier representation becomes an image of the richer one.

Example 3.2. Along the heavy-atom chain of ethanol we have carbon–carbon–oxygen. For dimethyl ether we have carbon–oxygen–carbon. Both graphs give $(2,6,1)^T$ in the carbon, hydrogen, oxygen basis. A calculation receiving only that vector has identical inputs for the two cases and therefore cannot distinguish them.

A graph may still leave out something we need. For example, a problem may depend on stereochemistry, conformation, electronic state, phase, or environment. We add the relevant information when that distinction matters. The purpose of an enrichment is to answer a question that the simpler representation leaves unresolved.

3.4 Isotopes and coarse representations

Isotopes provide another example. An elemental label groups atoms by nuclear charge, but an isotope-sensitive calculation needs finer labels. We can replace the carbon direction by separate carbon-12 and carbon-13 directions. Sending both of them

to the ordinary carbon direction defines a linear aggregation map.

After aggregation, the elemental count remains exact, while the isotope distinction has disappeared. We cannot reverse this loss using the elemental count alone. We call such a passage to a coarser representation *coarse graining*. Here the information removed by the map can be identified directly.

We can treat phase in a similar way. Liquid water and gaseous water have the same elemental image. When phase matters, we give them different species labels and let the composition map merge their elemental contributions. The elemental basis stays the same while the species vocabulary becomes more detailed.

Chapter 4 Species Space and the Composition Map

4.1 Two bases for two questions

The ethanol and dimethyl ether example suggests keeping species names as independent labels. Let us select a finite list $X_1, \dots, X_n$. We can include phase or isotope information in a label whenever it is needed for the problem.

Definition 4.1. The species space W is the real vector space with basis $u_1, \dots, u_n$, one basis vector for each selected species. The composition map $A: W \to E$ is the unique linear map satisfying

$$Au_j = a(X_j).$$

To construct the matrix, we place each species composition in its corresponding column. Specifying these images determines one linear map, because every species-space vector is a linear combination of the basis vectors. Molecular geometry and reaction mechanisms are not needed for this particular mapping.

Example 4.1. Let us return to hydrogen gas, oxygen gas, and water, in that order. With hydrogen and oxygen as the elemental row labels, we obtain

$$A = \begin{pmatrix} 2 & 0 & 2 \\ 0 & 2 & 1 \end{pmatrix}.$$

We can now compare two inventories. The vector $(2,1,0)^T$ contains two hydrogen molecules and one oxygen molecule. The vector $(0,0,2)^T$ contains two water molecules. Applying the matrix to either gives $(4,2)^T$. The species descriptions differ, although the elemental observation is the same.

4.2 Kernel and image

Some differences between species inventories disappear when we apply the composition map. We use the kernel to collect them.

Definition 4.2. The kernel of A is the set of species-space vectors sent to zero. Its image is the set of elemental vectors that arise as images of species-space vectors.

For example, subtract the first inventory in Example 4.1 from the second:

$$v = (-2,-1,2)^T, \qquad Av = 0.$$

The vector v records that difference. Under our chosen orientation, its negative coordinates describe consumption. We should expect nonzero kernel vectors whenever a species change leaves elemental totals unchanged. The vector represents the change, rather than a physical inventory on its own.

We also need to distinguish the full image from images of admissible inventories. As a linear span, the image includes signed combinations. Nonnegative species amounts give only the cone

generated by the columns of A. For example, a model with water as its only species cannot supply hydrogen without oxygen through a nonnegative amount of water.

4.3 Fibers and indistinguishability

Suppose an elemental observation is given and we want all species descriptions compatible with it. This leads to the fiber of the map.

Definition 4.3. The fiber over an elemental vector b is the set of all w satisfying $Aw = b$.

Proposition 4.1. If $Aw_0 = b$, the full fiber is $w_0 + \ker A$.

To see this, subtract the known solution. The equation $Aw = b$ is equivalent to $A(w - w_0) = 0$. We therefore obtain every solution by adding a kernel vector, and each such addition gives a solution. For physical inventories, we retain only the points with nonnegative species coordinates and any other required constraints.

We can now describe observational ambiguity precisely. Species vectors have the same elemental image exactly when their difference belongs to the kernel. The quotient $W/\ker A$ collects these indistinguishable vectors into classes and is isomorphic to the image of A. This identifies the distinctions available to elemental observation; the

original species labels remain necessary for questions about identity.

4.4 The map does not supply reactions

It is tempting to read every vector in $\ker A$ as a reaction. What we have established is elemental balance. A chemically specified model selects transformations from this space and adds direction, kinetics, and thermodynamic conditions. A balanced difference need not be an elementary step or even a plausible net reaction under the conditions being considered.

The word *possible* can therefore refer to different tests. A vector may pass the linear constraints and still fail a chemical requirement concerning energy, a pathway, a catalyst, phase, or time scale. When reporting a result, we should state which of these tests has actually been applied.

Chapter 6 introduces S, whose columns are selected reaction vectors. The condition $AS = 0$ connects these admitted changes with their elemental conservation.

Chapter 5 Balancing as a Nullspace Problem

5.1 Signed and unsigned formulations

Suppose we have already placed each species on the reactant or product side. We can record this choice in a diagonal matrix D, using -1 for reactants and $+1$ for products. Let c contain the nonnegative coefficients displayed in the equation. The corresponding signed reaction vector is then $\nu = Dc$.

Definition 5.1. The signed balancing matrix is $B = AD$. A balanced coefficient vector satisfies

$$Bc = 0, \qquad c \geq 0.$$

If all displayed species must participate, every coordinate of c must be positive. A zero coefficient removes that species from the equation. The all-zero vector satisfies the algebraic equation, but it describes no reaction. We therefore need these additional requirements when choosing a solution.

Keeping A and B separate helps us avoid counting the signs twice. The composition matrix contains nonnegative atom counts; the balancing matrix also records which side each species occupies. We can solve $A\nu = 0$ with sign restrictions on ν, or solve $Bc = 0$ with nonnegativity imposed on c.

5.2 A complete calculation

Example 5.1. Let us balance ethane combustion:

$$\mathrm{C_2H_6 + O_2 \rightarrow CO_2 + H_2O}.$$

Assign coefficients a, b, c, d in the displayed order. Counting carbon, hydrogen, and oxygen gives

$$2a = c, \qquad 6a = 2d,$$
$$2b = 2c + d.$$

From the first two equations we obtain $c = 2a$ and $d = 3a$. Substitution into the oxygen equation gives $b = 7a/2$. Taking $a = 2$ clears the fraction and gives $(2,7,4,6)^T$, the primitive positive integer coefficient vector. We can now write

$$\mathrm{2C_2H_6 + 7O_2}$$
$$\rightarrow \mathrm{4CO_2 + 6H_2O}.$$

Multiplying all coefficients by the same positive integer preserves balance. Choosing the smallest positive integers fixes the normalization. We still need separate information about the amount in a vessel, combustion conditions, and the mechanism.

5.3 Rational solutions and integer normalization

Proposition 5.1. A rational matrix has a kernel basis consisting of rational vectors.

We can see this from Gaussian elimination. Its operations stay within the rational numbers when the inputs are rational and division is by nonzero entries. Set one free variable to one and the others to zero to obtain each special solution. Their entries are rational, and clearing denominators gives integer vectors.

The signs require a further check. An individual kernel basis vector may have unsuitable signs, even when a combination of basis vectors gives a feasible balance. With several free directions, we must solve an inequality problem as well. If the rational matrix has a nonempty feasible region with the required strict signs, that region contains rational points. Scaling such a point yields an integer balance on the selected sides.

Proposition 5.2. Exactly one of the following statements holds: there is a vector $c > 0$ with $Bc = 0$, or there is an elemental covector y such that

$$y^T B \geq 0, \qquad y^T B \neq 0.$$

The inequalities are coordinatewise. This is Stiemke's alternative (Stiemke, 1915). The covector is nonpositive on every displayed reactant composition and nonnegative on every displayed product composition, with at least one strict value. The alternatives cannot both hold: pairing the second row with a strictly positive c gives a positive number, whereas $y^T Bc = 0$. The theorem supplies the converse as well. Thus the sign question has a

dual certificate; it is not settled merely by finding a nonzero kernel.

Example 5.2. For $H_2 \rightarrow H_2O$, use element order hydrogen, oxygen. Then

$$B = \begin{pmatrix} -2 & 2 \\ 0 & 1 \end{pmatrix}, \qquad y = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The oxygen readout gives $y^T B = (0,1)$. No positive coefficients can remove this oxygen deficit. In general, such a certificate rules out a balance using every displayed species on its assigned side. It does not rule out all balances obtained by dropping species or changing sides. Section 19.1 records the corresponding validation data.

5.4 Underdetermined equations

Example 5.3. Suppose we allow both carbon monoxide and carbon dioxide as products. For

$$aC + bO_2 \rightarrow cCO + dCO_2$$

we obtain $a = c + d$ and $2b = c + 2d$. We can vary both c and d, so elemental conservation leaves the product ratio undetermined. To select that ratio, we have to add information from outside the balancing equations.

This matters when a program returns several nullspace directions. Taking the first one may omit a required species or select an arbitrary member of the solution family. We should inspect the family and

its signs, then state any extra rule used to choose one balance.

For an integer composition matrix, exact arithmetic gives a direct check: the rational residual must vanish exactly. A numerical singular value close to zero is a different kind of result and needs a tolerance. We return to noisy measurements and uncertain compositions in Chapter 16.

5.5 Integer kernels and images

Clearing denominators produces integer kernel vectors, but in more than one dimension it need not produce a basis for all integer solutions. We now retain the integer domain of Section 2.2 explicitly.

Definition 5.2. For an integer composition matrix with n species and m elemental rows, let $\rho = \mathrm{rank}A$. The integer kernel and cokernel are

$$L_A = \ker A \cap \mathbb{Z}^n,$$
$$\mathrm{coker}_{\mathbb{Z}} A = \mathbb{Z}^m / A\mathbb{Z}^n.$$

The first is a lattice of rank $n - \mathrm{rank}A$. An integer basis generates every member by integer coefficients. Smith normal form uses invertible integer row and column operations to give diagonal entries $d_1, \ldots, d_\rho$, with $d_i > 0$ and $d_i \mid d_{i+1}$. Write this diagonalization as UAV. The columns of V corresponding to zero columns of UAV form an integer kernel basis. The nonzero entries describe the cokernel:

$$\operatorname{coker}_{\mathbb{Z}} A \cong \left(\oplus_{i=1}^{\rho} \mathbb{Z}/d_i\mathbb{Z}\right) \oplus \mathbb{Z}^{m-\rho}.$$

These are integer changes of coordinates, with determinant $+1$ or -1, so their inverses also preserve integers. Smith (1861) developed the underlying theory. Hermite normal form provides another route to integer lattice calculations.

Example 5.4. Return to species order hydrogen, oxygen, water and elemental order hydrogen, oxygen:

$$A = \begin{pmatrix} 2 & 0 & 2 \\ 0 & 2 & 1 \end{pmatrix}.$$

The greatest common divisor of the entries is one; that of the two-by-two minors 4,2, -4 is two. The Smith invariants are therefore 1,2. The integer kernel has basis $(-2,-1,2)^T$, while the cokernel is $\mathbb{Z}/2\mathbb{Z}$. Every integer species combination has an even hydrogen count, and every integer elemental pair with even hydrogen count lies in the integer image. Odd hydrogen totals remain excluded even though the map is onto over $\mathbb{Q}$.

Here the image allows signed species counts. Nonnegative molecular inventories form the smaller monoid $A\mathbb{N}^3$, so membership in the integer image alone does not establish physical feasibility. The primitive balance in Section 5.2 is the rank-one version of the integer-kernel calculation. Section 20.3 returns to the torsion recorded by these quotients.

Part II Reactions and Linear Algebra

We can collect several reactions into one map and examine their combined changes. This brings out different kinds of zero result. A species change may leave elemental totals unchanged, while a combination of reactions may leave the species inventory itself unchanged. We introduce the corresponding nullspaces and then use them to describe conserved quantities and feasible states.

Chapter 6 Reaction Spaces and Conservation Laws

6.1 The stoichiometric map

Let us select r reactions and choose an orientation for each. For reaction j, write the reactant counts as α_j and the product counts as β_j in species coordinates. Subtracting the former from the latter gives its net change, $\nu_j = \beta_j - \alpha_j$.

Definition 6.1. The reaction-coordinate space is $R = \mathbb{R}^r$. The stoichiometric map $S: R \to W$ has column ν_j for reaction j. Its image is the stoichiometric subspace.

We can now supply a reaction-coordinate vector ξ and calculate the species change $S\xi$. When ξ represents accumulated reaction extent, its entries have units of amount of substance. If we are counting discrete events, the entries are integers. The matrix is the same, but the interpretation of its input changes with the model.

Proposition 6.1. All selected reactions conserve the modeled elemental totals if and only if $AS = 0$.

To check this, inspect the columns of AS. Each is the residual $A\nu_j$ for one reaction. All columns vanish exactly when every selected reaction balances. We consequently have $\mathrm{im} S \subseteq \ker A$. The inclusion may

be strict, since a model can select only some of the formally balanced transformations.

6.2 A network with one cycle

Example 6.1. Consider formal species X, Y, Z with identical elemental composition. Suppose our model admits $X \to Y$, $Y \to Z$, and $X \to Z$. Writing their net changes as columns gives

$$S = \begin{pmatrix} -1 & 0 & -1 \\ 1 & -1 & 0 \\ 0 & 1 & 1 \end{pmatrix}.$$

Notice that adding the first two columns gives the third. The rank is therefore two, and $(1,1,-1)^T$ belongs to $\ker S$. We can combine the first two changes and subtract the third without changing the final species inventory.

We often call such a dependence a *cycle vector*. Here its third coordinate is negative. If all three reactions are irreversible in the displayed directions, we cannot carry out that combination as three forward events. The nullspace describes an algebraic cancellation; a circulating physical flux requires compatible directions as well.

6.3 Left nullspaces

We can ask another question: which totals remain unchanged by every selected reaction? This leads us to the left nullspace.

Definition 6.2. A linear conserved quantity for the reaction model is a covector ℓ satisfying $\ell^T S = 0$. Its value on an inventory is $\ell^T n$.

For an inventory reached by a change $n = n_0 + S\xi$, we obtain

$$\ell^T n = \ell^T n_0 + \ell^T S\xi = \ell^T n_0.$$

In Example 6.1, take $\ell = (1,1,1)^T$. It reads the total number, or total amount, of the three species. Each modeled reaction converts one species unit into another, so the total remains fixed. A basis of the left nullspace gives us all independent linear conserved quantities for this stoichiometric map.

Elemental totals are included among these quantities. Every row of A annihilates S when $AS = 0$. We may also find conserved combinations that arise because the selected mechanism is restricted. Adding a missing reaction can remove one of these extra invariants, so its scope should be stated with the model.

6.4 Three kernels with different meanings

We now have three kernels to keep separate. The kernel of A collects species changes invisible to elemental counting. The kernel of S collects reaction combinations with no net species change. The kernel of S^T gives observables preserved by every selected reaction. Their vectors belong to different

coordinate domains, even though all three calculations involve a nullspace.

Rank–nullity tells us how many independent quantities to expect. With n species and rank s, we have $n - s$ independent linear conservation laws. With r reaction columns, we have $r - s$ independent reaction dependencies. These are structural counts and do not assume that a system is at equilibrium.

This viewpoint is developed in reaction-network theory; see Feinberg (2019) and Baez and Biamonte (2011). Here $AS = 0$ connects it with our original elemental representation. It also provides a practical check: a nonzero column of AS identifies a selected reaction whose encoded elemental balance fails.

Definition 6.3. For reactions treated as irreversible in their displayed directions, the steady-state flux cone is

$$\mathcal{F} = \{v \in \mathbb{R}^r : v \geq 0,\ Sv = 0\}.$$

Its extreme rays generate every feasible flux by nonnegative combinations. In this all-irreversible setting, they are the support-minimal elementary flux modes of Schuster and Hilgetag (1994). They are generators, not necessarily a linearly independent basis. With reversible coordinates, the distinction between elementary modes and extreme rays requires further care.

Example 6.2. A subsystem with two input channels and two output channels for one species has $S =$

$(1,1,-1,-1)$. Its nullity is three, but its flux cone has four extreme-ray generators:

$$(1,0,1,0)^T, \quad (1,0,0,1)^T,$$
$$(0,1,1,0)^T, \quad (0,1,0,1)^T.$$

Each joins one input to one output. These are exchange channels of an open subsystem; its external reservoirs are omitted. By contrast, the irreversible network of Example 6.1 has no nonzero nonnegative steady flux despite its one-dimensional kernel. Signs change the answer.

Chapter 7 Amounts Extents and Feasible Regions

7.1 States move in affine spaces

Suppose we know the initial amount of each species. Collect the amounts in n using a common unit, such as moles. Starting from n_0, we can express the result of a net extent vector as

$$n = n_0 + S\xi.$$

Definition 7.1. The nonnegative stoichiometric compatibility class through n_0 is

$$\mathcal{C}(n_0) = (n_0 + \mathrm{im}S) \cap \mathbb{R}^n_{\geq 0}.$$

Geometrically, we have an affine subspace cut by the nonnegative orthant. A state reached by the closed model must lie in this intersection. To determine whether we can actually reach a particular point, we still need to examine irreversible directions, available reactants, rate laws, and boundaries.

7.2 Limiting reagents as interval endpoints

Example 7.1. Take $2\mathrm{H}_2 + \mathrm{O}_2 \rightarrow 2\mathrm{H}_2\mathrm{O}$ with initial amounts $(3,2,0)^T$ mol in the order hydrogen, oxygen, water. The inventory at a specified extent is

$$n(\xi) = (3 - 2\xi,\ 2 - \xi,\ 2\xi)^T.$$

Use moles for the numerical extent ξ under this reaction normalization. Requiring forward progress and nonnegative amounts gives $0 \le \xi \le 3/2$. At the upper endpoint we obtain $(0,1/2,3)^T$ mol. Hydrogen has been exhausted first, so it is the limiting reactant in this calculation.

We have found the largest extent allowed by the stated irreversible conversion and the initial amounts. A physical reaction need not reach that endpoint. Equilibrium, reverse reactions, or additional products may stop or alter the conversion. The stoichiometric bound is still useful, provided we do not interpret it as a prediction of the final physical state.

7.3 Normalization changes the extent coordinate

Now divide every coefficient in the reaction equation by two. The reaction vector becomes $\nu/2$. To obtain the same change in amounts, we must use the extent 2ξ. Thus the displayed equation helps define the numerical extent.

Proposition 7.1. Replacing a nonzero reaction vector ν by $c\nu$, with $c \neq 0$, leaves its one-dimensional subspace unchanged. The extent for the same change transforms to ξ/c.

The amount change agrees because $(c\nu)(\xi/c) = \nu\xi$. When $c < 0$, the reaction orientation also reverses. We should therefore report an extent together with

its reaction equation or normalization. The convention changes its numerical value while preserving the amount change, as in the chemical quantity conventions of IUPAC (2023).

7.4 Several extents and nonuniqueness

The network in Example 6.1 gives another ambiguity. A conversion from X to Z can use the direct reaction or the first two reactions. The corresponding extent vectors differ by a vector in $\ker S$. Even a complete measurement of the final species amounts may therefore leave the reaction history unresolved.

Example 7.2. Start with $(1,0,0)^T$ mol. We can use the extent vector $(0,0,1)^T$ mol or $(1,1,0)^T$ mol; both produce $(0,0,1)^T$ mol. The first describes the direct conversion, and the second describes the two successive transformations. Their common endpoint contains no information that distinguishes these routes.

For several reactions, we obtain a convex region from $n_0 + S\xi \geq 0$. If the coordinates are accumulated irreversible forward extents, we also impose $\xi \geq 0$. This tells us which endpoints pass the amount constraints. We must separately check whether an ordering of reactions can reach an endpoint while keeping every intermediate inventory feasible.

We can use this geometry to inspect yields, boundary species, and redundant reaction coordinates. It also helps us check simulations. If a simulated state leaves its compatibility class without an external flow, we should look for a numerical error or an omitted model term.

7.5 Discrete counts and Petri nets

For discrete events, we retain the input counts α_j and output counts β_j of every reaction. This is exactly the data of an ordinary weighted Petri net: species are places, reactions are transitions, and molecule counts are tokens (Murata, 1989; Baez and Biamonte, 2011). Reverse reactions can be represented by separate transitions.

Definition 7.2. A marking $M \in \mathbb{N}^n$ is a discrete species inventory. Transition j is enabled when $M \geq \alpha_j$. Firing it produces $M - \alpha_j + \beta_j$. After a sequence with firing-count vector $\sigma \in \mathbb{N}^r$, the state equation is

$$M' = M_0 + S\sigma.$$

Every reachable marking satisfies this equation. Its converse fails because the equation forgets the order and enabling conditions.

Example 7.3. Take the single transition $X + E \to Y + E$ and the marking $(1,0,0)^T$ in species order X, Y, E. The state equation with $\sigma = 1$ gives the nonnegative endpoint $(0,1,0)^T$, but no event can

occur: the catalyst place is empty. The net column $(-1,1,0)^T$ has canceled the required catalyst.

In the nonnegative integer convention, a T-invariant is a nonzero $\sigma \geq 0$ satisfying $S\sigma = 0$. A P-invariant is a nonzero integer place-weight vector $p \geq 0$ satisfying $p^T S = 0$. Signed integer solutions give the broader invariant lattices. Thus the real kernels of Chapter 6 underlie these constructions, while integrality and signs select their discrete forms. A T-invariant alone still does not supply an enabled firing sequence.

This completes the ordering condition in Section 7.4. Adding propensities supplies the stochastic model of Section 10.4; adding compartment labels supplies a larger place set. The input and output data, rather than the incidence matrix alone, preserve the enabling requirements.

Chapter 8 Dual Spaces Charge and Mass

8.1 Observables as covectors

An inventory is a vector. To read one of its coordinates, or a weighted total, we need a map from that vector to a number. Linear maps of this kind form the dual space.

Definition 8.1. The dual space E^* consists of linear maps from E to the scalar field. A dual basis functional e^X extracts the coordinate belonging to element X.

For example, if $v = 2e_{\mathrm{H}} + e_{\mathrm{O}}$, applying e^{H} returns two. We can combine such readouts with weights to obtain a total. The basis specifies the directions in which an inventory can vary; the dual basis specifies the corresponding independent coordinate observations.

Example 8.1. Let us use approximate mass weights 1 for hydrogen and 16 for oxygen. Their weighted readout is

$$m = 1e^{\mathrm{H}} + 16e^{\mathrm{O}}$$

and it gives 18 for water. These integer values are convenient approximations for this example. In a molar interpretation, the weights are approximate molar masses of the elements (in grams per mole).

For an accurate calculation, we would supply the appropriate isotopic composition and physical conventions instead of treating these integers as current standard atomic weights.

8.2 Pulling a measurement back to species

Suppose our input is a species inventory rather than an elemental inventory. We can first apply A and then take the elemental measurement. If m is the column of elemental weights, the resulting species weights are $A^T m$. We call this composition of a measurement with a map a *pullback*.

Proposition 8.1. If $AS = 0$, every functional of the form $m^T A$ is conserved by the reaction model.

Associativity gives the check directly: $m^T AS = 0$. For atomic-mass weights, this recovers ordinary stoichiometric mass balance. Its precision is that of the chemical bookkeeping model. An exact account of rest mass and energy would require information beyond these additive atomic weights.

A total mass also illustrates information loss. Different elemental vectors can give the same weighted sum. Except for a nonzero functional on a one-dimensional domain, a single scalar linear observation leaves nonzero directions in its kernel. We cannot generally reconstruct composition from mass alone.

Example 8.2. For neutral formulas containing carbon, hydrogen, oxygen, and nitrogen, choose carbon dioxide, water, and ammonia as reference products. The reduction-equivalent functional is

$$\gamma = 4e^{\mathrm{C}} + e^{\mathrm{H}} - 2e^{\mathrm{O}} - 3e^{\mathrm{N}}.$$

For $\mathrm{C}_a\mathrm{H}_b\mathrm{O}_c\mathrm{N}_d$, its value is $4a + b - 2c - 3d$ per formula unit. This degree-of-reduction convention is used in biochemical balance calculations (Roels, 1980, 1983); a value per carbon atom additionally divides by a, when $a > 0$. The unnormalized functional vanishes on CO_2, $\mathrm{H}_2\mathrm{O}$, and NH_3, while giving four for $\mathrm{CH}_2\mathrm{O}$. Several different compositions therefore become invisible to this one readout. The reference products matter, and ionic charge or additional elements require an extended convention.

8.3 Charge is additional information

Now compare an ion with a neutral species containing the same nuclei. Elemental counts cannot distinguish them. We therefore assign each species a signed charge number z_j, expressed in units of the elementary charge.

Definition 8.2. The augmented balance map is

$$\tilde{A} = \begin{pmatrix} A \\ z^{T} \end{pmatrix}.$$

The test $\tilde{A}\nu = 0$ checks both elemental totals and charge. We have added another additive observation to the codomain. The extra row records charge; it

does not add a chemical element to the elemental vocabulary.

Example 8.3. Consider $\mathrm{Fe}^{2+} \rightarrow \mathrm{Fe}^{3+}$. Iron count agrees, but the charge residual is $+1$. We can balance charge by placing an electron on the product side:

$$\mathrm{Fe}^{2+} \rightarrow \mathrm{Fe}^{3+} + e^{-}.$$

In this representation, the electron contributes zero to every elemental count and -1 to the charge row. The augmented residual vanishes. We have extended the species list and the observations made on it, while keeping the elemental basis unchanged.

The charge row can also be converted to an electrical quantity. Faraday's constant $F = N_A e$ is the charge magnitude per mole of electrons, approximately $96485\ \mathrm{C\ mol^{-1}}$. If the signed electron coefficient of a half-reaction is ν_e, a forward extent ξ transfers charge magnitude $F|\nu_e|\xi$ through the associated electron channel. The electron charge itself is $-F\nu_e\xi$; the sign of a reported circuit current follows its declared orientation. For the iron half-reaction above, one mole of extent corresponds to one mole of electrons. This interpretation assumes that the measured charge is assigned to that half-reaction; side reactions require their own accounting (IUPAC, 2023).

8.4 Conservation and electroneutrality

Charge conservation and electroneutrality answer different questions. Internal reactions in a closed system preserve total charge; an exchange term is needed to change it. Electroneutrality specifies zero total charge at the modeled scale. It may be a state constraint or an approximation used for that scale.

For example, $z^T S = 0$ preserves whatever charge the initial inventory has, including a nonzero value. Conversely, starting with $z^T n = 0$ does not correct a mechanism for which $z^T S \neq 0$. Such a mechanism can leave the neutral class unless we account for an external charge flow.

Oxidation numbers add a further level of bookkeeping. Their assignment depends on bonding rules and conventions. We should retain that interpretation when using them, rather than identify every oxidation number with a measured ionic charge or an independent elemental coordinate.

Chapter 9 Coordinates Projections and Lost Information

9.1 A new basis is not new chemistry

Suppose we describe the same elemental vector in a new basis. Let P contain the new basis vectors in the old coordinates. Writing the new coordinates as v' gives $v = Pv'$, and solving for them gives $v' = P^{-1}v$.

We can change the species coordinates as well. If $w = Qw'$, the composition matrix becomes

$$A' = P^{-1}AQ.$$

We have changed the description of the map. Its rank and kernel dimension stay the same. Its entries may change, as may lengths computed with an unadjusted dot product. Even coordinatewise nonnegativity depends on the basis used to describe the physical cone.

Example 9.1. Starting with the hydrogen and oxygen basis, choose

$$f_1 = e_{\mathrm{H}}, \qquad f_2 = 2e_{\mathrm{H}} + e_{\mathrm{O}}.$$

Water now has coordinates $(0,1)^T$, while the oxygen basis vector has coordinates $(-2,1)^T$. The negative entry has appeared through a change of basis, without any negative physical inventory. To

interpret the new coordinates, we must transform the physical cone along with them.

9.2 What counts as an inner product

We can declare the elemental basis orthonormal and use the ordinary dot product. This is a mathematical choice. Equal weighting and perpendicularity of hydrogen and oxygen directions do not follow from a chemical law.

More generally, we specify a positive-definite matrix G and use $\langle v, w\rangle_G = v^T G w$. Under the coordinate change $v = Pv'$, the same metric has matrix $G' = P^T G P$. Leaving the matrix unchanged would usually alter the measurement itself.

This becomes relevant when we compare formulas for similarity. Raw atom counts, normalized fractions, mass fractions, and learned descriptors can give different notions of closeness. We should choose the measure according to the question being asked. A small coordinate distance alone gives no conclusion about reactivity or molecular structure.

9.3 Projection as an observational choice

Sometimes we record only part of a species inventory, or a combination of its entries. We describe this selection by an observation map.

Definition 9.1. An observation map is a specified map $H: W \to Y$ from species amounts to recorded

quantities. When the observations are additive, H is linear.

Examples include dropping an elemental coordinate, aggregating isotope counts, or summing selected species amounts. Any change in the kernel of H is invisible to the resulting observation. We can therefore study what a measurement misses by inspecting its kernel.

Example 9.2. Suppose $H = (1,1)$ records only the sum of two isomer amounts. Both $(1,0)^T$ and $(0,1)^T$ give 1. We can distinguish them by adding the independent readout $(1,0)$. The combined matrix is

$$\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix},$$

which is invertible. We have supplied enough independent information to reconstruct the two coordinates. For a real instrument, we would still need to establish that the second readout has the assumed selectivity and accuracy.

9.4 Reduced models need compatible dynamics

A smaller observation does not automatically give a smaller dynamical model. Suppose $\dot{n} = Kn$ and we retain only $y = Hn$. Then $\dot{y} = HKn$. To express this as a closed equation $\dot{y} = \overline{K}y$ for all states, we need a matrix $\overline{K}$ such that

$$HK = \overline{K}H.$$

Equivalently, we require $K(\ker H) \subseteq \ker H$. Compare two states with the same observation. Their observed derivatives must agree as well; otherwise the reduced state would have two different futures at that instant. When this holds, we define the reduced map on the image of H, giving sufficiency.

Two species inventories with equal elemental totals may evolve differently. Their shared totals can remain constant throughout, but that observation alone does not describe the changing species. A conserved observation and a complete reduced kinetic model therefore provide different amounts of information.

Chapter 10 From Stoichiometry to Dynamics

10.1 The rate map

To follow a reaction in time, we need a rate as well as a direction. For a well-mixed system of constant volume, let x contain species concentrations and $v(x)$ contain reaction rates in concentration per unit time. We then write

$$\dot{x} = Sv(x).$$

Here the stoichiometric matrix acts linearly on the rates. The dependence of the rates on the state may be nonlinear. For an elementary mass-action model, reactant counts α_j give the monomial

$$v_j(x) = k_j \prod_i x_i^{\alpha_{ij}}.$$

We have now added a kinetic assumption. An overall balanced equation generally does not determine its own rate law. Intermediates, nonideal behavior, or an effective description can lead to other expressions. Horn and Jackson (1972) and Feinberg (2019) develop the mathematical framework; we use small examples to follow the relation between the structural and kinetic parts.

10.2 A reversible conversion

Example 10.1. Consider $X \rightleftharpoons Y$ with first-order constants k_+ and k_- for the forward and reverse directions. The conserved total is $c = x + y$, and the rate equations are

$$\dot{x} = -k_+ x + k_- y,$$

$$\dot{y} = k_+ x - k_- y.$$

We can use the conservation relation $y = c - x$ to reduce them to

$$\dot{x} = -(k_+ + k_-)x + k_- c.$$

Solving this scalar equation gives

$$x_* = \frac{k_- c}{k_+ + k_-},$$

$$x(t) = x_* + (x(0) - x_*)e^{-(k_+ + k_-)t}.$$

For a numerical example, take $k_+ = 2\,\mathrm{s}^{-1}$, $k_- = 1\,\mathrm{s}^{-1}$, $c = 3\,\mathrm{mol\,L}^{-1}$, and $x(0) = 3\,\mathrm{mol\,L}^{-1}$. The equilibrium concentration of X is then $1\,\mathrm{mol\,L}^{-1}$. Its excess over equilibrium is $2e^{-3t}$ in the stated concentration units, with t in seconds.

The kinetic matrix has eigenvalues zero and $-(k_+ + k_-)$. We can interpret them using the conservation classes. The zero eigenvalue allows different fixed totals, while the negative eigenvalue describes relaxation within one such class. The zero value alone does not indicate instability of that class.

10.3 Linearization of a nonlinear system

For nonlinear rates, we can examine behavior near a steady state. Write $f(x) = Sv(x)$, choose a steady state x_*, and consider a small perturbation h. The first-order approximation is

$$\dot{h} = Jh, \qquad J = S\,Dv(x_*).$$

Definition 10.1. The Jacobian J is the linear map of first derivatives of the vector field at the specified state.

We use its eigenvalues to study local linearized behavior on the relevant invariant subspace. Zero or purely imaginary eigenvalues need the usual further analysis. These are generally not eigenvalues of the stoichiometric matrix: S may not even be square. The rates and their derivatives are part of the calculation.

The conservation relations survive this approximation. From $\ell^T S = 0$ we obtain $\ell^T J = 0$. We should therefore test stability on perturbations compatible with the fixed totals. An ambient direction that changes the total inventory can otherwise be mistaken for a neutral motion available within the same physical class.

10.4 Counts and stochastic events

For small molecule counts, we may prefer to describe individual events. When reaction j occurs,

it changes the integer state N by ν_j. The occurrence time is random, with a state-dependent event rate called a propensity. Gillespie (1977) develops stochastic simulation in this setting, alongside the chemical master-equation description.

Consider a reaction consuming two identical molecules of X. The available unordered pairs number $N_X(N_X - 1)/2$. Using N_X^2 would include self-pairs and count distinct pairs twice. Different conventions absorb combinatorial factors into stochastic rate constants, so those conventions must be stated when comparing molecule-count and concentration models.

A further difference appears when we average nonlinear rates. The mean of N_X^2 depends on the variance as well as the mean count. In general, evaluating a rate at the mean state does not give its mean value. Sharing a stoichiometric matrix therefore does not make a deterministic trajectory the exact mean of a stochastic model.

Part III Tensor Algebra and Forms

So far, we have combined vectors and applied linear observations to them. Some questions require more information. We may want to compare two deviations, retain both species and location, or accumulate a quantity along a changing state. We introduce forms and tensor products through these uses, keeping track of what their arguments represent.

Chapter 11 Bilinear and Quadratic Forms

11.1 Two linear arguments

Suppose a comparison depends on two vectors and is linear in either one while the other is fixed. We can describe it by a bilinear form.

Definition 11.1. A bilinear form on a real vector space V is a map $b: V \times V \to \mathbb{R}$ that is linear separately in each argument. In coordinates it can be written

$$b(u, v) = u^T G v$$

with a fixed matrix G. We obtain a symmetric form when $G^T = G$. Positive definiteness requires $b(v, v) > 0$ for every nonzero v.

These properties need to be checked separately. A bilinear form may be nonsymmetric, indefinite, or degenerate, so bilinearity alone does not make it an inner product. We use the corresponding geometric language only after establishing its conditions. See Axler (2024, Chapters 6 and 9) for these distinctions.

Our elemental basis has not selected a chemical inner product. We must therefore explain the choice of G. It may encode measurement weights, come from local energy curvature, or simply define a

geometry for an example. The same algebra can support these different interpretations.

11.2 Quadratic measurements

We can also feed the same vector into both arguments. This gives a quadratic form.

Definition 11.2. The quadratic form associated with a bilinear form is $q(v) = b(v, v)$. Over the real numbers, only the symmetric part of G contributes to $v^T G v$.

To see what is lost, take $K^T = -K$. The scalar $v^T K v$ is equal to its transpose, $v^T K^T v = -v^T K v$, and hence must be zero. An antisymmetric part is therefore invisible to this repeated-argument measurement.

Example 11.1. Let us use $v = (h, o)^T$ for two dimensionless deviations and choose

$$G = \begin{pmatrix} 1 & 1/2 \\ 1/2 & 2 \end{pmatrix}.$$

The resulting expression is $q(v) = h^2 + ho + 2o^2$. We can check its sign by completing the square:

$$q(v) = (h + o/2)^2 + 7o^2/4.$$

The expression is positive for every nonzero deviation, so the form is positive definite. At $(h, o) = (1, -1)$ we obtain 2. A negative coordinate is appropriate here because we are recording deviations. We have specified a weighted

comparison for this example; a chemical energy interpretation would need an additional model.

11.3 Local energy curvature

For a physical example of a quadratic form, suppose a thermodynamic model supplies a twice differentiable free energy $G_f(n)$ at fixed external conditions. Near the reference state n_*, we can write

$$G_f(n_* + h) \approx G_f(n_*) + \mu_*^T h + 1/2\, h^T H_f h,$$

where μ_* is the first derivative and H_f is the Hessian at the reference state. We have retained terms through second order. Continuity of the relevant second derivatives makes the Hessian symmetric.

Now restrict the perturbations to the modeled reaction directions, $h = S\eta$. Substitution gives the quadratic term

$$1/2\, \eta^T \left(S^T H_f S\right) \eta.$$

We have pulled the species-space curvature back to reaction coordinates. At a constrained stationary state, positive curvature on every nonzero allowed state perturbation gives a local minimum criterion. We must test the actual perturbation: a nonzero η can still satisfy $S\eta = 0$ when reaction coordinates are redundant. The resulting zero value then records a redundant description of the zero perturbation.

Example 11.2. Consider a single ideal mixture at fixed temperature and pressure, with positive species amounts and $N = \sum_i n_i$. Its Gibbs energy is

$$G_f(n) = \sum_i n_i \, \mu_i^\circ + RT \sum_i n_i \ln(n_i/N).$$

Differentiating gives $\mu_i = \mu_i^\circ + RT\ln(n_i/N)$ and the Hessian entries

$$\left(H_f\right)_{ij} = RT\left(\frac{\delta_{ij}}{n_i} - \frac{1}{N}\right).$$

Here δ_{ij} is one when $i = j$ and zero otherwise. For a perturbation h,

$$h^T H_f h = RT\left[\sum_i \frac{h_i^2}{n_i} - \frac{(\sum_i h_i)^2}{N}\right] \geq 0.$$

Cauchy–Schwarz applied to $h_i/\sqrt{n_i}$ and $\sqrt{n_i}$ gives the inequality. Equality holds exactly when h is proportional to n. Scaling every amount preserves composition, so this one direction has zero curvature. Every other direction has positive curvature. In particular, a nonzero change with $\sum_i h_i = 0$ has positive curvature. A closed compatibility class with a strictly positive conserved weight also excludes the scaling direction.

For $X \rightleftharpoons Y$, let x, y now denote positive amounts, not the concentrations of Chapter 10. With $S = (-1,1)^T$, the pullback is

$$\eta^T S^T H_f S \eta = RT\eta^2 \left(\frac{1}{x} + \frac{1}{y}\right).$$

For amounts $x = 1$ mol and $y = 2$ mol, the coefficient is $3RT/2$ per mole. We have obtained a concrete positive quadratic form on the conversion direction. It belongs to the stated ideal-mixture potential; a nonideal model requires its own Hessian. Section 15.4 compares the associated equilibrium equations with conservation classes.

11.4 Distances and uncertainty weights

An uncertainty calculation gives another use of the same algebra. Suppose two observations differ by d and a positive-definite covariance matrix Σ describes their uncertainty on a fixed coordinate scale. We can form the discrepancy

$$d^T \Sigma^{-1} d$$

which accounts for uncertainty and correlation. The same numerical deviation receives a different weight in a poorly measured direction than in a tightly measured one.

Before taking the inverse, we should check the covariance model. Exact constraints can make Σ singular. We can then work with independent coordinates or justify a generalized inverse for the intended calculation. Obtaining a number from a routine does not by itself resolve which directions are being compared.

These examples share the form of a bilinear calculation, but their sources differ. We chose one metric to compare composition vectors, obtained another form from a specified potential, and used an uncertainty model for a third. Keeping the source of the form visible prevents us from transferring a physical interpretation merely because the matrices look alike.

Chapter 12 Tensor Products and Chemical Context

12.1 Two kinds of labels

Suppose we need to distinguish water in one compartment from water in another. A species label by itself cannot do this. Let us introduce a compartment space C with basis c_1, c_2, while retaining the species space already constructed.

Definition 12.1. The tensor product $W \otimes C$ has basis vectors $u_i \otimes c_a$, one for each species–compartment pair, with bilinearity imposed in the two factors.

We can now record an inventory using both labels:

$$N = \sum_{i,a} n_{ia}\; u_i \otimes c_a .$$

If W has dimension n and C has dimension q, there are nq basis pairs. The product records which species occurs in which compartment. We have combined two kinds of labels through the usual tensor-product construction; no chemical bond is implied by the product symbol.

12.2 Pure tensors and correlations

A particularly simple tensor is $u \otimes c$. Its coefficients are $u_i c_a$, giving a matrix of rank at most one. For a nonzero nonnegative inventory of this kind, each occupied compartment has the same relative species profile, with only its scale changing. General compartment inventories need not have this factorization.

Example 12.1. Consider the following array, with species as rows and compartments as columns:

$$N = \begin{pmatrix} 2 & 1 \\ 0 & 3 \end{pmatrix}$$

The first compartment contains two units of the first species and none of the second. The second contains one and three. Its determinant is six, so its rank is two and it cannot be a single outer product. Replacing it by a product of marginal profiles would remove this association between species and location.

Here we are using ordinary matrix rank for a tensor with two factors. With three or more factors, different definitions of tensor rank need not agree. Decompositions can help us analyze multiway chemical data, but a factor still needs evidence for its interpretation as a physical component. Kolda and Bader (2009) survey these methods, including their use in chemometrics.

12.3 Composition and aggregation by contraction

We can act on one label while retaining the other. The map $A \otimes I_C$ converts species composition within each compartment. To sum compartments instead, let $\sigma \in C^*$ give one on every compartment basis vector. Contracting the compartment factor gives

$$N_{\text{total}} = (I_W \otimes \sigma)N.$$

Proposition 12.1. We obtain the same elemental total whether we first calculate composition in each compartment and then sum, or first sum the species inventories and then calculate composition:

$$A(I_W \otimes \sigma)N$$
$$= (I_E \otimes \sigma)(A \otimes I_C)N.$$

For a basis pair $u_i \otimes c_a$, both routes give Au_i. Linearity extends the result to every inventory. This supplies a useful check when a data pipeline can aggregate material either before or after converting species amounts to elemental totals.

12.4 Transport and local reactions

For example, moving species i from compartment 1 to compartment 2 has direction

$$u_i \otimes (c_2 - c_1).$$

Contracting over compartments sends this direction to zero. We can see the transfer in the separate compartments but not in their aggregate species amount. A local chemical reaction changes species identity instead, while preserving the elemental totals encoded by the model.

Suppose every compartment has the same local linear reaction dynamics K, and a species-independent linear operator L describes transport. We can combine them as

$$\mathcal{K} = K \otimes I_C + I_W \otimes L.$$

This expression uses species as the first tensor factor. A different vectorization order changes the displayed Kronecker matrices, so the software representation must state that order. If transport depends on species, volumes differ, reactions depend on the compartment, or kinetics is nonlinear, we need to extend the construction accordingly.

12.5 Tensor language and physical tensors

We should also distinguish the spaces behind different uses of tensor language. Our species–compartment tensor carries chemical and location labels. A spatial diffusion tensor relates gradients to fluxes and transforms under changes of spatial coordinates. A quantum state may use a tensor product of state spaces. The algebra is related, while the factors and their interpretation differ.

For an array to serve as a useful tensor representation, we specify its spaces, coordinate transformations, units, and permitted contractions. This tells us which operations on the array correspond to operations on the modeled objects. Naming the array a tensor without that information would leave those questions unanswered.

Chapter 13 Polynomial Kinetics as Tensor Contraction

13.1 Multilinear maps and diagonal evaluation

The product of two concentrations in a rate law suggests a map with two inputs. We first allow the inputs to vary independently and then set them equal.

Definition 13.1. A multilinear map is linear separately in each of its arguments. A bilinear map $B: W \times W \to W$ can be evaluated on the repeated argument (x, x) to produce a quadratic vector-valued function.

For a quadratic kinetic field, we can express this in coordinates as

$$f_i(x) = \sum_{j,k} T^i_{jk}\, x_j x_k.$$

Each coefficient has one output index and two input indices. Together these coefficients represent a tensor. Changing the bases changes its coordinate array according to the transformation law of the underlying map.

The two-input map is linear in either input separately. Once we use the same state in both slots, $x \mapsto B(x, x)$ is generally quadratic. This repeated-

argument operation explains why a multilinear tensor can represent nonlinear dynamics without making the original species evolution linear.

13.2 A bimolecular example

Example 13.1. Consider the association model $X + Y \to Z$ with rate kxy. In the order X, Y, Z, its reaction vector is $\nu = (-1, -1, 1)^T$. We can use the symmetric bilinear map

$$B(a, b) = \frac{k}{2}\nu(a_X b_Y + a_Y b_X).$$

Setting the arguments equal gives $B(x, x) = kxy\nu$. The two cross terms contribute equally, so the factor one-half gives the required coefficient. We could instead put the whole coefficient on one input order and use a nonsymmetric map. Both choices give the same diagonal polynomial. A convention is therefore needed to select a unique tensor representation.

The units also come from the rate model. For concentrations in mol L^{-1}, the constant k has units $\mathrm{L\,mol^{-1}\,s^{-1}}$ when rates are measured in concentration per second. These units belong to the kinetic tensor entries and distinguish them from dimensionless stoichiometric counts.

13.3 Repeated reactants and polynomial degree

Repeated reactants require us to state the rate convention. For the deterministic model $2X \to Y$, choosing $v = kx^2$ gives $\dot{x} = -2kx^2$ and $\dot{y} = kx^2$. There is no additional compulsory one-half in the xx coefficient when k is defined by this convention. The stochastic pair-counting construction of Chapter 10 uses a different interpretation.

For a third-order elementary rate, we use three concentration factors, or a trilinear map evaluated at (x, x, x). A general polynomial field can be separated into constant, linear, quadratic, and higher-degree parts. Each degree has its own tensor order and associated units. Keeping the parts separate makes the rate expression easier to inspect.

Saturation and inhibition can lead to nonpolynomial rate laws. We may approximate such a law locally by a polynomial, but we should state the approximation and where it is valid. The tensor representation describes a supplied law; atom balance alone does not supply that law.

13.4 Conservation as a tensor constraint

We can retain the conservation property while changing the rate representation. If $f(x) = Sv(x)$ and $AS = 0$, then $Af(x) = 0$ wherever the rates are defined. In a symmetric quadratic representation,

applying the elemental map to the output index must therefore annihilate the effective quadratic coefficients.

Proposition 13.1. If a bilinear map B satisfies $AB(a, b) = 0$ for all a, b, its diagonal field $f(x) = B(x, x)$ preserves elemental totals.

Substituting the same state into both arguments gives $Af(x) = AB(x, x) = 0$. This suggests a way to construct a fitted field: place its outputs in $\ker A$, or make them pass through a selected matrix S. Conservation is then guaranteed algebraically. Positive concentrations, accurate rates, and thermodynamic consistency still require their own checks.

For positivity, consider a boundary state with $x_i = 0$. An admissible concentration field must not point through that boundary into negative values of this coordinate. We can choose polynomial coefficients that conserve every elemental total and still violate this requirement. The two properties are therefore checked separately.

13.5 Tensor powers do not close dynamics automatically

We may try to make the polynomial field linear by introducing variables for products such as $x_i x_j$. Differentiate them using the product rule. A quadratic original field generally produces cubic terms in these new equations. Adding those terms as

further coordinates can produce still higher orders. The process need not close after finitely many additions.

Tensor lifting is useful for exposing these relationships and constructing approximations. An exact finite-dimensional linear closure requires special structure. If we truncate the lifted system, we should state where we truncated and compare its trajectories and invariants with those of the original model.

Chapter 14 Exterior Algebra and Alternating Forms

14.1 Alternation as a mathematical choice

Sometimes we want a construction that vanishes when two directions coincide. We can impose this condition on a bilinear form.

Definition 14.1. An alternating bilinear form ω assigns a scalar to two vectors, is linear in each argument, and satisfies $\omega(v, v) = 0$ for every vector v.

Apply the repetition rule to $u + v$ and expand. Over any field, we obtain $\omega(u, v) = -\omega(v, u)$. Swapping the arguments therefore changes the sign. If one argument is a multiple of the other, the value is zero.

For vectors themselves, we introduce the exterior product $u \wedge v$. Its corresponding rules are

$$u \wedge v = -v \wedge u, \qquad v \wedge v = 0.$$

We should pause before giving the product a chemical interpretation. Two atoms of the same element still contribute two atoms to a composition, whereas a vector wedged with itself gives zero. Exterior algebra is therefore suited to questions about independent directions and orientation. For multiplicity, we retain composition sums or an appropriate symmetric construction.

14.2 A two-dimensional calculation

Example 14.1. Use a two-dimensional basis e_1, e_2 and take $u = 2e_1 + e_2$ and $v = e_1 + 2e_2$. Expanding their wedge and removing the repeated terms gives

$$u \wedge v = 3e_1 \wedge e_2.$$

The coefficient is the determinant of the two columns. Since it is nonzero, the vectors are independent. With an orthonormal basis in a chosen metric, its absolute value also gives the parallelogram area. Without that metric, we have an oriented coordinate-area coefficient, with no automatic physical-area interpretation.

Proposition 14.1. The wedge of k vectors in a finite-dimensional vector space is nonzero exactly when those vectors are linearly independent.

If one vector is a combination of the others, expansion produces terms with repeated vectors, all of which vanish. For independent vectors, we can extend the family to a basis. Their wedge is then a nonzero basis element of the relevant exterior power. This gives the rank test in any finite dimension.

14.3 Reaction directions and rank

Return to Example 6.1. Wedging its first two reaction vectors gives a nonzero result, so we have two independent species-change directions. Wedging all three gives zero because the third is the

sum of the first two. We have recovered the rank calculation through exterior algebra.

This observation concerns the structure of the reaction columns. A slow reaction may add an independent direction, while a fast reaction may have a net change already spanned by other columns. Energy, rate, and causal importance require different information.

In matrix coordinates, rank s means that some $s \times s$ minor is nonzero and all larger minors vanish. These determinants express the corresponding exterior-power information. They are convenient for small symbolic examples. For a larger numerical matrix, elimination or singular-value methods may be preferable, with tolerances specified for approximate data.

14.4 A special three-coordinate shortcut

Example 14.2. For hydrogen, oxygen, and water, the two composition rows are $r_H = (2,0,2)$ and $r_O = (0,2,1)$. Their three-coordinate cross product is

$$r_H \times r_O = (-4,-2,4).$$

The result is orthogonal to both rows. We can therefore transpose it into a kernel column and divide by two to obtain $(-2,-1,2)^T$. This gives the primitive water balance.

We can use the cross product as a shortcut in this three-coordinate example. It depends on the

selected Euclidean structure and orientation. With a different number of species, the shortcut does not extend unchanged. The general balancing operation remains the kernel calculation.

14.5 Forms versus vectors

A k-vector belongs to an exterior power of the vector space. A k-form receives k vectors and returns a scalar. Their natural pairing does not identify them; such an identification needs additional structure, for example a metric.

Reaction directions are vectors, while coordinate differentials are covectors. In Chapter 15, a thermodynamic model supplies a differential form whose pairing with a reaction direction expresses a scalar change at a particular state.

For bilinear forms, alternating multilinear forms, and tensor products, see Axler (2024, Chapter 9). For exterior powers, wedge products of vectors, and differential forms, see Sjamaar (2017).

Chapter 15 Differential Forms and Thermodynamic Change

15.1 Forms at a state

Suppose the amount vector n varies over an open region. At a particular state, a small displacement h is a tangent vector. We use dn_i to read its ith component. The readout is a covector at that state.

Definition 15.1. A differential one-form is a state-dependent linear functional on small changes. In coordinates,

$$\theta = \sum_i a_i\,(n)\,dn_i.$$

Along a differentiable path $n(t)$, we accumulate its values by integrating:

$$\int \theta = \int \sum_i a_i\,\big(n(t)\big)\dot{n}_i(t)\,dt.$$

The form assigns a value to each infinitesimal displacement along the path. To determine the path itself, we need a rate law. We can therefore separate how the system moves from what we evaluate during that motion.

15.2 Exact forms and path dependence

Some accumulated quantities can be found from the endpoints alone. We describe them using exact forms.

Definition 15.2. A one-form is exact if it is the differential dF of a scalar function F. It is closed if its exterior derivative is zero.

With the required smoothness, exactness implies closedness. To reverse that implication, we also need information about the domain. It holds on a convex open domain and under suitable more general topological conditions, but can fail on other domains. See Sjamaar (2017).

Example 15.1. Take $\theta = 2x\,dx + 2y\,dy = d(x^2 + y^2)$ on the plane. We can integrate from $(0,0)$ to $(1,2)$ simply by subtracting the endpoint values of the potential. The result is five for every admissible path between those points.

Now take $\omega = x\,dy$ and follow the counterclockwise unit square. The successive edges contribute zero, one, zero, and zero, giving a loop integral of one. Its exterior derivative, $d\omega = dx \wedge dy$, agrees with the oriented-area calculation. This example shows path dependence mathematically; we have not supplied a chemical heat model.

15.3 Chemical potentials as a one-form

For a thermodynamic application, consider a differentiable Gibbs free energy $G(T, p, n)$ with only the standard pressure-volume work variables. Its fundamental differential is

$$dG = -\mathcal{S}\, dT + V\, dp + \sum_i \mu_i\, dn_i.$$

We use $\mathcal{S}$ for entropy to distinguish it from the reaction map S. Holding temperature and pressure fixed leaves

$$dG = \mu^T dn.$$

The chemical potentials thus give an exact one-form on the smooth composition domain of this model. Integrating it produces a Gibbs free-energy difference. Heat transfer generally requires a different interpretation. At a phase boundary or another loss of smoothness, we must check the conditions for using the local differential.

We can express the same form in reaction coordinates. Substitute $dn = S\, d\xi$ to obtain its pullback:

$$dG = (S^T \mu)^T d\xi.$$

We recognize $\Delta_r G_j = \nu_j^T \mu$ as the reaction Gibbs energy for the chosen normalization. With the affinity convention $\mathcal{A}_j = -\Delta_r G_j$, its sign is reversed. Rescaling a reaction therefore rescales the value per

unit extent, while the paired total free-energy change stays the same.

15.4 A reaction-level example

Example 15.2. For $X \rightleftharpoons Y$, use the direction $(-1,1)^T$. Pairing it with the chemical potentials gives $\Delta_r G = \mu_Y - \mu_X$. A forward infinitesimal extent contributes this difference multiplied by the extent to the free-energy change.

We can make the chemical potentials explicit through an idealized activity description:

$$\mu_i = \mu_i^\circ + RT\ln a_i.$$

Here a_i is dimensionless and uses a specified standard state. For an unconstrained reversible direction, equilibrium requires $\Delta_r G = 0$. The standard chemical-potential difference then determines the activity ratio. We need a reference concentration if concentrations are used to form activities; a logarithm of a dimensional concentration would not be meaningful.

This criterion comes from thermodynamics rather than from the numerical method used to integrate the rate equations. A kinetic model intended to approach that equilibrium must satisfy the corresponding relation between its forward and reverse constants. We cannot choose those constants arbitrarily and also impose an unrelated equilibrium constant.

For a network of reversible directions, put $z = \ln a$ coordinatewise. The condition $S^T \mu = 0$ becomes

$$S^T z = -S^T \mu^\circ / (RT).$$

Its algebraic solution set is $-\mu^\circ/(RT) + \ker S^T$, an affine set in log-activity coordinates. The conservation-covector space of Section 6.3 has reappeared. Activities must still obey their physical model: ideal mole fractions, for example, sum to one. Taking logarithms is a nonlinear change of coordinates, not a linear basis change of Chapter 9.

Proposition 15.1. For a fixed-volume, detailed-balanced mass-action model with a positive equilibrium x^*, the positive equilibrium set satisfies

$$\ln(x/x^*) \in \ker S^T.$$

The ratio is coordinatewise and dimensionless. Each positive stoichiometric compatibility class has exactly one such equilibrium. This is the detailed-balanced case of Horn and Jackson (1972); the intersection geometry is also expressed by Birch's theorem (Birch, 1963; Henriksson and colleagues, 2025).

We can check uniqueness directly. If two distinct positive equilibria x, u were in the same class, then $x - u \in \mathrm{im} S$, while $\ln(x/u) \in \ker S^T$. Their pairing would be zero. Strict increase of the logarithm instead gives

$$\sum_i (x_i - u_i)\ln(x_i/u_i) > 0.$$

The contradiction proves at most one intersection; the cited theorem supplies existence under its hypotheses. Here concentrations are proportional to ideal dilute activities at fixed volume. The result does not follow for arbitrary kinetics or arbitrary nonideal activity maps. For $X \rightleftharpoons Y$, the two descriptions reduce to a fixed positive ratio $y/x = K$ and the conserved total $x + y = c$, giving $x = c/(1 + K)$ and $y = Kc/(1 + K)$.

15.5 Dissipation and its assumptions

For a closed, isothermal, isobaric reaction model under the usual thermodynamic assumptions, the chemical entropy production is $\mathcal{A}^T\dot{\xi}/T$. Its nonnegativity provides a consistency check on the dissipative model. For an open system, we must also include the relevant material flows and work terms before drawing the corresponding conclusion.

Near a suitable equilibrium, we may use the linear response law $\dot{\xi} = L\mathcal{A}$. A positive-semidefinite symmetric part of L makes the quadratic production nonnegative. Symmetry of L itself needs further physical assumptions; it does not follow from atom balance. Rao and Esposito (2016) give the broader thermodynamic treatment of open reaction networks. Here we use the relation to connect affinity, rates, and bilinear forms.

Chapter 16 Measurements Inverse Problems and Uncertainty

16.1 An observation model

Suppose we have instrument readings and want to infer the species amounts that produced them. A local linear model for the readings is

$$y = Hn + \varepsilon.$$

The matrix H describes the response and ε the error. Its rows may represent elemental totals, calibrated spectral responses, or other additive signals. The algebra works with the supplied response; calibration and physical validity have to be established for the instrument and conditions being modeled.

Definition 16.1. A state is identifiable within a specified admissible set if no other state in that set gives the same exact observation.

Full column rank of H gives a unique unrestricted inverse on its image. Constraints can sometimes select a unique admissible state even when full rank is absent. We should also check sensitivity: an invertible matrix may strongly amplify small errors, making a unique algebraic answer unreliable as a measurement result.

16.2 A worked reconstruction

Example 16.1. Let two species amounts a, b produce

$$y_1 = a + b, \qquad y_2 = a + 2b.$$

With $y = (3,5)^T$, subtraction gives $b = y_2 - y_1 = 2$ and then $a = 2y_1 - y_2 = 1$. Both amounts are nonnegative. If the recorded values are instead $(3,7)^T$, solving gives $a = -1$ and $b = 4$. The negative amount tells us that these observations do not fit the exact nonnegative two-species model.

Several explanations remain possible: measurement error, an omitted species, an incorrect response matrix, or failure of the linear approximation. The calculation locates a mismatch. Choosing its cause requires additional observations or checks of the model.

16.3 Constraints and residuals

For noisy observations, we can choose a state by minimizing a weighted residual while imposing the appropriate physical constraints:

$$\min_{n \geq 0} \ (Hn - y)^T Q (Hn - y),$$

$$\text{subject to } An = b.$$

We choose a positive-semidefinite weight Q, often from an uncertainty model. The elemental equality is a hard constraint only when those totals are treated

as exact. If they are uncertain, their errors also belong in the reconstruction model.

The minimization supplies a selection rule. It may prefer one state even when the observations admit several. Regularization, priors, and sparsity assumptions likewise add information to the selection. We should distinguish the part determined by the measurements from the part chosen by these additional assumptions.

With zero-mean errors of known positive-definite covariance Σ, take $Q = \Sigma^{-1}$ as in Section 11.4. If H has full column rank and no inequality or elemental equality is imposed, the generalized least-squares estimate is

$$\hat{n} = (H^T \Sigma^{-1} H)^{-1} H^T \Sigma^{-1} y,$$
$$\operatorname{Cov}(\hat{n}) = (H^T \Sigma^{-1} H)^{-1}.$$

The generalized Gauss–Markov conclusion is that this is the best linear unbiased estimator: every other linear unbiased estimator has covariance larger by a positive-semidefinite matrix (Aitken, 1936). Gaussian errors are not required. Multiplying the observation covariance through the estimator gives the displayed covariance, connecting the residual weights with Section 16.4. Hard constraints or active nonnegativity bounds require the corresponding constrained calculation.

Example 16.2. Use scaled readings and amounts with

$$H = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{pmatrix}, \quad y = \begin{pmatrix} 1 \\ 2 \\ 4 \end{pmatrix},$$

$$\Sigma = \mathrm{diag}(1,1,4).$$

The third reading is less precise. Exact arithmetic gives

$$\hat{n} = \begin{pmatrix} 7/6 \\ 13/6 \end{pmatrix}, \qquad H\hat{n} - y = \begin{pmatrix} 1/6 \\ 1/6 \\ -2/3 \end{pmatrix},$$

$$\mathrm{Cov}(\hat{n}) = \frac{1}{6}\begin{pmatrix} 5 & -1 \\ -1 & 5 \end{pmatrix}.$$

The residual is nonzero because the three readings are inconsistent with an exact two-amount state. Nevertheless, $H^T\Sigma^{-1}(H\hat{n} - y) = 0$, which independently checks the weighted optimum. The estimate is positive, so nonnegativity alone does not change this example. The negative off-diagonal covariance records the dependence introduced by the shared sum reading.

16.4 Covariance transforms with the map

We can also propagate uncertainty through a linear map. If a random amount vector has mean m and covariance Σ, its deterministic linear observation has mean Hm and covariance

$$\Sigma_y = H\Sigma H^T.$$

Expand the covariance of $H(n-m)$ and take the constant matrices outside the expectation to obtain this formula. The cross-correlation terms remain present. Treating all error bars as independent would remove them and may change the result.

Example 16.3. Take independent amount measurements with variances 0.04 and 0.09 in squared amount units. Their sum has variance 0.13. If the covariance is -0.03 instead, the sum has variance $0.04 + 0.09 - 0.06 = 0.07$. Thus the uncertainties of individual coordinates do not by themselves determine the uncertainty of their total.

A singular covariance can also have a meaningful interpretation. Exact conservation confines fluctuations to a lower-dimensional subspace, leaving zero variance in a conserved direction. We should represent that constraint rather than introduce artificial variance merely to obtain an invertible matrix.

16.5 Structural and dynamical inference

Elemental observations leave all states in one fiber of A indistinguishable. Time-series data can help, but their interpretation uses a dynamic model. Different mechanisms may share a stoichiometric subspace and even the same observed net rates. We then need other experiments or conditions to separate them.

Flux balance analysis provides a related example (Orth, Thiele, and Palsson, 2010). It combines

steady-state balances with direction and capacity bounds, often adding an objective. Suppose an input of at most five splits into two nonnegative outputs. Maximizing the first output assigns five to it and zero to the other. Conservation alone allows all intervening splits. We can see directly which part of the answer came from the objective.

Before choosing an objective, we can inspect the elementary generators of the feasible steady-state flux cone from Definition 6.3. An objective selects among their nonnegative combinations; it does not turn those generators into a unique reaction history. Capacity bounds further cut the cone into a constrained feasible region.

Part IV Applications Limits and Development

We can now apply several constructions to the same example and follow how their meanings connect. We begin with ions and catalysts, then examine reaction graphs, exchange between systems, and the interpretation of algebraic models. Finally, we return to the original idea and consider how the representations can be developed further.

Chapter 17 Integrated Chemical Examples

17.1 Neutralization and the charge row

Example 17.1. Let us use species order H^+,OH^-, H_2O and observe hydrogen, oxygen, and charge. We obtain the augmented matrix

$$\tilde{A} = \begin{pmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 1 & -1 & 0 \end{pmatrix}.$$

Multiplying by $(-1, -1, 1)^T$ gives zero. We can inspect the three residuals separately: hydrogen gives $-1 - 1 + 2 = 0$, oxygen gives $0 - 1 + 1 = 0$, and charge gives $-1 + 1 + 0 = 0$.

Notice a dependence between the rows in this restricted example. Subtracting twice the oxygen row from the hydrogen row gives the charge row. Charge therefore adds no rank to this particular matrix. With other species and charges, the same observation can become independent. We determine independence from the selected columns and rows, rather than from the names of the quantities alone.

Suppose the two initial ion amounts are 2 and 3 in a common unit. An extent of two consumes the first ion, leaves one unit of hydroxide, and adds two units of water to the modeled inventory. This is a balance calculation. For an aqueous equilibrium, we would

also need the solvent, activities, water autoionization, and experimental boundary conditions.

17.2 A redox balance

Example 17.2. Consider the following net ionic equation:

$$\mathrm{MnO_4^-} + 5\mathrm{Fe^{2+}} + 8\mathrm{H^+} \rightarrow \mathrm{Mn^{2+}} + 5\mathrm{Fe^{3+}} + 4\mathrm{H_2O}.$$

We count one manganese, five iron, four oxygen, and eight hydrogen atoms on each side. The charge totals are $-1 + 10 + 8 = 17$ and $2 + 15 = 17$. Every augmented residual is therefore zero.

We can also obtain the balance from two half-reactions. The manganese half-reaction consumes five electrons, while each iron half-reaction produces one. Multiply the latter by five and add. The electron coordinate cancels. In vector terms, we have added changes in an enlarged species space and eliminated a bookkeeping species from the net expression.

The cancellation explains the coefficients without identifying a molecular mechanism or a rate. It also shows why a zero elemental vector does not make a species irrelevant: the electron contributes to charge balance even though it contributes to no elemental row. No laboratory procedure follows from this bookkeeping calculation.

17.3 A catalytic mechanism

Example 17.3. Let us examine the model mechanism

$$X + E \rightleftharpoons C, \qquad C \to Y + E,$$

with catalyst E and bound complex C. Assume that X and Y share their elemental composition and that $a(C) = a(X) + a(E)$. With species order X, Y, E, C, we obtain three oriented reaction columns:

$$S = \begin{pmatrix} -1 & 1 & 0 \\ 0 & 0 & 1 \\ -1 & 1 & 1 \\ 1 & -1 & -1 \end{pmatrix}.$$

We can identify two conserved totals:

$$x + y + c = x_0 + y_0 + c_0,$$

$$e + c = e_0 + c_0.$$

Check them by multiplying S by the rows (1,1,0,1) and (0,0,1,1). Although the catalyst cancels from the net conversion $X \to Y$, its free and bound amounts can change. Their distribution remains part of the mechanism and affects its rate.

For a constant-volume concentration model, assign the mass-action rates $k_1 xe$, $k_{-1} c$, and $k_2 c$. The complex then satisfies

$$\dot{c} = k_1 xe - (k_{-1} + k_2)c.$$

Write the conserved catalyst total as $e_T = e + c$. When a quasi-steady approximation is justified, we set $\dot{c} \approx 0$ and use $e = e_T - c$. Solving gives

$$c \approx \frac{e_T x}{K_M + x}, \qquad K_M = \frac{k_{-1} + k_2}{k_1}.$$

The resulting product-formation rate is $k_2 e_T x / (K_M + x)$. We have solved the imposed quasi-steady equation algebraically. The decision to impose it still needs a suitable time-scale separation and initial behavior. Eliminating the complex from the expression retains those conditions as assumptions of the reduced model.

17.4 Reaction energies and linear dependence

Example 17.4. Take standard molar free energies $(3,1,0)^T$ for three model species, all in the same energy-per-amount unit. The changes along $X \to Y$, $Y \to Z$, and $X \to Z$ are -2, -1, and -3, respectively.

The third direction is the sum of the first two. Evaluating the same energy covector on it therefore gives the sum of their changes. The numbers illustrate the linear relation; they are not measured values for named substances.

More generally, suppose $Sc = 0$ and assign reaction energies by $g = S^T \mu^\circ$. We must then have $c^T g = 0$. If a proposed set of energies violates this relation, it cannot come from one species potential under the stated conditions and normalization. This checks

potential consistency. A driven open kinetic cycle requires a separate entropy-production analysis.

For the corresponding reversible reactions, use dimensionless equilibrium constants with

$$\Delta_r G^\circ = -RT\ln K.$$

The same consistency test now becomes

$$\ln K \in \mathrm{im} S^T = (\ker S)^\perp,$$

$$\prod_j K_j^{c_j} = 1 \quad \text{whenever } Sc = 0.$$

These are the Wegscheider conditions (Wegscheider, 1901). For the dependency $(1,1,-1)^T$ of Example 6.1, they give $K_1 K_2 = K_3$. Thus equilibrium constants cannot be assigned independently to dependent reaction directions.

For reversible mass-action pairs written using the same dimensionless activity convention, detailed balance gives $K_j = k_j^+ / k_j^-$. The cycle identity then constrains those rate-constant ratios. Concentration-based constants can carry different units in the forward and reverse directions; their standard-concentration factors must be included before identifying the ratio with a dimensionless thermodynamic constant. This closes the connection between the kernel of S, a species potential, and the kinetic requirement in Section 15.4.

Chapter 18 Reaction Graphs and Composed Models

18.1 Complexes are multisets of species

A reaction side contains a multiset of species. In reaction-network terminology, we call this formal sum a *complex*. For example, $X + E$ and $Y + E$ are complexes even though each names two species. This use of the word differs from the chemically bound species C in the catalytic example.

Definition 18.1. The complex-composition matrix Y_G has one column for each distinct reaction-network complex, expressed in species coordinates. The incidence matrix I_G has a -1 at the source complex and a $+1$ at the target complex of each oriented reaction.

For each column, the product counts minus the reactant counts give

$$S = Y_G I_G.$$

We can read the factorization in two steps. Incidence changes the complex labels, and Y_G translates that change into species counts. This is a standard reaction-network construction (Feinberg, 2019; Baez and Biamonte, 2011).

18.2 Factorizing the catalytic example

Example 18.1. Return to the catalytic mechanism and order its complexes as $X + E$, C, and $Y + E$. For species order X, Y, E, C, the matrices are

$$Y_G = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix},$$

$$I_G = \begin{pmatrix} -1 & 1 & 0 \\ 1 & -1 & -1 \\ 0 & 0 & 1 \end{pmatrix}.$$

Multiplication of $Y_G I_G$ reproduces Example 17.3. The net changes agree, but the factorization retains more information: it identifies the formal complexes joined by each reaction. We cannot always reconstruct that information from S alone.

For example, $X + E \to Y + E$ and $X \to Y$ have the same net vector. In the first expression, the catalyst occurs on both sides and cancels. It still appears in the mass-action input, so the two expressions need not have the same rate law. Net change can hide a participant that matters to dynamics.

18.3 A structural count and its limits

We can associate a structural count with the complex graph. Its *deficiency* is the number of complexes minus the number of connected components of the underlying undirected graph

minus the stoichiometric rank. Example 18.1 gives $3 - 1 - 2 = 0$.

Before applying a theorem to that count, we must check its other hypotheses. The displayed network is not weakly reversible, since no directed path leads from $Y + E$ back to C. For any kinetics whose reaction rates are positive at every strictly positive composition, the non-weakly-reversible branch of the Deficiency Zero Theorem therefore rules out a steady state in which every species concentration is strictly positive (Feinberg, 2019). This conclusion does not locate boundary steady states or calculate a transient trajectory.

We can compare this with a program’s dependency graph. The graph constrains possible interactions but does not determine running time by itself. Similarly, the chemical complex graph acquires behavioral consequences through a kinetic interpretation and its assumptions. The useful analogy concerns the relation between structure and behavior.

18.4 Opening the system

Suppose material can enter or leave the modeled system. We add the external flows explicitly:

$$\dot{n} = S v(n) + B u(t).$$

The columns of B specify exchange directions and $u(t)$ gives their rates. Applying the composition map, with $AS = 0$, leaves

$$\frac{d}{dt}(An) = ABu(t).$$

The internal reactions cancel from the elemental derivative, leaving the boundary contribution. If a reservoir is omitted from the internal species list, a displayed subsystem reaction may appear unbalanced. We can interpret it as exchange only when the boundary convention is explicit. The same expression would be an error if presented as a closed internal balance.

Example 18.2. Let a compartment receive X at rate u and export Y at rate w, with internal conversion of X to Y. If both species have the same elemental composition, their total amount changes at rate $u - w$. At steady state, inflow and outflow match, although internal throughput can continue.

18.5 Composing compartments

We can assemble several compartments using the tensor labels of Chapter 12. Each block has local reactions; transport joins blocks; external exchange crosses the chosen outer boundary. Summing the compartment inventories cancels the internal transport contributions while retaining the external ones.

This suggests a practical checking order. We first verify local reaction balances, then match removal and addition in every transport column. Finally, we identify the boundary terms that remain. The global balance should follow from those local checks. A failure directs attention to an omitted, duplicated, or incorrectly signed contribution.

By analogy with a software interface contract, each component declares what it consumes, produces, and preserves. We can then check a composition of components against those declarations. The analogy concerns this checking method. It does not impose elemental chemistry on software transitions or assert that every chemical process is a finite-state machine.

In the discrete model of Section 7.5, combine places and local and transport transitions. Input and output arcs retain enabling requirements. This correspondence is exact; the software-contract comparison concerns its use.

Chapter 19 Model Meaning and Validation

19.1 Labels and interpretation

Consider receiving a matrix with no row or column labels. We cannot tell which equation it represents. A reproducible model therefore records element and species order, compositions, charge numbers, reaction orientation, and units. If the species list changes, every dependent array must change consistently.

Definition 19.1. A balance certificate consists of the ordered balance matrix, a proposed signed coefficient vector, and its exactly computed residual. A zero residual certifies the encoded linear balance for that vector.

We can verify that the supplied vector balances the supplied matrix. The certificate does not establish that an input formula is correct, that a named species exists, or that the reaction is chemically feasible. Those checks concern the inputs and their interpretation. The situation is similar to implementing a specification correctly while still having to verify the specification itself.

Definition 19.2. An infeasibility certificate for a balance using all displayed species consists of the ordered signed balancing matrix B, a covector y, and

the exactly computed row $y^T B \geq 0$ with at least one positive entry. Proposition 5.2 shows why this rules out $Bc = 0$ with $c > 0$. We verify the signs as well as the multiplication. Like a balance certificate, it concerns the encoded species and side assignments; changing those inputs changes the question.

19.2 Exact arithmetic

An elemental composition matrix records integer counts. Its balance equations can therefore be treated over the rational numbers. A vanishing residual is then an exact identity in the chosen representation. Clearing denominators and removing a common integer factor changes the normalization of a reaction direction without changing its span.

This distinction matters when we interpret a result. Rounding can make a nonzero residual appear small, while an exact zero still depends on the correctness of the species list and its compositions. Exact arithmetic removes one source of uncertainty; it does not supply the chemical interpretation of the input.

The water example in Chapter 4 illustrates the sign convention. In $Av = 0$, the reaction vector carries the signs. In $Bc = 0$, the selected sides are already included in the balancing matrix. Mixing the two conventions would duplicate or reverse signs. The same chemical statement can have either

representation, provided its meaning remains explicit.

For all integer solutions, use the lattice methods of Section 5.5. Clearing the denominators of a rational nullspace basis may generate only a proper sublattice. Smith or Hermite methods retain divisibility information; real coordinates lose the torsion recorded by an integer cokernel.

19.3 More than one solution

A kernel basis describes all solutions of the linear balance equations. If there is more than one independent direction, its basis vectors need not individually satisfy the desired reaction signs. Selecting a feasible combination is a further problem. The distinction belongs to the model itself: conservation specifies a space of possibilities, while additional conditions select among them.

In the carbon monoxide/carbon dioxide example, we can add a product ratio, impose bounds, or restrict the admissible integer coefficients. Each option supplies an extra selection rule. We should keep that rule separate from the elemental equations so that its effect can be inspected and changed.

We also select arithmetic according to the input. Integer compositions support exact elimination. An empirical response matrix may be noisy or nearly rank-deficient, requiring tolerances and uncertainty analysis. A numerically small residual should be

reported with that qualification, rather than presented as an exact conservation identity.

19.4 Validation by independent checks

After obtaining a vector, we can check it by direct multiplication. This tests the result without repeating the elimination that produced it. For a reaction network, calculate $AS = 0$ and each claimed conserved relation $\ell^T S = 0$. For an extent example, independently compute the endpoint and its elemental totals.

The same checking should include units. A rate in moles per second becomes a concentration rate only after volume is accounted for. If volume changes, differentiating $x = n/V$ introduces a dilution term as well as the amount-rate term. An interface that labels a quantity only as “rate” leaves this distinction unresolved.

Example 19.1. Starting with $n = Vx$ and differentiating, we obtain

$$\dot{x} = \frac{\dot{n}}{V} - \frac{x\dot{V}}{V}.$$

Thus even when $\dot{n} = Sv$ is supplied, volume change contributes separately. Concentration can change while the elemental amount remains conserved. We should check which quantity was measured before interpreting such a change as a balance failure.

19.5 The organization of a model

A model connects its vocabulary with its mathematical structure. Species identities, formulas, phase and isotope labels, and units determine what its coordinates mean. The maps A and S describe composition and selected changes, while observation and exchange maps relate the modeled system to its surroundings. Constraints then specify which algebraic possibilities are admitted.

We can follow a reported result back through these relationships. A residual concerns particular equations. Those equations concern particular species, with a declared orientation and choice of units. A balanced chemical string compresses this structure, but examining the underlying maps reveals which statement has actually been established.

Adding charge, isotope distinctions, compartments, or uncertainty changes some part of the representation. We should be able to identify that change and follow its consequences. An extra observation may distinguish states that were previously indistinguishable; a new constraint may remove an otherwise balanced direction. The explanation of the result lies in these relationships.

Chapter 20 Scope Originality and Further Directions

20.1 What has been developed

I began with elements as basis vectors and stoichiometric coefficients as scalars. We have now placed that correspondence in a counting monoid and a larger vector space, while introducing a separate basis for species. In species space we distinguish ethanol from dimethyl ether. The composition map assigns both the same vector.

Each extension answered a question raised by the preceding representation. We introduced reaction coordinates for selected changes and dual vectors for additive observations. Tensor products retained species together with context. Alternating forms tested independent directions, and differential forms evaluated changes along paths. The contribution is to show which chemical distinctions each map preserves and when we need a richer description.

Aris (1965) used distinct spaces for species and reactions, an atomic matrix, and a matrix of stoichiometric coefficients. My 2009 announcement records how I came to the elemental representation at school (Vostokov, 2009). The dates 1986–1987 describe the origin of this book. The mathematical methods have an earlier history.

20.2 Exact correspondences and useful analogies

Some correspondences are exact once the representation is fixed. A formula's atom counts give vector coordinates. A selected atom-conserving reaction belongs to the composition kernel. A linear conserved observation annihilates the stoichiometric map. We can prove these statements from the definitions and verify them in the examples.

Other connections are analogical. We compared a balance certificate with a software invariant check and a reaction interface with a component contract. We can also view a thermodynamic covector as a measurement of a state displacement. Such comparisons suggest questions and constructions. To use a law from either setting, we still have to establish its conditions in the new one.

The examples also showed where a proposed identification fails. Tensor products did not supply chemical bonds, and wedge products did not count repeated atoms. A nullspace contained no reaction mechanism. Coordinate distance alone gave no guarantee of chemical similarity. These failures help us choose the next representation rather than ask the existing one for information it does not contain.

20.3 Enriching composition systematically

One direction is to organize richer descriptions and their projections. We can retain isotope counts, species identity, molecular graphs, stereochemical labels, and environmental or spatial information as needed. Each map should identify the distinctions it removes. Calculations could then move between representations while recording what remains observable.

We can also compare selected reaction directions with all balanced ones. When $AS = 0$, the quotient $\ker A/\mathrm{im} S$ describes balanced directions modulo those generated by the reaction subspace. It exposes an algebraic limitation of the selected mechanism. Whether an omitted direction can be realized chemically is a further question.

The sequence $R \xrightarrow{S} W \xrightarrow{A} E$ has zero composite and therefore the algebraic structure of a chain complex. We can develop homological language from that observation. A spatial or topological meaning would require additional construction; the elemental map has not thereby become an ordinary geometric boundary operator.

Example 20.1. Add hydrogen peroxide to the water system, using species order $\mathrm{H_2}, \mathrm{O_2}, \mathrm{H_2O}, \mathrm{H_2O_2}$, but retain water formation as the only selected reaction. Then

$$A = \begin{pmatrix} 2 & 0 & 2 & 2 \\ 0 & 2 & 1 & 2 \end{pmatrix},$$

$$v = (-2,-1,2,0)^T, \qquad w = (-1,-1,0,1)^T.$$

Direct multiplication gives $Av = Aw = 0$. The rank of A is two, so v, w form a basis of its real kernel. With S consisting only of column v, the quotient $\ker A/\mathrm{im} S$ has dimension one, represented by the class of w. It records the missing balanced peroxide direction without asserting that this direction is an available chemical step.

Over integers, v, w also generate the whole kernel: the oxygen equation forces the water coefficient to be even, after which both remaining coefficients are determined. Thus $\ker_{\mathbb{Z}} A/\mathrm{im}_{\mathbb{Z}} S \cong \mathbb{Z}$ here. If a discrete model instead admits only the doubled event $2v$, the quotient becomes

$$\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}.$$

The class of v is nonzero but twice it vanishes: torsion in homology at species counts. The cokernel of Section 5.5 is the homology at elemental counts when the next map is zero. Rescaling preserves a reaction's real direction (Section 7.3), but can change its integer image. These quotients alone do not decide nonnegative reachability.

20.4 Constrained learning and observation design

For a learned kinetic model, we can constrain outputs to ker$A$ and obtain encoded conservation by construction. Factoring through S restricts them to selected reactions; nonnegative rates with suitable boundary behavior can impose further conditions. We can then test on independent data whether these restrictions improve prediction, generalization, or the detection of inconsistent inputs.

Observation design gives another direction. Suppose two candidate species or mechanisms remain indistinguishable. We can inspect the kernel of the present observation map and ask which additional readout removes the relevant ambiguity. The admissible states and uncertainty also matter. An ideal row of numbers must eventually be connected with an assay that is physically realizable and affordable.

Compartment-resolved and multiway measurements provide opportunities for tensor models. We can test whether retaining associations changes the answer, whether a low-rank approximation preserves nonnegativity and conservation, and whether its factors have a supported chemical interpretation. Fitting the observations is one check among these.

20.5 Further investigation

The same chemical system can be examined at successive levels of representation. Elemental counts constrain its composition, species labels retain identity, and kinetic assumptions specify possible evolution. Comparing these levels can reveal which conclusions survive a change of representation and which require information that the earlier description omitted. Counterexamples are useful here because they locate the limits of a proposed correspondence.

To investigate these extensions, we need test cases with declared inputs, assumptions, and expected properties. Exact balancing, constrained reconstruction, kinetic fitting with conservation, and compartment aggregation can be compared for correctness, sensitivity to missing species, interpretability, and reproducibility. The small cases in this book supply starting checks. They do not constitute measured evidence of improved performance.

Elemental basis vectors and scalar coefficients remain the starting representation. The maps around them preserve selected information and expose what is missing. Further development begins with a question that this representation cannot answer and constructs the additional structure it needs.

Appendix A Notation and Terminology

A.1 Spaces and structural maps

We use n species, r reactions and stoichiometric rank s. In Section 5.5, ρ denotes the composition-map rank. These conventions distinguish the main constructions; an explicit local definition takes precedence when a symbol is reused.

E is the elemental composition space. Its distinguished basis vectors e_a represent the selected elements, not atoms at particular positions. A composition vector $v \in E$ records elemental multiplicities or amounts, according to the declared coefficient domain and units. The notation $a(X)$ denotes the composition of species X. In Chapters 17–18, the locally defined species E is the catalyst, rather than the elemental space.

W is the species space. Its distinguished basis vectors u_i represent the species in a model. Two basis vectors can have the same elemental composition. The species list can distinguish isomers, charge states, phases, or isotopic labels when the problem requires them.

R is the reaction-coordinate space. A basis vector in R names one oriented reaction channel. Its coefficient is an extent, an extent rate, or an abstract

coordinate, as specified in context. The letter $\mathbb{R}$, with a different typeface, denotes the real numbers.

C is a compartment space in Chapter 12. Its basis labels locations or compartments; it is unrelated to the chemical symbol C for carbon. The identity map on a space is written I with a subscript when useful. In Chapters 17–18, the species C is the bound complex; this is separate from the compartment space and from elemental carbon.

$A: W \to E$ is the composition map. Column i of its matrix is the elemental composition of species i. $S: R \to W$ is the stoichiometric map. Its columns are signed reaction vectors, negative for consumed species and positive for produced species. The central consistency condition for a closed, element-balanced reaction model is

$$AS = 0.$$

$\tilde{A}$ denotes an augmented composition matrix when charge or other additive labels are appended. Such a row is an additional constraint only when it is linearly independent of the existing rows for the chosen species list.

Y_G maps reaction-complex coordinates into species coordinates. Here a complex is a multiset such as $X + E$, not necessarily a bound chemical complex. I_G denotes a directed graph incidence matrix when it must be distinguished from an identity matrix. The factorization is $S = Y_G I_G$. This matrix is distinct from

the formal species Y and from the observation codomain Y in Definition 9.1.

B denotes the signed balancing matrix in Chapter 5 and the exchange matrix in Section 18.4. In Chapter 13 it denotes a bilinear map; the ordered elemental basis in Section 1.1 uses a calligraphic basis symbol. These are locally defined objects with different domains.

A.2 States rates and measurements

n is a vector of species amounts. x is a vector of concentrations when a kinetic model uses concentration coordinates. N is used for discrete molecule counts or a species-by-compartment array; its meaning is stated locally. None of these quantities should be substituted for another without the appropriate conversion.

ξ is a vector of reaction extents. With a fixed stoichiometric normalization, an amount change is $S\xi$. A rate vector $v(x)$ supplies reaction rates, so that the concentration equation $\dot{x} = Sv(x)$ applies under the fixed-volume convention used in Chapter 10. The amount equation instead uses extent rates.

H denotes an observation matrix. y is an observed or predicted measurement vector and ε is a measurement error. A Hessian is distinguished as H_f or described explicitly. A covariance matrix is written Σ.

G can denote a Gram matrix in a discussion of bilinear forms or Gibbs energy in a thermodynamic discussion. The argument or surrounding definition distinguishes them. μ is the vector of chemical potentials; $\Delta_r G = S^T \mu$ is the vector of reaction Gibbs energies. Affinity is $\mathcal{A} = -S^T \mu$ under the sign convention of Chapter 15.

c denotes the balancing coefficient vector in Chapter 5, the conserved total concentration in Section 10.2, the bound-complex concentration in Section 17.3, or a reaction-dependency vector in Section 17.4. Subscripted compartment basis vectors in Chapter 12 use the same base letter. The local definition distinguishes these uses.

A.3 Algebraic operations

$\ker T$ is the set of vectors that a linear map T sends to zero. $\operatorname{im} T$ is the set of its outputs. The rank of a map is the dimension of its image. A superscript T denotes matrix transpose, not time; time is the variable t.

V^* is the dual space of linear real-valued functionals on V. Pairing a covector with a vector produces a scalar. For a linear map $T: V \to U$, the dual map $T^*: U^* \to V^*$ acts by composition. In corresponding coordinate bases its matrix is the transpose of the matrix of T.

$\otimes$ denotes a tensor product. It combines independent types of labels or assembles multilinear

constructions. ∧ denotes an exterior product, which changes sign when adjacent vector factors are exchanged. Neither symbol, by itself, describes a chemical bond.

dF is the differential of a scalar function F. A differential form ω assigns alternating multilinear data to tangent directions at each state. The exterior derivative is $d\omega$, and the pullback of a form through a map f is $f^{*}\omega$.

A.4 Units and coefficient domains

We count atoms per formula unit with dimensionless nonnegative integers. Molecule counts are also nonnegative integers. For macroscopic quantities, we use amounts in moles or concentrations obtained by dividing amount by volume. Stoichiometric numbers are dimensionless, while a molar extent has units of amount. Rescaling a reaction equation changes the numerical extent for the same physical change.

For example, consider kx^{p} in a concentration-rate equation. For a total order p, the units of k are concentration to the power $1-p$ divided by time. Before transferring a rate constant between models, we should therefore compare their conventions and units.

We use $\mathbb{N}^{m}$, $\mathbb{Z}^{m}$, $\mathbb{Q}^{m}$, and $\mathbb{R}^{m}$ for different stages of a calculation, with $\mathbb{N}$ including zero. Nonnegative integers describe finite counts, integers allow signed

differences, rationals support exact balancing, and reals support continuous amounts and dynamics. Enlarging the domain supplies new mathematical operations. We still check which resulting vectors can represent physical states.

Dimensional analysis uses a related linear map. Put the exponents of the base units of each quantity into a column of a dimension matrix D_u. A monomial $\prod_j q_j^{a_j}$ is dimensionless exactly when $D_u a = 0$. For rational unit exponents, we can perform the calculation over the rational numbers. With k quantities and rank r, there are $k - r$ independent exponent directions, giving the dimensionless groups of the Buckingham π theorem under its dimensional-homogeneity assumptions (Buckingham, 1914). For length L, time t, and speed u, the columns $(1,0)^T, (0,1)^T, (1,-1)^T$ give the group ut/L. As with balancing, the kernel supplies admissible combinations; it does not supply the empirical relation between them.

A.5 A compact glossary

Affine space. A translate of a vector subspace. A compatibility class has an affine part because it is anchored at an initial inventory.

Alternating form. A multilinear scalar-valued function that vanishes when two arguments coincide. Interchanging arguments reverses its sign

over any field. The converse requires characteristic different from two, as in the real case.

Basis. An independent spanning family. The elemental basis is selected to record elemental totals; it is not deduced from a metric.

Bilinear form. A scalar-valued function linear in each of two vector arguments separately.

Chemical species. A distinct entity in the model's vocabulary. Its identity can include more information than its elemental formula.

Cokernel over integers. The quotient of an integer target group by the image of an integer map. Its torsion can record divisibility restrictions invisible over the rationals.

Complex in a reaction graph. A multiset of species appearing on one side of a reaction. This graph-theoretic usage differs from a coordinated or bound chemical complex.

Composition map. The linear map that forgets species identity while retaining elemental totals.

Conservation covector. A linear functional ℓ satisfying $\ell^T S = 0$. In a closed model it is constant along every modeled reaction trajectory.

Contraction. Pairing a tensor index with a compatible dual index and summing. Aggregation over compartments is an example.

Deficiency. For a reaction graph, the number of complexes minus the number of linkage classes minus the stoichiometric rank. Its use in theorems requires their other hypotheses.

Elementary flux mode. A nonzero feasible steady flux with inclusionwise minimal support. For all-irreversible coordinates, these modes correspond to extreme rays of the flux cone.

Exact form. A differential form that is the exterior derivative of another form. For a one-form, exactness means it is dF for a scalar potential.

Extent. A coordinate measuring progress along an oriented, normalized reaction vector.

Fiber. All inputs that a map sends to one specified output. A composition fiber groups species inventories with equal elemental totals.

Identifiability. Uniqueness of an inferred quantity under the stated model, observations, and admissible domain.

Inner product. A symmetric positive-definite bilinear form over a real vector space. A basis alone does not select one.

Kernel. The inputs invisible to a linear map. Its interpretation depends on the map: composition, reaction, and observation kernels answer different questions.

Lattice. An integer linear structure inside a real vector space. Integer stoichiometric changes belong to such a structure after a species basis is fixed.

Mass action. A kinetic assumption under which an elementary reaction rate has a specified product dependence on reactant concentrations or, in a stochastic model, a specified dependence on available molecular combinations.

Multiset. A collection with multiplicities. A chemical formula becomes a multiset of element labels after structural information is forgotten.

Petri net. A discrete model with places, transitions, and specified nonnegative integer input and output counts. A transition can fire only when its input tokens are present.

Positive cone. The set of nonnegative linear combinations of a chosen family. The nonnegative orthant is tied to the physical species basis.

Quadratic form. A homogeneous degree-two scalar function obtained by evaluating a bilinear form twice on the same vector.

Quotient space. A space whose vectors are equivalence classes modulo a subspace. Quotienting by a kernel identifies inputs that a map cannot distinguish.

Rank. The number of independent output directions of a linear map. It is not a count of independent reaction mechanisms established by experiment.

Stoichiometric subspace. The image of S, containing every modeled reaction displacement.

Tensor. An element of a tensor product, or equivalently an appropriate multilinear construction in finite dimensions. An array represents a tensor only after its spaces and transformation rules are specified.

Torsion. A nonzero element of an additive quotient group that becomes zero after multiplication by some positive integer.

Weak reversibility. The property that every directed reaction in the complex graph lies on a directed cycle. It does not require each reaction to have an explicit reverse edge.

Appendix B Mathematical Background

B.1 Linear maps before matrices

We can construct a matrix by recording the action of a map on basis vectors. Suppose $T: V \to U$ is linear and $v_1, \ldots, v_n$ is a basis of V. Every input has a unique expression $v = \sum_i a_i\, v_i$. Applying linearity gives

$$T(v) = \sum_i a_i\, T(v_i).$$

The basis images determine every other output. After choosing an output basis, we place the coordinates of $T(v_i)$ in column i. This is the reason for arranging species formulas as columns in the composition matrix.

Now consider two successive maps. First S converts reaction coordinates into species changes; then A converts those changes into elemental changes. Their combined matrix is AS. The relation $AS = 0$ says that every selected reaction change disappears at the elemental level. Since this concerns the maps themselves, it survives compatible coordinate changes.

We cannot generally reverse the order. The expression SA usually has incompatible source and target spaces and is undefined. Before multiplying arrays, we should identify the domain and codomain

of each map and check that the proposed composition makes sense.

B.2 Row reduction with an explicit kernel

Let us calculate the kernel of the water composition matrix, with species ordered as hydrogen gas, oxygen gas, and water:

$$A = \begin{pmatrix} 2 & 0 & 2 \\ 0 & 2 & 1 \end{pmatrix}.$$

Writing $Av = 0$ gives $2v_1 + 2v_3 = 0$ and $2v_2 + v_3 = 0$. We can divide each row by its leading coefficient to obtain

$$\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1/2 \end{pmatrix}.$$

The first two columns contain pivots. We can choose the third coordinate freely. Put $v_3 = t$; then $v_1 = -t$ and $v_2 = -t/2$. The full kernel is therefore

$$\ker A = \{t(-1, -1/2, 1)^T : t \in \mathbb{R}\}.$$

Multiplying the rational generator by two gives $(-2, -1, 2)^T$. Its negative is another primitive integer generator. We choose between their orientations according to the intended reactant and product sides. The kernel calculation itself does not establish which physical reaction occurs.

Each elementary row operation is invertible, so it preserves the solutions of the homogeneous system. We can rescale a row by a nonzero number, add a

multiple of another row, or exchange rows. Small pivots need care in floating-point arithmetic. Exact rational arithmetic avoids that rounding issue for these small integer examples, although intermediate fractions may become large in bigger calculations.

B.3 Why rank and nullity add up

To understand the dimension count, begin with a basis $k_1, \ldots, k_d$ of $\ker T$. Add vectors $v_1, \ldots, v_r$ until we have a basis of the whole domain. The outputs $T(v_1), \ldots, T(v_r)$ span the image: the kernel part contributes zero to every output.

We can also check their independence. If $\sum_j a_j\, T(v_j) = 0$, then $\sum_j a_j\, v_j$ belongs to the kernel and can be expressed using the k_i. Since the extended domain family is independent, every a_j must vanish. The added vectors therefore count the rank:

$$\dim V = \dim\ker T + \operatorname{rank} T.$$

For a matrix with m rows, n columns, and rank r, the right-nullspace dimension is $n - r$ and the left-nullspace dimension is $m - r$. We then attach the interpretation appropriate to the map. For S, right-nullspace vectors describe dependent reaction changes, while left-nullspace vectors describe conserved species combinations.

B.4 Fibers quotients and feasible slices

Suppose $Tw = b$ has a known solution w_0. Subtracting that solution from another gives $T(w - w_0) = b - b = 0$, so the difference lies in the kernel. Conversely, adding a kernel vector to w_0 leaves the output fixed. We have therefore described the whole fiber:

$$T^{-1}(b) = w_0 + \ker T.$$

This set is affine. It need not include zero, and adding two of its points generally doubles their common output. To obtain physically admissible solutions, we intersect it with the required nonnegative cone, integer lattice, or other constraints. The intersection may be empty and need not be a vector space.

The quotient $V/\ker T$ replaces each fiber by one equivalence class. The rule $[v] \mapsto T(v)$ is well-defined, linear, one-to-one, and onto $\mathrm{im}T$. This gives the isomorphism used in Chapter 4. We have matched classes with outputs, rather than recovered the individual inputs that the map merged.

For example, an elemental measurement may identify a class of mixtures without identifying one mixture. The quotient describes the information that remains. To distinguish members of the class, we need a further observation, a justified restriction on admissible states, or both.

B.5 Dual maps and change of basis

A covector gives a linear numerical readout. For vector coordinates v and covector column ℓ, the value is $\ell^T v$. Substitute the coordinate change $v = Pv'$ to obtain $\ell^T Pv'$. To preserve that readout in the new coordinates, we use $\ell' = P^T \ell$.

The vector uses the inverse coordinate change, while the covector uses the transpose. These transformations preserve their pairing. Mass weights and species amounts consequently do not follow the same coordinate rule. The transpose in $A^T m$ is an instance of this general dual construction.

We can check a bilinear form in the same way. Start with $b(v, w) = v^T G w$ and substitute $v = Pv'$ and $w = Pw'$. The new matrix is $G' = P^T GP$. Replacing it by the identity would generally choose another form. Unless the application provides a preferred metric, orthogonality, distance, and angle remain additional choices.

B.6 Tensor products in coordinates

Take bases $u_1, \ldots, u_n$ and $c_1, \ldots, c_q$. Their pairs $u_i \otimes c_\alpha$ form a basis of $W \otimes C$. We can store the general coefficients $N_{i\alpha}$ in an n-by-q array. A pure tensor has the special factorization $a_i b_\alpha$; the general array need not have it. Thus an array representation does not imply separability of its data.

The universal property makes this construction useful. A bilinear rule $B: W \times C \to U$ determines a unique linear rule on $W \otimes C$, taking $w \otimes c$ to $B(w, c)$. We can therefore replace two independently linear arguments by one argument in a larger space. The representation changes without adding a physical assumption.

For compartment aggregation, let $\sigma \in C^*$ sum compartment coordinates. Then $I_W \otimes \sigma$ contracts that factor. If A converts species to elements, $A \otimes I_C$ acts on the species factor instead. We can apply these operations in either compatible order, as checked in Chapter 12.

B.7 Differentiation and pullback

Consider a differentiable scalar function F of coordinates $x_1, \dots, x_n$. At a chosen point, its differential gives the linear part of the change:

$$dF = \sum_i \frac{\partial F}{\partial x_i}\, dx_i.$$

The differential is a covector there. A curve $x(t)$ supplies the tangent $\dot{x}(t)$, and pairing them gives the chain rule, $dF(\dot{x}) = dF\big(x(t)\big)/dt$. A general one-form also reads tangent vectors, but a global scalar potential need not exist for it.

For the reaction parametrization $n(\xi) = n_0 + S\xi$, the derivative is S. Substitution of $dn = S\, d\xi$ into $\mu^T dn$ yields $(S^T\mu)^T d\xi$, the pullback used in Chapter 15.

The chemical potentials μ may vary with state. The transpose still arises from the derivative of the coordinate map.

With twice continuous differentiability, equal mixed partial derivatives give $d(dF) = 0$. Hence exact forms are closed. To infer exactness from closedness, we also inspect the domain. An open convex domain is sufficient; holes can obstruct a global potential. A chemical use of the converse should therefore identify the state region involved.

B.8 A practical order of calculation

We begin by naming the objects represented and choosing their bases and coefficient domains. We also attach units to amounts, rates, and observations. A map can then be built by specifying the image of each basis vector. Before composing maps, we compare their source and target spaces.

Next we choose the calculation that answers the question. A kernel can identify balances, a left kernel can identify conserved totals, and a rank can count distinguishable directions. A pullback expresses an observation through another map. We then impose physical admissibility and identify any missing information, such as structure, kinetic parameters, thermodynamic data, or calibration.

This order allows us to inspect where each conclusion enters. The algebra follows from the represented objects and maps, while additional

chemical assumptions remain visible as separate inputs.

Bibliography

The following sources provide the mathematical foundations, chemical conventions, reaction-network theory, and related mathematical constructions used in the book. Vostokov (2009) provides the earlier public account of this project. Web references were checked on 22 October 2026 (Version 7), and later additions on 5 November (Versions 8 and 9) and 7 November (Aris, Version 13).

Aitken, A. C. 1936. “On least squares and linear combination of observations.” *Proceedings of the Royal Society of Edinburgh* 55: 42–48. A foundation for generalized least squares with correlated observations. doi:10.1017/S0370164600014346.

Aris, Rutherford. 1965. “Prolegomena to the rational analysis of systems of chemical reactions.” *Archive for Rational Mechanics and Analysis* 19: 81–99. An early vector-space treatment of chemical species and reactions. doi:10.1007/BF00282276.

Axler, Sheldon. 2024. *Linear Algebra Done Right*. 4th edition. Springer. The chapters on linear maps, duality, and multilinear algebra provide a mathematical foundation for the constructions used here. doi:10.1007/978-3-031-41026-0.

Baez, John, and Jacob Biamonte. 2011. “Network Theory Part 17.” 12 November. An introduction to reaction complexes, stoichiometric subspaces, and

structural properties of reaction networks. Author's exposition.

Birch, M. W. 1963. "Maximum likelihood in three-way contingency tables." *Journal of the Royal Statistical Society: Series B (Methodological)* 25 (1): 220–233. Original statistical source for the intersection theorem used in Section 15.4. doi:10.1111/j.2517-6161.1963.tb00504.x.

Buckingham, Edgar. 1914. "On physically similar systems; illustrations of the use of dimensional equations." *Physical Review* 4: 345–376. The dimensional-analysis construction used in Appendix A. doi:10.1103/PhysRev.4.345.

Feinberg, Martin. 2019. *Foundations of Chemical Reaction Network Theory*. Applied Mathematical Sciences 202. Springer. A systematic treatment of the network framework introduced in Chapters 6 and 18. doi:10.1007/978-3-030-03858-8.

Flowers, Paul, Klaus Theopold, Richard Langley, and William R. Robinson. 2019. *Chemistry 2e*. OpenStax, Rice University. A general chemistry reference for formulas, stoichiometry, and chemical quantities. Open textbook.

Gillespie, Daniel T. 1977. "Exact stochastic simulation of coupled chemical reactions." *The Journal of Physical Chemistry* 81 (25): 2340–2361. Foundational work on discrete stochastic reaction events. doi:10.1021/j100540a008.

Henriksson, Oskar, Carlos Améndola, Jose Israel Rodriguez, and Polly Y. Yu. 2025. "Maximum

likelihood estimation of log-affine models using detailed-balanced reaction networks." *Journal of Mathematical Biology* 91: 34. Relates the equilibrium intersection geometry to Birch's theorem. doi:10.1007/s00285-025-02262-5.

Horn, F., and R. Jackson. 1972. "General mass action kinetics." *Archive for Rational Mechanics and Analysis* 47: 81–116. An early general treatment of mass-action systems and their thermodynamic structure. doi:10.1007/BF00251225.

IUPAC. 2023. *Quantities, Units and Symbols in Physical Chemistry*. 4th edition, abridged version. Royal Society of Chemistry. The Green Book provides conventions for physical quantities and units. doi:10.1039/9781839163180.

Kolda, Tamara G., and Brett W. Bader. 2009. "Tensor decompositions and applications." *SIAM Review* 51 (3): 455–500. A reference for multiway arrays, tensor representations, and decomposition methods. doi:10.1137/07070111X.

Murata, Tadao. 1989. "Petri nets: Properties, analysis and applications." *Proceedings of the IEEE* 77 (4): 541–580. A reference for markings, firing rules, invariants, and reachability. doi:10.1109/5.24143.

Orth, Jeffrey D., Ines Thiele, and Bernhard Ø. Palsson. 2010. "What is flux balance analysis?" *Nature Biotechnology* 28: 245–248. An introduction to constrained steady-state flux calculations. doi:10.1038/nbt.1614.

Rao, Riccardo, and Massimiliano Esposito. 2016. “Nonequilibrium thermodynamics of chemical reaction networks: Wisdom from stochastic thermodynamics.” *Physical Review X* 6: 041064. A treatment of thermodynamic structure in chemical reaction networks. doi:10.1103/PhysRevX.6.041064.

Roels, J. A. 1980. “Application of macroscopic principles to microbial metabolism.” *Biotechnology and Bioengineering* 22 (12): 2457–2514. Elemental balances and degree-of-reduction conventions. doi:10.1002/bit.260221202.

Roels, J. A. 1983. *Energetics and Kinetics in Biotechnology*. Elsevier Biomedical Press, Amsterdam. ISBN 978-0-444-80442-6. A broader treatment of biochemical balances and energetics.

Schuster, Stefan, and Claus Hilgetag. 1994. “On elementary flux modes in biochemical reaction systems at steady state.” *Journal of Biological Systems* 2 (2): 165–182. Introduces elementary modes and their relation to sign-constrained steady fluxes. doi:10.1142/S0218339094000131.

Sjamaar, Reyer. 2017. *Manifolds and Differential Forms*. Lecture notes, Cornell University, version December 2017. A mathematical reference for the differential-form language introduced in Chapter 15. Author’s notes.

Smith, Henry J. Stephen. 1861. “On systems of linear indeterminate equations and congruences.” *Philosophical Transactions of the Royal Society of*

London 151: 293–326. Integer linear equations and the normal-form theory used in Section 5.5. doi:10.1098/rstl.1861.0016.

Stiemke, Erich. 1915. “Über positive Lösungen homogener linearer Gleichungen.” *Mathematische Annalen* 76: 340–342. The theorem of alternatives used for balancing infeasibility. doi:10.1007/BF01458147.

Strang, Gilbert. 2010. *Linear Algebra*. MIT course 18.06, Spring 2010. MIT OpenCourseWare. Lectures and materials on matrix spaces, eigenvalues, orthogonality, and positive-definite matrices. Course materials.

Vostokov, Dmitry. 2009. “Vector Space Chemistry.” *Software Diagnostics Library*, 23 January. Original announcement of this book and account of its starting idea. Original post.

Vostokov, Dmitry. 2026. *Applied Set Theory: Foundations and Practice for Computing, Data, and Engineering*. Version 12, 12 September. Author’s manuscript. A companion treatment of sets, relations, functions, and multisets, and the structural model for this book.

Wegscheider, Rudolf. 1901. “Über simultane Gleichgewichte und die Beziehungen zwischen Thermodynamik und Reactionskinetik homogener Systeme.” *Monatshefte für Chemie* 22: 849–906. Equilibrium and kinetic consistency relations. doi:10.1007/BF01517498.

Index

Page references identify the principal discussions.

affine spaces 43, 129
affinity 85, 87
alternating forms 78, 81
amounts of substance 43, 120
atom counts 15, 20
balancing equations 31, 32
basis changes 53, 130
bilinear forms 62, 63
Birch theorem 86
catalytic mechanism 97, 101
chain complexes 112
charge 50, 95
chemical potentials 84, 85
closed forms 83, 131
coarse representations 25, 54
compartments 68, 103
compatibility classes 43
complexes in reaction graphs 100, 101
composition map 27, 28
concentration and volume 108, 120
conservation laws 39, 75
constrained learning 114
contraction 70, 73
coordinate dependence 53, 130
covariance 91
cross product 79
cycles 39, 98

deficiency 101
degree of reduction 50
differential forms 82, 83, 131
dimensional analysis 121
dissipation 87
duality 48, 130
electroneutrality 52
elemental basis 15, 19
elementary flux modes 41, 92
empirical formulas 24
energies of reactions 85, 98
equilibrium in log coordinates 86
ethane combustion 32
exact arithmetic 106
exact forms 83, 131
extents 43, 44, 45
exterior algebra 78, 79
Faraday constant 51
feasibility 21, 43
fibers 29, 129
flux balance analysis 92
flux cones 40, 92
formula units 16, 24
free vector spaces 19
Gauss–Markov estimator 90
Gibbs energy 84, 85
graph incidence 100, 101
Hessians 65
identifiability 88, 92
image of a map 28, 128
infeasibility certificates 32, 105
inner products 54, 62

integer kernels and images 35, 107
integer normalization 32
inverse problems 88, 89
isomers 18, 29
isotopes 25, 112
Jacobians 59
kernels 28, 40, 127
kinetic closure 55, 76
lattices 20, 35
left nullspaces 39, 128
limiting reagents 43
linearization 59
mass action 57, 74
mass functionals 48, 49
measurement operators 54, 88
metrics 54, 66
molecular identity 25, 27
multisets 23, 100
neutralization 95
nonnegative cones 20, 21
observation design 114
open systems 102
original formulation 15, 110
path dependence 83
Petri nets 46
polynomial kinetics 73, 75
positive definiteness 62, 63
projections 54
pullbacks 49, 84, 131
quadratic forms 63, 64
quasi-steady approximation 98
quotient spaces 29, 112, 129

rank and nullity 39, 128
reachability 45, 46
reaction graphs 100, 101
reaction spaces 38
redox balance 96
row reduction 127
Smith normal form 35
species basis 27
state functions 83, 98
Stiemke alternative 32
stochastic events 59, 75
stoichiometric map 38
stoichiometric subspace 38, 43
tensor products 68, 69, 130
tensor rate coefficients 73, 75
thermodynamics 64, 84, 87
torsion 36, 113
transport 70, 103
uncertainty 66, 91
units 108, 120
validation by independent checks 108
water balance 17, 31, 127
weak reversibility 101
wedge product 78, 79
Wegscheider conditions 99