

A TEMPORAL–STRUCTURAL FRAMEWORK
ANALOGOUS TO GENETIC EVOLUTION

Linguistics in the Time Cycle

Moninder Singh Modgil
Dnyandeo Dattatray Patil

Cosmos Research Labs
Centre for Ontological Science, Meta Quanta Physics and Omega Singularity

Dedication

This work is humbly dedicated to

Shiv Baba

the Supreme Soul — the eternal Point of Light,
the Ocean of Knowledge and Peace —
from whom every word proceeds
and to whom it returns.

Om Shanti

Preface

Time is the moving image of eternity; and language, its moving voice.

This book grew out of a short paper, *Linguistics in the Time Cycle*, in which we sketched a single conjecture: that language is not merely something that *changes* in time but something that is *organised by* time, across many nested scales at once. The paper proposed an analogy—language as a temporally regulated system, parallel in its architecture to the way genetic information is expressed, regulated, and inherited through biological time. What began as a set of tables and suggestive parallels seemed to invite a fuller treatment. The present monograph is that treatment.

The argument is interdisciplinary by necessity. To speak of language in time is to borrow at once from chronobiology and phonology, from population genetics and historical linguistics, from Bakhtin's poetics of the chronotope and from the mathematics of dynamical systems. We have tried throughout to keep the analogy disciplined: an analogy is a scaffold for hypotheses, not a proof, and the value of the LTC framework lies in the concrete, testable questions it makes it natural to ask. Each chapter therefore closes by pointing toward empirical or computational work that could confirm or refute the picture it paints.

The book is organised into five parts. Part I lays the conceptual foundations: the central thesis, the genetic analogy, and the architecture of nested time cycles. Parts II, III, and IV follow language down through the three temporal scales—the micro-cycles of cognition and circadian rhythm, the meso-cycles of the lifespan and social time, and the macro-cycles of history and culture. Part V turns to method and to the future: experimental designs, computational models, and the shape of a research programme.

A word about how to read what follows. It is written for no single discipline, and therefore, we hope, for several. The linguist will find the biology elementary and the chronobiologist will find the linguistics so; each is asked to be patient with the passages that restate what they already know, for the argument of this book lives precisely in the joints between fields, and those joints cannot be exhibited without briefly displaying the bones on either side. We have tried to keep technical prerequisites light and to define terms as they arise; the

mathematics, where it appears at all, is confined to a few sections that a reader may skip without losing the thread of the whole. Those who prefer to glimpse the argument's destination before setting out on its route may begin with the *Plan of the Book* and the concluding chapter, and then double back to whichever parts draw them most. However it is read, the book asks only that its central image be entertained seriously for the length of an argument—the image of language not as a monument standing outside time, but as a living rhythm moving within it.

We are grateful to the many scholars whose work we lean on here, and whose ideas appear in these pages recombined into an unfamiliar order. Whatever is sound in the synthesis is owed to them; the speculative leaps, and any errors, are ours.

Dr. Moninder Singh Modgil

PhD Physics, IIT Kanpur, India

B. Tech. (Hons.), Aeronautical Engineering, IIT Kharagpur, India

research@cosmosresearchlabs.org

Mr. Dnyandeo Dattatray Patil

Electrical Engineering and Artificial Intelligence

Cosmos Research Labs

Centre for Ontological Science, Meta Quanta Physics and Omega Singularity

Contents

Dedication	ii
Preface	iii
List of Figures	x
List of Tables	xii
Plan of the Book	xiii
I Foundations of a Temporal Linguistics	1
1 Language as a Temporal System	2
1.1 The Temporal Turn	2
1.2 Beyond Synchrony and Diachrony	3
1.3 The Central Thesis	4
1.4 An Analogy Worth Its Keep	6
1.5 Three Kinds of Temporal Question	6
1.6 Why Cycles, and Not Merely Time	8
1.7 On the Discipline of Analogy	10
2 From GTC to LTC: The Genetic Analogy	12
2.1 The Central Dogma of Molecular Biology	13
2.2 A Central Dogma of Language	13
2.3 The Correspondence in Detail	15
2.4 Memes, Linguemes, and the Unit of Replication	16
2.5 Where the Analogy Breaks	17
2.6 What the Analogy Licenses	18
2.7 Summary	19
3 The Innate Kernel: Universal Grammar, the Soul, and the Tree of Languages	21
3.1 Chomsky's Kernel	21
3.2 The Search for the Kernel in the Genes	23
3.3 Relocating the Kernel: The Soul	23

3.4	Xenoglossy: Language Across Births	24
3.5	From One Kernel to Many Tongues	25
3.6	The Cycle of Time and the Breaking of Connection	26
3.7	The Tree of Languages and the Groups of Tongues	27
3.8	Summary	27
4	The Architecture of Time Cycles	30
4.1	The Spectrum of Linguistic Time	30
4.2	The Fastest Cycles: Speech Rhythm	31
4.3	Entrainment: The Synchrony of Dialogue	33
4.4	Nesting: Cycles within Cycles	34
4.5	Frequency, Repetition, and Entrenchment	34
4.6	Coupling and the Actuation Problem	35
4.7	Summary	36
II	The Micro-Cycles: Cognition, Rhythm, and the Day	38
5	Circadian and Diurnal Variation in Language	39
5.1	The Day as a Linguistic Cycle	39
5.2	The Chronobiology of Cognition	40
5.3	Time-of-Day Effects on Language Processing	40
5.4	Diurnal Patterns in Language Production	41
5.5	Chronotype and the Synchrony Effect	42
5.6	Seasonal and Longer Cycles	44
5.7	Toward a Linguistic Expression Clock	45
5.8	Summary	45
6	Sleep, Consolidation, and the Linguistic Clock	47
6.1	The Architecture of a Night's Sleep	47
6.2	Sleep Consolidates the Words of the Day	48
6.3	Why Language Needs Two Memory Systems	49
6.4	Does Sleep Improve Your Grammar?	50
6.5	Sleep Loss and the Fraying of Language	51
6.6	The Nightly Reorganisation and the Expression Clock	52
6.7	Summary	52
7	Recursion and Repetition as Temporal Syntax	54
7.1	Two Faces of Temporal Structure	54
7.2	Recursion: Structure Nested in Time	55
7.3	The Real-Time Limits of Recursion	56
7.4	Repetition as Temporal Marker	57

7.5	Meter and Mantra: Language as Oscillator	57
7.6	Structural Priming: Repetition Below Awareness	58
7.7	Computational Models: Recursion versus Iteration	59
7.8	Summary	60
III The Meso-Cycles: Lifespan and Society		61
8	Social Time and Register Shifts	62
8.1	Register as Adaptation to the Social Moment	62
8.2	Language Across the Lifespan	63
8.3	Ritual and Institutional Temporalities	64
8.4	Digital Chronotopes and Register Fluidity	65
8.5	Linguistic Capital and the Marketplace	65
8.6	Sociolinguistic Phase Transitions	67
8.7	Summary	67
9	Acquisition, Transmission, and Linguistic Heredity	69
9.1	The Timetable of Acquisition	69
9.2	Learning from the Stream	70
9.3	Iterated Learning: Transmission as a Structuring Force	71
9.4	Language Born from Broken Transmission	72
9.5	The Reconstruction Each Generation	74
9.6	Summary	74
IV The Macro-Cycles: History and Culture		76
10	Diachronic Genealogy Trees and Computational Phylogenetics	77
10.1	Language Families as Lineages	77
10.2	The Comparative Method	78
10.3	The Lexical Clock and Glottochronology	78
10.4	Frequency and the Conservation of Words	80
10.5	Computational Phylogenetics	80
10.6	Trees, Networks, and the Problem of Borrowing	81
10.7	Archaeogenetics: Aligning Languages, Genes, and Pots	82
10.8	Applications and Visualisations	83
10.9	Summary	83
11	Chronotopic Pragmatics and Discourse	85
11.1	The Chronotope	85
11.2	Deixis and the Origo	86

11.3	The Grammar of Time Across Languages	86
11.4	Narrative and the Ordering of Time	87
11.5	Discourse Across Sociopolitical Cycles	88
11.6	Chronotopes of Digital Media	90
11.7	Orders of Indexicality: The Small Utterance and the Large Scale	90
11.8	Summary	91
12	Language and Cultural Evolution	92
12.1	Language as a Cultural System	92
12.2	The Meme and the Lingueme	93
12.3	Linguistic Ecology and Transmission Biases	93
12.4	Chronocultural Co-Evolution: Technology and Language	94
12.5	Brain–Language Co-Evolution	95
12.6	Temporal Bottlenecks: Death, Creolization, and Revival	96
12.7	The Convergence of the Scales	97
12.8	Summary	97
V	Methods, Models, and Futures	99
13	Experimental LTC Studies	100
13.1	Isolating the Linguistic Clock	100
13.2	Measuring Language Over Time	102
13.3	Neural Oscillations and the Linguistic Hierarchy	102
13.4	Sleep and Second-Language Learning	103
13.5	Longitudinal and Panel Designs	104
13.6	A Programme of LTC Experiments	104
13.7	Summary	104
14	Computational Modeling and AI Systems	107
14.1	Agent-Based Models of Language Change	107
14.2	Neural Sequence Models and the Representation of Time	108
14.3	Iterated Learning and the Evolution of Structure in silico	108
14.4	Corpus Models of Semantic Change	109
14.5	Large Language Models as Temporal Artifacts	111
14.6	Summary	112
15	Toward a Unified Temporal Grammar	114
15.1	The Argument in Brief	114
15.2	The Unity of the Scales	115
15.3	The Analogy Assessed	116
15.4	What a Unified Temporal Grammar Would Be	117

15.5 Open Problems and a Research Programme	118
15.6 The Two Readings	119
15.7 Envoi	119

Appendices	121
A Glossary of Core LTC Terms	121
B The Full GTC–LTC Correspondence	122
C Mathematical Notes: Coupled Oscillators and Entrainment	125
D Corpora, Datasets, and Software for Temporal Linguistics	128
E Glossary of Brahma Kumaris Terms	129

List of Figures

1.1	The three nested scales of linguistic time	5
1.2	The GTC–LTC correspondence	7
2.1	The central dogma, twice	14
2.2	Vertical versus horizontal transmission	18
3.1	Two candidate seats for the linguistic kernel	24
3.2	The tree of languages across the cycle of time	28
4.1	The spectrum of linguistic timescales	31
4.2	Coupling between the scales	36
5.1	Diurnal variation in linguistic measures	42
5.2	Chronotype and the synchrony effect	44
6.1	The architecture of a night’s sleep	48
6.2	Complementary learning systems for language	50
7.1	Recursion versus repetition	55
7.2	Meter as a linguistic oscillator	59
8.1	Vernacular use across the lifespan	64
8.2	The register landscape	66
9.1	The timetable of first-language acquisition	70
9.2	Iterated learning and the transmission bottleneck	72
10.1	The lexical clock	79
10.2	The Indo-European family tree	81
11.1	Temporal deixis and the origo	87
11.2	Narrative time	89
12.1	Chronocultural co-evolution	95
12.2	Brain–language co-evolution	96
13.1	Isolating the linguistic clock	101

13.2 Neural oscillations track nested linguistic units	103
14.1 How sequence models encode time	109
14.2 Diachronic semantic drift	110
15.1 The unified temporal grammar	116
15.2 A research roadmap for temporal linguistics	118

List of Tables

1.1	The three families of temporal question	8
2.1	The GTC–LTC correspondence	15
3.1	The principal language families	29
4.1	Linguistic units and processes by timescale	32
5.1	Predicted diurnal effects on linguistic measures	43
6.1	Sleep and linguistic memory	49
7.1	Devices of recursion and repetition	56
8.1	Registers and their temporalities	63
9.1	Modes of linguistic transmission	73
10.1	Biological and linguistic phylogenetics	82
11.1	Time-encoded pragmatic devices	88
12.1	Transmission biases and their linguistic effects	94
13.1	A programme of LTC experiments	105
14.1	Computational approaches to the time cycle	111
15.1	The LTC framework at a glance	115
B.1	The full GTC–LTC correspondence	123

Plan of the Book

This book advances a single thesis across fifteen chapters: that language is not merely something that changes *in* time but something that is organised *by* time, on many nested and coupled scales at once. The argument is cumulative, and it is built in five parts. Part I lays the conceptual foundations—the thesis itself, the analogy with genetics, the question of where linguistic knowledge comes from, and the formal architecture of nested time cycles. Parts II, III, and IV then descend that architecture scale by scale—through the micro-cycles of cognition and the day, the meso-cycles of the lifespan and society, and the macro-cycles of history and culture. Part V turns from theory to method, and to the future. A reader interested in only one scale may read the corresponding part on its own, but the central contention—that the scales are coupled, and that language is what it is because of that coupling—emerges only when they are seen together.

Chapter 1 opens the case. It argues that the temporality of language has been systematically underweighted since Saussure severed the synchronic from the diachronic, and that the true temporal structure of language is not two-valued but a layered spectrum of scales, all present at once in every utterance. It states the book's central thesis—that language is a temporally *regulated* system, its forms expressed, transmitted, and transformed through a hierarchy of coupled cycles—and introduces the three commitments packed into it: regulation, hierarchy, and coupling. It previews the genetic analogy that motivates the whole framework, insists on the discipline that any analogy demands, and explains why the framework speaks of *cycles* rather than merely of timescales: language, at every level, recurs, oscillates, and returns.

The rest of Part I completes the foundations. Chapter 2 develops the genetic analogy in full, building a detailed correspondence between the central dogma of molecular biology and a proposed “central dogma” of linguistic expression. Chapter 3 takes up the origin of the linguistic “code”—Chomsky's innate kernel—follows the search for it into the genes, proposes that the kernel may be seated in the soul, and traces how a single kernel branches into the tree of the world's languages. Chapter 4 then formalises the three-scale architecture itself, using the linguistic phenomena of rhythm, entrainment, nesting, and coupling to make the notions precise.

Parts II through IV descend the hierarchy. Part II treats the micro-cycles: the circadian and diurnal modulation of language (Chapter 5), sleep and the linguistic “expression clock” (Chapter 6), and recursion and repetition as a temporal

syntax (Chapter 7). Part III treats the meso-cycles of social and biographical time: register and the lifespan (Chapter 8), and acquisition, transmission, and linguistic heredity (Chapter 9). Part IV treats the macro-cycles: diachronic genealogy trees and computational phylogenetics (Chapter 10), chronotopic pragmatics and discourse (Chapter 11), and the integration of LTC with theories of cultural evolution (Chapter 12).

Part V turns from theory to practice. Chapter 13 surveys experimental designs that could test LTC predictions; Chapter 14 treats computational modelling, from agent-based simulations of language change to time-aware neural language models; and Chapter 15 draws the threads together into a single unified picture and a research programme for a temporal linguistics.

Five appendices close the volume and may be consulted at any point. Appendix A is a glossary of the core LTC terms; Appendix B sets out the full GTC–LTC correspondence in a single table; Appendix C gathers the mathematical notes on coupled oscillators and entrainment on which the formal parts of the argument rest; Appendix D catalogues the corpora, datasets, and software available for empirical and computational work; and Appendix E is a glossary of the Brahma Kumaris terms drawn on where the book relates its temporal linguistics to that system’s teaching on the cycle of time.

Part I

Foundations of a Temporal Linguistics

Language as a Temporal System

A language is not a thing that endures; it is a thing that happens, again and again, and is never quite the same twice.

Ask what a language *is* and the grammar books answer with an inventory: a set of sounds, a lexicon, a system of rules for combining them. The answer is not wrong, but it is a photograph of something that is always in motion. No one has ever met a language standing still. What we encounter, always, are acts of speaking and writing—events that unfold in seconds, that differ from morning to evening, that shift across a lifetime, and that drift, over centuries, so far from their origins that the descendants of a single tongue become mutually unintelligible. Language is, in the most literal sense, a temporal phenomenon. It lives in time the way a melody lives in time: take time away and nothing of it remains.

This book is an extended argument that we should take that observation far more seriously than linguistics usually does—not as a background truism but as an organising principle. The claim of *Linguistics in the Time Cycle* (hereafter LTC) is that language is structured by time at every scale on which it exists, and that these scales are not independent. The fast rhythms of cognition, the slow tides of history, and the intermediate arc of a human life are coupled together into a single, multi-layered temporal system. To describe that system—to give it a vocabulary, a set of analogies, and a research programme—is the task of the chapters that follow.

1.1 The Temporal Turn

Modern linguistics was founded, in part, on a deliberate act of temporal abstraction. When Ferdinand de Saussure [43] separated the *synchronic* study of language—language as a system at one moment—from the *diachronic* study of its history, he did so to secure the autonomy of structural analysis. A chess position, Saussure observed, can be described completely without any knowledge of the moves that produced it. So too, he argued, a language state forms a coherent sys-

tem that can be analysed on its own terms, without reference to the changes that brought it about. The move was enormously productive. It made possible the structuralist programme and, through it, much of twentieth-century linguistics.

But the abstraction came at a price, and the price was time. By treating the synchronic system as primary and the diachronic as a separate, almost extra-systemic matter, Saussure encouraged a habit of mind in which change is what happens *to* an otherwise timeless structure, rather than something intrinsic to what language is. Even the metaphor betrays the assumption: the chess position is static; only the player, from outside, sets it in motion. Yet language has no player outside it. Its “positions” are never held; the board is being rearranged continuously, by every speaker, in the very act of using it. The synchronic snapshot is a useful fiction, but we should not mistake the fiction for the thing.

Over the past half-century a quiet “temporal turn” has been reintroducing time into the centre of linguistic theory. Labov [76] showed that variation and change are not noise to be idealised away but structured, quantifiable, and observable in progress. Usage-based and cognitive linguists [25, 34] argued that grammatical structure *emerges* from the repeated events of language use, so that frequency and recency—properties defined over time—are constitutive of grammar rather than peripheral to it. Perhaps most explicitly, the “Five Graces” group recast language itself as a *complex adaptive system*: a population of interacting speakers whose behaviour is shaped by their history of past interactions, exhibiting “perpetual dynamics,” emergence, and phase transitions [10]. In each of these programmes the same realisation recurs: the properties we most want to explain are properties of a system evolving in time.

LTC is an attempt to carry that realisation to its conclusion. If time is constitutive of language, then a theory of language should be, at its core, a theory of *how many different timescales interact*. That is the gap this book tries to fill.

1.2 Beyond Synchrony and Diachrony

The synchronic/diachronic dichotomy, for all its utility, is too coarse. It recognises only two temporal stances—the frozen present and the flowing past—and it treats them as methodologically opposed. But the temporality of language is not two-valued; it is layered. Consider a single English speaker uttering a single sentence. That utterance is shaped simultaneously by:

- the *millisecond* dynamics of articulation and the *second-to-minute* dynamics of planning and working memory;
- the *diurnal* state of the speaker—whether it is morning or late evening, whether she is alert or fatigued;

- the *biographical* stratum—her age, the generation whose slang she absorbed as a teenager, the register she acquired at work;
- and the *historical* stratum—the centuries of sound change and semantic drift that gave the words their present shape.

All of these are “present” in the utterance at once. None of them is the Saussurean synchrony, and none is simply diachrony either. They are distinct *timescales*, each with its own characteristic rhythm, each leaving its fingerprint on the same linguistic act. A theory adequate to language must be able to speak of all of them together, and—crucially—of how they are coupled. The evening fatigue that simplifies a speaker’s syntax [103] operates on a scale of hours; the age-grading that makes her vocabulary sound dated operates on a scale of decades; yet both are audible in the same conversation. The LTC framework is built precisely to hold these scales in one view.

1.3 The Central Thesis

The central thesis of this book can be stated compactly. *Language is a temporally regulated system: its forms are expressed, modulated, transmitted, and transformed through a hierarchy of coupled time cycles, and its structure at any moment is the joint product of processes running at every scale at once.*

Three commitments are packed into that sentence, and it is worth unfolding them.

First, *regulation*. To say that language is temporally *regulated*, not merely temporally *located*, is to claim that time does active work. It is not that linguistic events happen to fall at various points on a timeline; it is that *when* something happens partly determines *what* happens. The metaphorical richness of speech may rise and fall with the emotional arousal of the day; the survival of a new word may depend on the generational moment at which it appears; the meaning of a sentence may hinge on the historical epoch that supplies its connotations. Time gates linguistic expression, much as it gates other biological and cognitive processes.

Second, *hierarchy*. The relevant timescales are not a flat list but a nested hierarchy—micro within meso within macro—in which faster cycles run their course many times inside a single turn of the slower ones. A speaker lives through some twenty-five thousand diurnal cycles in a lifetime, and a language persists through dozens of lifetimes. This nesting is not incidental; it is the source of much of the system’s characteristic behaviour, because events at one scale can accumulate into changes at the next.

Third, *coupling*. The scales are not sealed off from one another. A micro-scale bias—say, a tendency to reduce unstressed syllables when speaking quickly—




Figure 1.1. The three nested scales of linguistic time. Every linguistic act is expressed at the centre, “here and now,” but is modulated simultaneously by micro-cycles (circadian and cognitive rhythms, spanning milliseconds to a day), meso-cycles (the lifespan and generational turnover, spanning years to decades), and macro-cycles (the historical and cultural evolution of the language, spanning centuries). The scales are nested and coupled rather than independent.

can, repeated across millions of utterances and thousands of speakers, become a macro-scale sound change. Conversely, a macro-scale fact—the prestige of a particular dialect—reaches down to shape the micro-scale choices of an individual speaker in a job interview. The upward path is a kind of statistical summation; the downward path is a kind of constraint. Understanding how these couplings work is, in a sense, the whole of the LTC programme.

1.4 An Analogy Worth Its Keep

There is a system, already well understood, that exhibits exactly this architecture: the machinery by which genetic information is expressed and inherited. Genes are not expressed uniformly; they are switched on and off in time, entrained to circadian and seasonal rhythms, regulated by promoters and inhibitors, mutated and drifted and selected across generations, and transmitted by heredity down the long macro-scale of evolution. The parallel to language is striking enough to be more than decorative. It was the point of departure for the original LTC paper, which mapped the framework onto a companion model, *Genetics in the Time Cycle* (GTC), and it will occupy the whole of the next chapter.

Figure 1.2 sets out the correspondence in its simplest form. Read it as a claim about *architecture*, not about substance: language is not made of nucleotides, and nothing here suggests it is. The claim is that the two systems are organised the same way—a code expressed through time, regulated by context, varied by mutation and drift, and inherited across generations—and that the mature conceptual toolkit of evolutionary and regulatory biology can therefore be borrowed, with care, to sharpen questions about language.

An analogy earns its keep in science when it is *generative*—when it makes us ask questions we would not otherwise have asked, and suggests methods we would not otherwise have tried. Judged by that standard the genetic analogy has already proved its worth in historical linguistics, where the mathematics of molecular phylogenetics has been imported wholesale to build quantitative family trees of languages, and where linguistic and genetic lineages are now routinely compared to reconstruct human prehistory [92]. LTC proposes to extend that borrowing from the macro-scale, where it is already established, to the micro- and meso-scales, where it is largely unexplored. If circadian biology can be imported to ask how gene expression varies across the day, the same importation can ask how *linguistic* expression varies across the day—and the empirical tools of chronobiology come with it.

1.5 Three Kinds of Temporal Question

If the framework is to be more than a suggestive picture, it must generate questions of a definite shape. The genetic analogy is helpful here precisely because molecular biology already sorts its questions into recognisable families—questions of *expression*, of *transmission*, and of *transformation*—and each has a linguistic counterpart that the LTC lens brings into focus.

Questions of *expression* concern the fast scales: how a fixed underlying competence is realised differently from moment to moment. In biology this is the domain of gene regulation—why the same genome yields different proteins




Figure 1.2. The master correspondence between the Genetic Time Cycle (GTC) and the Linguistic Time Cycle (LTC). Each genetic process on the left has a linguistic counterpart on the right, and in both systems expression is regulated by time (centre). The analogy is structural: it licenses the transfer of *questions* and *methods* from biology to linguistics, not the literal identification of genes with words.

in different tissues and at different hours. In language it is the domain of performance: why the same speaker, with the same grammar, produces longer and more complex sentences at eleven in the morning than at eleven at night [103, 111], why metaphor comes more readily under emotional arousal, why a word on the tip of the tongue surfaces after a night's sleep. These are questions about how time *gates* the expression of a stable competence, and Part II is devoted to them.

Questions of *transmission* concern the middle scales: how linguistic material passes from one mind, and one generation, to the next. In biology this is heredity, with its bottlenecks, its fidelity and its errors. In language it is acquisition and diffusion—but with the crucial disanalogy noted above, that transmission is horizontal as well as vertical. A child inherits the phonology of her parents but the slang of her peers, and an adult can borrow wholesale from a language she meets in mid-life. The iterated passage of language through the bottleneck of learning is now known to be a powerful shaping force in its own right, capable

Table 1.1. The three families of temporal question that the LTC framework distinguishes, with the timescale each interrogates and its characteristic instance in biology and in language. The rows correspond to Parts II–IV of the book.

Question	Timescale	Biological instance	Linguistic instance
Expression	milliseconds to a day	gene regulation; circadian expression	diurnal variation in syntax, lexicon, metaphor
Transmission	years to a lifetime	heredity; fidelity and error	acquisition and diffusion; iterated learning
Transformation	centuries to millennia	mutation, drift, selection, speciation	sound change, semantic drift, language splits

of manufacturing grammatical regularity out of noisy input. Part III takes up these questions of transmission across the lifespan.

Questions of *transformation* concern the slow scales: how the system itself changes over historical and cultural time. In biology this is evolution proper—mutation, drift, selection, speciation. In language it is sound change, semantic drift, grammaticalisation, and the branching of families into daughters. Here the analogy is oldest and best developed, for historical linguists have borrowed the mathematics of phylogenetics directly, estimating the ages of language splits from rates of lexical replacement much as geneticists date divergences from molecular clocks [3, 92]. Part IV pursues these questions of transformation.

The three families are not watertight, and the most interesting phenomena sit at their seams—where an expression-scale bias, repeated and transmitted, becomes a transformation-scale change. But the taxonomy is useful as a map. It tells us, for any linguistic phenomenon, which timescale to interrogate first, and which region of the biological toolkit is likely to repay borrowing. Much of the work of the chapters ahead is to populate this map with concrete cases.

1.6 Why Cycles, and Not Merely Time

The framework is called Linguistics in the Time *Cycle*, and the choice of word is deliberate. Much of what has been said so far would be equally at home in a theory of linguistic *timescales*—a flat catalogue of durations from the millisecond to the millennium. But the claim of this book is stronger and more specific: that the temporal organisation of language is fundamentally *cyclic*, built out of processes that recur, oscillate, and return, rather than merely elapse. Linear time

is a line along which events are strung; cyclic time is a wheel that turns, bringing the same phase around again and again. Language, I argue, lives on the wheel more than on the line.

The evidence for this is woven through the scales. At the micro-scale the circadian rhythm is, by definition, a cycle: the same trough of alertness returns every twenty-four hours, and with it the same depression of syntactic performance [103, 111]. Metre, rhyme, and the breath-group of speech are cycles nested inside the utterance itself. At the meso-scale the generational turnover of slang is a cycle of rise, saturation, and obsolescence that each cohort re-enacts; the life of a speaker moves through recurring ritual occasions—the wedding, the funeral, the annual festival—each summoning its own register. At the macro-scale, historical linguists have long noted that certain changes recur: languages cycle between analytic and synthetic morphology, cases erode and are rebuilt from postpositions, honorific systems inflate and collapse and inflate again. These are not the same event repeated, but they are the same *shape* of event, returning at a characteristic period. That is what makes them cycles.

Why should this matter? Because a cyclic process invites a different mathematics and a different intuition than a linear one. Cycles have phase, frequency, and amplitude; they can be entrained by external signals, as the body clock is entrained by daylight; they can fall into or out of synchrony with one another; and coupled cycles of different periods produce the rich, quasi-periodic behaviour that dynamical systems theory was built to describe. If language is organised by nested cycles, then the tools for understanding it are the tools of oscillation and entrainment—and Chapter 3 will take up exactly that formalism. It also matters because cycles link the scales in a particular way: a fast cycle, sampled once per turn of a slow cycle, leaves a trace that accumulates. The daily return of evening fatigue, integrated over a life and a community, is one of the quiet engines of the slow drift we call change.

There is a deeper resonance here, too, with the intellectual tradition from which the “time cycle” framing is drawn—one in which time is understood not as an arrow shot once from past to future but as a great wheel, within which individuals, societies, and even languages pass through recurring seasons. One need not adopt any particular metaphysics of eternal return to take the formal point seriously: whether or not history rhymes, *language* demonstrably does, at every scale we can measure. The task of a temporal linguistics is to measure those rhythms, and to understand how the wheels within wheels turn together.

1.7 On the Discipline of Analogy

A word of caution is in order at the outset, because analogical frameworks are easy to abuse. The history of the human sciences is littered with metaphors that hardened into dogma—society as an organism, the mind as a steam engine, then as a telephone switchboard, then as a digital computer. Each captured something and then, over-extended, obscured as much as it revealed. The genetic analogy at the heart of LTC is subject to the same danger. Words are not genes; a language is not literally a genome; “linguistic mutation” is a figure of speech and must earn any explanatory weight it is asked to bear.

I therefore adopt three rules of discipline, and try to keep to them throughout. *First, the analogy is a source of hypotheses, not of conclusions.* Its job is to suggest what to look for; whether the thing is there must be settled by evidence, not by the elegance of the parallel. *Second, disanalogies matter as much as analogies.* Language differs from the genome in ways that are themselves informative—most obviously, linguistic “heredity” is horizontal and promiscuous, since we inherit words from peers and strangers, not only from parents, and a speaker can acquire wholesale from an unrelated language in a way no organism can. Each such difference marks a place where the biological toolkit must be modified rather than merely applied. *Third, every chapter must cash out in something testable*—an experiment, a corpus study, a simulation, a measurable prediction. An analogy that generates no such commitments is literature, not science, and while there is nothing wrong with literature, it is not what this book is trying to be.

With those rules in hand, the analogy becomes a scaffold rather than a cage. It organises an enormous, scattered literature—chronobiology, sociolinguistics, historical phylogenetics, discourse analysis, computational modelling—under a single set of questions, and it lets findings in one area suggest experiments in another. That organising power, and not any claim of literal identity, is the justification for the framework.

Coda: Linguistics and the Brahma Kumaris

The framework’s central move—to treat time as a cycle rather than a line—has a direct precedent in the Brahma Kumaris teaching from which this book takes its name. In that teaching, world history is not an arrow but a wheel, the *kalpa*, an eternally repeating cycle [7, 128]. The insistence of this chapter that language *recurs*—that its structures return, phase upon phase—echoes that cosmology at the linguistic scale. And the image of language “expressed here and now” at the centre of the nested cycles (Figure 1.1) accords with the Brahma Kumaris view of the soul as the present, conscious agent that acts through the body within the great Drama [23]. Where the tradition speaks of a cosmic wheel, this book finds

smaller wheels within it; the claim that language keeps time is, in this light, one turning of a much larger one. We return to these resonances, chapter by chapter, as the temporal architecture is filled in.

From GTC to LTC: The Genetic Analogy

The formation of different languages and of distinct species, and the proofs that both have been developed through a gradual process, are curiously parallel.

—Charles Darwin, *The Descent of Man* (1871)

Darwin was not the first to notice the resemblance, but he put it most memorably. Writing in *The Descent of Man*, he observed that languages and species alike form branching genealogies, accumulate small changes, preserve vestigial “rudiments,” and are winnowed by a struggle for existence among competing forms [40]. The parallel was not idle: the young Darwin had been struck by the philologists’ demonstration that the languages of Europe and India descend from a common ancestor, and the image of a language family tree helped him to conceive of descent with modification in the living world. The analogy that this book proposes to extend is therefore not a modern conceit imported into linguistics from biology; it is a two-way traffic nearly as old as evolutionary theory itself, in which each field has repeatedly lent the other its metaphors and, more recently, its mathematics.

This chapter develops the analogy in full. It is the load-bearing chapter of Part I, for everything that follows—the circadian expression of language, the transmission of grammar across generations, the branching of families—is a working-out, at one scale or another, of the single correspondence set out here. I begin with the biological system to be borrowed from, the central dogma of molecular biology; propose a parallel “central dogma” for language; lay out the detailed correspondence term by term; and then, with equal care, mark the places where the analogy fails. For it is the failures, as much as the parallels, that tell us how to use the framework well.

2.1 The Central Dogma of Molecular Biology

The organising principle of molecular genetics is a directional flow of information, christened by Francis Crick [33] the *central dogma*: information passes from nucleic acid to nucleic acid, and from nucleic acid to protein, but not back out of protein. In its familiar form the flow runs DNA → RNA → protein. A gene, a stretch of DNA, is first *transcribed* into a messenger RNA that carries a complementary copy of its sequence; the messenger is then *translated*, codon by codon, into the chain of amino acids that folds into a protein. DNA is the archive, stable and protected; RNA is the transient message; protein is the expressed product that does the cell's work.

Two features of this system matter most for the analogy to come. The first is the separation of *code* from *expression*. The genome is not the organism; it is a set of instructions, read out differently in different cells and at different times. The same DNA yields a liver cell and a neuron because *gene regulation*—the action of promoters, enhancers, repressors, and, over longer times, epigenetic marks—governs which genes are transcribed, when, and how much [71]. Expression is thus gated: the code is fixed, but its realisation is exquisitely sensitive to context and to time, including the circadian and seasonal cycles that will occupy Chapter 4. The second feature is *variation and inheritance*. The code is copied with high but imperfect fidelity; mutations introduce novelty; the resulting variants are transmitted to descendants and sorted by selection and drift. Together these give the system its evolutionary dynamics.

It is worth dwelling on how thoroughly *temporal* genetic regulation is, because this is the feature the LTC framework most wants to borrow. Gene expression is not merely switched on or off; it is scheduled. A dedicated molecular clock, built from a feedback loop among genes such as PER and CLOCK, drives the circadian oscillation of expression across roughly a twenty-four-hour cycle, and this clock in turn gates thousands of downstream genes so that the same genome presents a different transcriptional face at dawn and at midnight [71]. Longer rhythms—seasonal, developmental, and life-historical—layer further schedules on top. The lesson for language is immediate. If the genome is a code whose expression is timed by nested oscillators, and if language shares the genome's architecture of regulated expression, then we should expect linguistic expression to be timed too: to vary across the day, the lifespan, and the epoch. The whole of Part II is, in effect, the search for language's PER and CLOCK.

2.2 A Central Dogma of Language

Language exhibits a strikingly similar architecture, and it is worth stating the parallel as sharply as possible. Corresponding to the genome is the speaker's




Figure 2.1. The central dogma of molecular biology (top) and a proposed parallel for language (bottom). In each case a stable *code* is turned into a transient *message* and then into an overt *expression*, and in each case the flow is regulated by context and by time. The vertical correspondences—code, message, expression—define the backbone of the GTC–LTC analogy.

grammatical competence—the internalised system of phonological, morphological, syntactic, and semantic knowledge that Saussure called *langue* and Chomsky *competence* [29, 43]. Corresponding to transcription is the encoding of an intended meaning into a structured *utterance plan*; corresponding to translation is the *articulation* of that plan into overt speech or text. The stable code is the grammar; the transient message is the plan; the expressed product is the utterance. Figure 2.1 sets the two flows side by side.

The parallel is more than cosmetic, because the same two features that mattered in biology reappear. First, code is separated from expression: a speaker’s competence is relatively stable across a day, yet its *realisation* varies enormously with fatigue, attention, arousal, audience, and the hour of the clock. This is the linguistic counterpart of gene regulation, and it is precisely what makes a *temporal* linguistics possible—for if expression were a rigid read-out of compe-

Table 2.1. The detailed correspondence between the Genetic Time Cycle (GTC) and the Linguistic Time Cycle (LTC). Each evolutionary or regulatory concept on the left has a linguistic counterpart on the right; the third column names the mechanism through which the linguistic version operates.

Genetics (GTC)	Linguistics (LTC)	Linguistic mechanism
Codon / gene	Phoneme / morpheme / construction	units combined by grammar
Genome	Grammatical competence	internalised system (<i>langue</i>)
Transcription	Encoding of an utterance plan	mapping meaning to form
Translation	Articulation	plan realised as speech or text
Gene regulation	Syntax, pragmatics, timing	context gates expression
Mutation	Innovation / neologism	error, analogy, reanalysis
Genetic drift	Lexical diffusion	chance propagation of variants
Selection	Social prestige, utility	differential replication of forms
Replicator	Lingueme (token in an utterance)	the thing that is copied
Interactor	Speaker	the thing that interacts
Heredity	Acquisition and diffusion	transmission across learners

tence, there would be nothing for time to modulate. Second, language shows variation and inheritance: competence is acquired anew by each speaker, copied with high but imperfect fidelity from the speech of others, subject to innovation, and transmitted onward. The utterance, not the grammar, is where variation is generated and exposed to selection—a point to which we now turn.

2.3 The Correspondence in Detail

To make the analogy do work we must be specific about what corresponds to what. Table 2.1 lays out the mapping term by term, extending the compact scheme of the original LTC paper. Several rows repay comment.

The rows deserving most emphasis are the two that name the *replicator* and the *interactor*. In David Hull [70]’s generalised account of selection—designed to abstract Darwinism away from its biological particulars—evolution requires

entities that are copied (replicators) and entities that interact with the environment in a way that biases copying (interactors). In the organism, genes are the replicators and bodies the interactors. William Croft [34] applied this scheme to language with a single decisive move: the linguistic replicator, he argued, is not the word in the dictionary or the rule in the grammar but the *lingueme*—a token of linguistic structure as actually produced in an utterance. Speakers are the interactors. Language change is then literally a selection process over linguemes: variant pronunciations, constructions, and meanings are thrown up in the flux of utterances, and some are copied more than others because of who uses them and to what effect. This “utterance selection” has since been given an explicit mathematical form borrowed from population genetics [9], and it supplies the micro-level engine beneath the macro-scale trees of Chapter 9.

Locating the replicator in the utterance, rather than in the grammar, is what makes the whole framework temporal. The utterance is an event; it happens at a time; and its production is gated by all the rhythms—circadian, biographical, historical—catalogued in Chapter 1. Selection over utterances is therefore selection distributed across the time cycles, and the analogy with regulated gene expression is at its tightest precisely here.

2.4 Memes, Linguemes, and the Unit of Replication

The claim that language evolves by the differential replication of units raises at once the hardest question in any evolutionary account of culture: what, exactly, is the unit? In biology the question has a clean answer—the gene, a delimitable stretch of sequence—and it is the cleanness of that answer that gives population genetics its power to count, measure, and predict. Culture is less obliging. When Richard Dawkins [41] coined the *meme* as the cultural analogue of the gene—a tune, a catchphrase, a fashion, copied from mind to mind—he offered a suggestive image but not a countable unit. Where does one meme end and the next begin? Is a whole melody a meme, or each phrase, or each note? Without a principled answer the analogy risks becoming, as critics justly complained, a metaphor in search of a measurement.

Croft’s *lingueme* is best understood as an attempt to solve this problem for the special case of language, and its solution is instructive. Rather than seeking the replicator in the abstract system—in a phoneme-as-type or a rule-as-type, whose boundaries are theorist’s artefacts—Croft locates it in the concrete event: a lingueme is a *token* of linguistic structure in a particular utterance [34]. Tokens are individuated naturally, by the occasions of their production; they can, in principle, be counted in a corpus; and their frequencies can be tracked through time. The move mirrors a shift in biology from thinking about the gene as an

abstract Mendelian factor to thinking about it as a physical, copiable molecule. It trades the intuitive appeal of types for the countability of tokens, and countability is what an empirical science of language change requires.

Two consequences follow, both of which shape the rest of this book. First, the choice of unit fixes what can be measured. If linguemes are utterance tokens, then the raw data of linguistic evolution are corpora—time-stamped records of utterances—and the tools are those of frequency, entropy, and diffusion applied across time. This is why Part IV leans so heavily on diachronic corpora and Part V on computational text analysis: they are the instruments that make the replicator visible. Second, the unit problem never fully closes. Linguistic structure is organised at many levels simultaneously—feature, phoneme, morpheme, word, construction—and replication can occur at any of them, so that the “same” change may be described as one lingueme or many depending on the grain of analysis. Far from a defect, this multi-level character is exactly what the nested time cycles of Chapter 3 are designed to represent: different levels of linguistic structure replicate at different characteristic rates, just as different genes are transcribed on different schedules. The unit of replication, in language as in the cell, is not a single thing but a hierarchy of things, each with its own clock.

2.5 Where the Analogy Breaks

An honest analogy advertises its limits. Language departs from the genetic model in at least three ways, and each departure is informative rather than merely inconvenient.

The first and most important is the *mode of transmission*. Genetic inheritance is overwhelmingly vertical: organisms receive their genes from parents and pass them to offspring, and lineages, once split, do not ordinarily exchange genes. Linguistic transmission is promiscuously *horizontal*. A child acquires the phonology of her caregivers but the slang of her peers; adult speakers borrow words, sounds, and constructions wholesale from languages they merely encounter. Where a species tree is (mostly) a branching structure, a language genealogy is a tree overgrown with a web of contact and borrowing (Figure 2.2). This is why phylogenetic methods imported into linguistics must be modified to detect and accommodate horizontal transfer, and why the naive reading of a language tree as pure descent can mislead.

The second departure is that linguistic change can be *Lamarckian*. Acquired characteristics are inherited: an innovation a speaker coins in her lifetime—a new word, a clipped pronunciation—can pass directly into the language and be transmitted onward, with no equivalent of the Weismann barrier that in biology insulates the germ line from the body’s acquired traits. Change can even be




Figure 2.2. A key disanalogy. Genetic information descends vertically, from ancestor to descendant (solid lines). Linguistic material also passes horizontally, through borrowing and contact between coexisting languages (dashed arrows). A realistic model of linguistic “heredity” is therefore a tree crossed by a network.

directed, as when speakers deliberately reform a spelling or revive a dead tongue, a kind of intentional intervention with no counterpart in blind genetic mutation.

The third departure is the looseness of the *genotype–phenotype distinction*. In the organism the separation of gene from body is sharp; in language the “code” (competence) and its “expression” (the utterance) are continuously and reciprocally coupled, since the utterances a speaker hears are exactly what build and reshape her competence. The flow that the central dogma forbids in biology—from expressed product back to stored code—is, in language, the very mechanism of acquisition. This is a profound disanalogy, and we should keep it in view: language is in this respect *more* evolutionarily fluid than the genome, not less.

2.6 What the Analogy Licenses

Granting these limits, what does the correspondence buy us? Its first and best-proven dividend is methodological. Once linguistic descent is seen as formally akin to genetic descent, the whole apparatus of computational phylogenetics

becomes available. Gray and Gray and Atkinson [62] famously applied Bayesian phylogenetic methods to a matrix of 87 Indo-European languages scored for 2,449 lexical items, and dated the family's root to roughly 7,800 to 9,800 years before present—a quantitative result, with error bars, on a question that had resisted a century of qualitative argument. The same toolkit has since shown that languages, like some lineages of organisms, evolve in *punctuational bursts* concentrated around events of splitting [5], and has been used to compare linguistic and genetic lineages to reconstruct prehistoric migration [4, 92]. Chapter 9 takes up these methods in detail.

But the deeper promise of the analogy, and the wager of this book, is that the borrowing need not stop at the macro-scale, where descent and phylogeny live. If language is a regulated system in the same sense that gene expression is a regulated system, then the biology of *expression*—chronobiology, with its circadian clocks and its careful experimental designs—should be importable to the micro-scale, to ask how the read-out of linguistic competence varies across the hours of the day. And if the utterance is a replicator transmitted across learners, then the biology of *inheritance*—with its bottlenecks and its fidelity—should illuminate the meso-scale of acquisition. The correspondence of Table 2.1 is, in this sense, a set of promissory notes, one per row, to be redeemed in the chapters ahead.

There is, finally, a payoff running in the opposite direction, from language back to the general science of cultural change. Because language is the best-documented and most finely structured of all human cultural systems—recorded in vast corpora, segmented into discrete units, and traceable across millennia—it serves as the premier test-bed for the wider programme of cultural evolution [86, 92]. Whatever we learn about how linguistic replicators are expressed, transmitted, and transformed in time is likely to generalise, with suitable care, to other domains of culture that are copied and selected but far harder to measure. The LTC framework thus sits at a productive junction: it imports the mathematics of biological time into linguistics, and it exports the resulting temporal discipline to the study of culture at large.

2.7 Summary

Language and the genome share an architecture: a stable code, turned by regulated processes into a transient message and an overt expression, varied by innovation, and transmitted with modification across generations. That shared architecture, noticed by Darwin and formalised by Hull and Croft, is what licenses the transfer of biological methods to linguistic questions. The transfer must be disciplined by the disanalogies—horizontal and Lamarckian transmis-

sion, and the porous boundary between competence and performance—but these are themselves signposts, marking where language is its own kind of evolving system. With the correspondence in hand, we can make its central notion precise. The next chapter formalises the architecture of nested, coupled time cycles on which the whole framework rests.

Coda: Linguistics and the Brahma Kumaris

This chapter separated a stable linguistic *code* from its regulated *expression*. The Brahma Kumaris teaching draws a strikingly parallel distinction at the level of the person: the soul (*atma*) is the stable bearer of *sanskars*—impressions and tendencies—while the body is the instrument through which those *sanskars* are expressed in thought, word, and deed (*mansa, vacha, karmana*) [23, 128]. The *sanskar*, in this scheme, plays the role the genome plays in biology and competence plays in language: a carried code, expressed variably in time and transmitted onward. The next chapter presses this parallel further, asking whether the linguistic code itself might be borne, like a *sanskar*, by the soul. Here we note only that the architecture of *code and expression*, which the genetic analogy supplies from biology, the Brahma Kumaris supply from their account of the self.

The Innate Kernel: Universal Grammar, the Soul, and the Tree of Languages

A child born in a community learns the language of that community—but the plan of the language he learns is not something he invents; it is something he already, in a sense, knows.

Chapter 2 located the linguistic “code” in the speaker’s grammatical competence, the analogue of the genome. But this only pushes the question back a step. Where does that competence come from? A newborn cannot be taught grammar, for she has as yet no language in which the teaching could be conducted; and yet within a few years every normal child, in every community on Earth, has acquired a system of staggering complexity, on the basis of fragmentary and imperfect evidence. The oldest and deepest problem in the study of language is how this is possible. This chapter takes up that problem, follows the mainstream answer to the edge of what genetics can currently deliver, and then proposes—as the ontological hypothesis of the LTC framework—that the missing piece may lie not in the genome but in the soul. From there it turns to how a single inherited kernel could give rise to the thousands of distinct languages we observe, and offers a model of that divergence in which geographical isolation and the cycle of time together grow the tree of tongues.

3.1 Chomsky’s Kernel

The modern form of the problem, and its most influential answer, are due to Noam Chomsky. Against the behaviourist view that language is learned from scratch by imitation and reinforcement, Chomsky argued that the linguistic input available to a child is far too sparse and too degenerate to determine the rich grammar she ends up with. This is the celebrated *poverty of the stimulus*: children

routinely master constructions they have never been corrected on and could not have induced from the data alone [30]. The only way to close the gap, Chomsky held, is to suppose that much of the structure is already present in the child before experience begins—an innate *Universal Grammar*, a species-wide kernel of linguistic knowledge, engaged and parameterised by exposure to a particular language through what he called the *language acquisition device*.

Chomsky was not alone, and the thesis gathered a formidable following. Eric Lenneberg [81], in *Biological Foundations of Language*, grounded the kernel in maturation, proposing a genetically timed *critical period* for acquisition. Steven Pinker [96] popularised the idea as a “language instinct,” as much a part of human biology as the spider’s web-spinning is of the spider’s. In its most distilled later form, Hauser, Chomsky, and Hauser et al. [67] proposed that the uniquely human core of the faculty—the faculty “in the narrow sense”—might reduce to a single computational capacity, *recursion*, the ability to embed structures within structures without limit. Whatever its exact content, the kernel was taken to be innate, universal, and biologically given.

The thesis has never wanted for critics, and intellectual honesty requires naming them. Usage-based and constructivist linguists, foremost among them Michael Tomasello [124], have argued that general-purpose learning mechanisms, applied to rich statistical structure in the input, can account for acquisition without a dedicated grammar module. The debate between these camps is among the liveliest in the cognitive sciences and is far from settled. But both sides agree on the datum that matters here: the newborn is not a blank slate. Something is already present that makes language possible. The question is what that something is, and where it resides.

It is worth registering how the innateness thesis has itself evolved under pressure, for the movement is telling. Chomsky’s “Principles and Parameters” model refined the kernel into a set of universal principles together with a small number of switches, or *parameters*, that each language sets one way or another—a picture that elegantly reconciles a single innate endowment with manifest surface diversity. Yet even the strongest candidate for a universal, recursion, has been contested empirically: Daniel Everett [53] argued that the Amazonian language Pirahã lacks embedded recursion altogether, a claim that ignited fierce debate precisely because so much theoretical weight had been placed on that single feature. The lesson is not that the kernel is illusory but that its exact content is unexpectedly hard to pin down—which is itself a reason to keep an open mind about where, and in what medium, it might be inscribed.

3.2 The Search for the Kernel in the Genes

If the kernel is innate and biological, it should ultimately be written in the genome, and the natural next step is to look for it there. Chomsky himself repeatedly gestured toward genetics as the eventual home of Universal Grammar. For a brief period it seemed the search might succeed spectacularly. In 2001, Lai and colleagues identified a gene, *FOXP2*, mutated in the members of a London family—the “KE family”—who suffered a severe, heritable disorder of speech and language [79]. The press promptly christened it “the grammar gene.”

The reality proved more sobering, and the sobering is instructive. *FOXP2* turned out to be a transcription factor—a regulatory gene that switches many other genes on and off—expressed in the brain but also in the heart, the lungs, and the gut, and present in nearly identical form in mice and songbirds. Its disruption impairs the fine motor sequencing of speech rather than grammar as such, and it is, at most, one component in a vast and still-obscure developmental network. There is no “gene for grammar,” and after two decades of molecular genetics none has appeared. The kernel that Chomsky’s arguments seem to require has not been found in the DNA, and it is fair to say that the genetic programme, for all its brilliance, has not delivered a clear account of where Universal Grammar is stored or how it is encoded. This is not a refutation of innateness—absence of evidence in so young a science is weak evidence of absence—but it is an honest impasse, and it is the impasse from which the present framework departs.

3.3 Relocating the Kernel: The Soul

Here the LTC framework makes its most distinctive, and frankly its most speculative, proposal. The impasse of the genetic programme may reflect not a gap in our data but a mistake about the address. Suppose the kernel of linguistic knowledge is real, innate, and prior to experience, exactly as Chomsky argued—but suppose it is carried not by the body’s genome but by the *soul*: the conscious, non-material self that, on the ontology this book adopts, is the true seat of cognition and identity. On this view the newborn brings her linguistic endowment with her not from her parents’ chromosomes but from her own continuous existence as a conscious being; the genome builds the instrument, the vocal tract and the neural circuitry, but the musician who already knows how to play is the soul (Figure 3.1).

This is, of course, a metaphysical hypothesis, and it will not persuade a reader who rejects the underlying ontology. It is offered here not as an established result but as a research proposal in the programme of what our institution calls ontological science: a hypothesis to be judged, like any other, by the phenomena

WHERE DOES THE KERNEL LIE?




Figure 3.1. Two candidate seats for the innate linguistic kernel. Mainstream theory seeks it in the genome, where no clear “grammar gene” has been found (dashed). The LTC framework proposes instead that the kernel is carried by the soul (solid), a hypothesis for which the phenomena of xenoglossy are taken as suggestive evidence.

it can order and the predictions it can motivate. Its appeal is that it takes the Chomskyan arguments wholly seriously—the kernel is real and innate—while declining to insist that innateness must mean “encoded in DNA.” And unlike the purely genetic account, it makes contact with a class of reported phenomena that the genetic account cannot easily touch: cases in which linguistic knowledge appears in an individual who could not have acquired it in this life at all.

3.4 Xenoglossy: Language Across Births

Xenoglossy is the reported ability to speak or write a language that one has never learned. The phenomenon was investigated most systematically by the psychiatrist Ian Stevenson of the University of Virginia, who over several decades documented cases in careful detail and drew a crucial distinction [117, 118]. *Recitative* xenoglossy is the mere ability to reproduce fixed strings—words or

phrases—in an unknown tongue, which might in principle be explained by forgotten exposure. *Responsive* xenoglossy is far stronger: the ability to converse intelligibly in the unknown language, answering novel questions with novel and appropriate replies, which cannot be reduced to rote recall. In his best-known case, an Indian woman, Uttara, who in ordinary life spoke Marathi, would in an altered state manifest a personality, Sharada, who conversed fluently in Bengali—a language Uttara had never studied.

To this literature one of the present authors (M. S. M.) can add a case from his own family, offered as testimony rather than as controlled evidence. His wife, Mamta, as a young child of four or five years, would spontaneously fill page after page in the *Sindhi* script—a script she had never been taught and to which, so far as her family could establish, she had had no exposure. The episode ceased as she grew older, as such cases characteristically do. Within the ontology proposed here the interpretation is natural: the linguistic knowledge of a previous life, carried by the soul across the interval of death and rebirth, surfaced briefly in the present one before the demands of the current life overwrote it.

Two cautions must be stated plainly. First, xenoglossy is not accepted by mainstream science; the cases are contested, hard to verify, and impossible to reproduce on demand, and sceptics offer deflationary explanations for many of them. Second, a single family anecdote proves nothing on its own. The claim of this section is therefore modest and conditional: *if* responsive xenoglossy occurs even once, then linguistic knowledge can exist in a person independently of any learning or genetic inheritance in this life—and that would be direct evidence that the kernel's true carrier is the soul, exactly as Figure 3.1 proposes. The phenomenon is offered as a signpost toward a research programme, not as a proof.

3.5 From One Kernel to Many Tongues

Grant, for the sake of the argument, that every soul carries the same universal kernel. An immediate puzzle follows. If the endowment is universal and identical, why are the languages of the world so profoundly different? Why does the child in Chennai grow into Tamil and the child in Reykjavík into Icelandic, when both begin from the same seed?

The answer is the one historical linguistics has given since the nineteenth century, now recast in the present framework. The kernel fixes what all languages share—the deep design, the space of possible grammars—but it underdetermines the surface, leaving a vast range of parameters to be set by the community. Once two populations are separated, their languages drift independently: sounds shift, words are coined and borrowed and lost, constructions are reanalysed, and the

accumulated changes, sampled generation by generation, carry the two varieties steadily apart. Given enough time and enough separation, mutual intelligibility is lost and one language has become two [26]. The comparative method, and its quantitative descendant lexicostatistics [120], reconstruct this process in reverse, using shared retentions to infer common ancestry and rates of change to estimate the depth of the split—the same logic that, in Chapter 2, we saw scaled up into computational phylogenetics.

The two essential ingredients, then, are *geographical isolation* and *change over time*. Neither alone suffices: isolation without change would freeze a language in place, and change without isolation would sweep a whole speech community along together. It is their combination that generates diversity—and it is here that the temporal framework of this book meets the oldest questions of historical linguistics, and that a specific cosmology of time can be brought to bear.

3.6 The Cycle of Time and the Breaking of Connection

The framework of this book takes its name and its temporal architecture from a particular understanding of time as cyclic, and it is worth making the debt explicit, for it supplies a striking account of *when* the divergence of languages began. In the cosmology taught by the Brahma Kumaris, world history is an eternally repeating cycle—a *kalpa*—divided into four ages of declining spiritual power: Satyug (the Golden Age), Treta (the Silver Age), Dwapar (the Copper Age), and Kaliyug (the Iron Age), standing in the proportion 4 : 3 : 2 : 1 [7]. In the first two ages soul-consciousness is at its height and souls remain, in the framework's terms, *connected*—to their source and to one another. The decline into duality and separation begins precisely at the onset of Dwapar.

The LTC framework proposes to read the linguistic record against this template. In the connected ages, on this reading, humanity shares effectively a single tongue: a unified expression of the common kernel, sustained by the very connectedness of souls. The breaking of that connection at the beginning of Dwapar is then also the beginning of linguistic divergence. As soul-consciousness wanes and populations disperse, the two ingredients of the previous section—isolation and time—come fully into play, and the single stem begins to branch. Kaliyug, the age of maximal separation, is correspondingly the age of maximal linguistic diversity: the several thousand mutually unintelligible languages of the present are, in this picture, the fully unfolded crown of a tree whose trunk stood in the Golden Age. The reader who does not share this cosmology may treat it as a heuristic frame; but even as a frame it has the virtue of binding the mechanism of divergence to a definite temporal origin.

It is worth noting that the notion of a single ancestral tongue is not foreign to

mainstream inquiry, even if it lies at its speculative edge. The question of *monogenesis*—whether all human languages descend from one original language—has been debated since antiquity, and a minority of historical linguists have attempted to reach beyond the securely reconstructed families toward deeper groupings such as Nostratic, and ultimately toward a single “Proto-World” [107]. Most linguists regard such deep reconstructions as unprovable, because the signal of common ancestry decays below detectability after roughly ten millennia. The LTC reading does not depend on the comparative method reaching so far; it supplies, from the cycle of time, an independent reason to expect a single stem—and locates its branching at a definite moment, the breaking of connection at the onset of Dwapar.

3.7 The Tree of Languages and the Groups of Tongues

Whatever one makes of its origin, the branching structure itself is not in dispute: the world’s languages fall into a manageable number of great *families*, or groups, each descended from a common ancestral language. Figure 3.2 renders this as a single tree mapped onto the cycle of time—one connected tongue in the upper ages, the connection breaking at Dwapar, and the familiar families fanning out through Dwapar and Kaliyug.

The principal families are set out in Table 3.1. Two of them, Indo-European and Sino-Tibetan, between them account for the native speech of more than half of humanity; the six largest cover the overwhelming majority [48]. Each is itself a tree: Indo-European branches into the Romance, Germanic, Slavic, Indo-Iranian, and other groups, and each of these branches again, so that the structure is fractal, the same branching logic repeating at every scale—exactly the multi-level, multi-rate replication that Chapter 2 anticipated and that the nested time cycles of this book are built to represent.

3.8 Summary

The oldest problem in linguistics is the acquisition of a boundlessly complex system from impoverished data, and the most influential answer is Chomsky’s: an innate, universal kernel of linguistic knowledge. The search for that kernel in the genome has not succeeded; no “grammar gene” has been found. The LTC framework proposes, as an ontological hypothesis, that the kernel is carried instead by the soul, and it points to the contested phenomena of xenoglossy—including a case from the authors’ own family—as the kind of evidence that would, if sound, decide the question in its favour. From a single shared kernel the thousands of actual languages arise through the joint action of geographical isolation and change over time; and read against the cycle of time, that divergence begins with




Figure 3.2. The tree of languages mapped onto the cycle of time. A single connected tongue in the Golden and Silver ages branches, from the breaking of connection at the onset of Dwapar, into the great language families, which diversify further through Kaliyug under the joint action of geographical isolation and change over time.

the breaking of soul-connection at the onset of Dwapar, growing the tree of language families whose crown is the diversity of the present age. The mechanism of that branching—how small changes accumulate, generation upon generation, into the splitting of tongues—is a matter of the slow macro-cycles of history, to which Part IV will return. First, however, we must make the architecture of the time cycles themselves precise.

Coda: Linguistics and the Brahma Kumaris

The proposals of this chapter are the book’s closest approach to the Brahma Kumaris teaching, and it is worth marking the correspondences explicitly. The claim that the linguistic kernel is borne by the soul restates, in the vocabulary

Table 3.1. The principal language families (groups) of the world, with approximate counts of member languages and native speakers and representative members. Figures are order-of-magnitude and follow standard catalogues [48].

Family (group)	Langs.	Speakers	Representative members
Indo-European	~440	2.9 bn	English, Hindi, Russian, Spanish
Sino-Tibetan	~400	1.4 bn	Mandarin, Tibetan, Burmese
Afro-Asiatic	~380	0.5 bn	Arabic, Hebrew, Hausa
Niger-Congo	~1500	0.44 bn	Swahili, Yoruba, Zulu
Austronesian	~1200	0.39 bn	Malay, Tagalog, Malagasy
Dravidian	~85	0.25 bn	Tamil, Telugu, Kannada

of linguistics, the tradition's doctrine that the soul carries its *sanskars* across births [7, 23]; xenoglossy is then simply a *sanskar* of language surfacing out of turn. The account of divergence, too, follows the tradition's own timeline: unity of tongue in the connected ages of Satyug and Treta, and the onset of separation—linguistic and otherwise—at the breaking of connection that begins Dwapar [128]. Even the image of the language tree has its counterpart in the *Kalpa Vriksha*, the world tree whose single trunk and many branches the Brahma Kumaris use to depict the unfolding of humanity through the cycle. The reader is free to take these as metaphor; within the framework they are read as one structure seen from two disciplines.

The Architecture of Time Cycles

Speech is rhythm within rhythm within rhythm—the beat of the syllable inside the swing of the phrase inside the turn of the conversation.

The previous chapters have spoken freely of “micro,” “meso,” and “macro” cycles, of nesting and of coupling, as though these terms were already precise. They are not yet. This chapter makes them so—but it does the work with the materials of linguistics rather than of physics. Where a mathematician would reach for coupled differential equations, we shall reach for the syllable, the turn at talk, the accommodation of one speaker to another, and the slow riddle of how a language changes at all. The claim to be established is that the temporal organisation of language is genuinely hierarchical: that language runs on many clocks at once, that the fast clocks are nested inside the slow, and that the scales are coupled so that events on one become causes on another. This architecture is the skeleton on which the rest of the book hangs.

4.1 The Spectrum of Linguistic Time

Consider how many different durations are simultaneously relevant to a single conversation. The articulation of a single consonant occupies a few tens of milliseconds; a syllable, a couple of hundred; a word, a few hundred more. The gap a listener leaves before taking her turn averages, astonishingly, only about two hundred milliseconds [119]. A whole utterance lasts seconds; a conversation, minutes; and the register the speakers adopt may drift across the hours of the day. Beneath all of this run slower processes still: the years over which each speaker acquired her language, the decades across which her generation reshaped it, the centuries of sound change that gave the words their forms, and the millennia over which the language diverged from its sisters. These durations span more than twelve orders of magnitude, from the hundredth of a second to the hundred centuries.

Figure 4.1 arranges these processes on a single logarithmic axis, and Table 4.1 sets out the same information with representative linguistic units and examples.




Figure 4.1. The spectrum of linguistic timescales, on a logarithmic axis spanning more than twelve orders of magnitude. Each labelled event is a linguistic process with its own characteristic duration; the bands group them into the micro-, meso-, and macro-cycles of the LTC framework. A single conversation implicates the entire left-hand region at once, while resting on the slow processes to the right.

The picture makes vivid a fact easily lost in prose: linguistic time is not one scale but a near-continuous spectrum of them. The three-fold division into micro, meso, and macro is a convenience imposed on this continuum—useful, because the processes do cluster, but not a claim that there are exactly three clocks. There are many, and the labels simply name three regions of a spectrum.

4.2 The Fastest Cycles: Speech Rhythm

At the fast end of the spectrum, language is overtly periodic. Speech has rhythm—a recurrence of prominences in time that listeners feel as strongly as they feel the beat of music. Classical phonetics divided languages into rhythmic types: *stress-*

Table 4.1. Linguistic units and processes arranged by characteristic timescale, with the corresponding LTC band and an illustrative example. Durations are order-of-magnitude.

Unit / process	Duration	Band	Example
Phoneme	~50–200 ms	micro	release of a stop
Syllable / turn-gap	~200 ms	micro	inter-turn silence
Word	~400 ms	micro	lexical retrieval
Utterance / turn	seconds	micro	a spoken sentence
Conversation	minutes	micro	a single exchange
Style-shift	hours (diurnal)	micro	formal → casual
Acquisition	years	meso	learning an L1
Generational change	decades	meso	new slang cohort
Sound change	centuries	macro	a vowel shift
Family divergence	millennia	macro	Latin → Romance

timed languages such as English, in which stressed syllables are said to recur at roughly equal intervals, and *syllable-timed* languages such as Spanish, in which it is the syllables themselves that seem evenly spaced [1, 95]. The strict version of this “isochrony” does not survive measurement—the intervals are not in fact equal—but something real underlies the intuition, and it is best captured not by a ruler but by an *oscillator*. On this view a speaker carries an internal rhythmic cycle, and the tendency is for stressed syllables to be attracted toward its beats; the result is a statistical, not a mechanical, regularity [37]. The important point for us is the form of the explanation: the fastest linguistic cycle is modelled as an entrainable oscillation, exactly the vocabulary the LTC framework needs, arrived at here on purely phonetic grounds.

That the rhythm classes are real, even if isochrony is not, has been confirmed by turning the intuition into measurement. Rather than timing intervals between stresses, phoneticians now quantify the *variability* of consonantal and vocalic intervals across an utterance: stress-timed languages, packing consonants into complex clusters and reducing unstressed vowels, show high variability, while syllable-timed languages show low [97]. These rhythm metrics recover the traditional groupings as regions of a continuous space rather than as discrete types—the same lesson, at the fast end, that Figure 4.1 taught for the spectrum as a whole. Rhythm is gradient, but it is measurably there, and it is periodic

enough to serve as language's metronome.

Nor is there only one fast cycle. Prosodic phonology recognises a whole *hierarchy* of rhythmic units nested one inside another: the syllable within the foot, the foot within the phonological word, the word within the phonological phrase, and so up to the intonational phrase and the utterance [89]. Each level has its own characteristic span, and each nests within the one above. The rhythm of speech is thus already, in miniature, the whole architecture this chapter describes: cycles within cycles, the faster completing many periods inside a single turn of the slower.

4.3 Entrainment: The Synchrony of Dialogue

If rhythm is language's internal clock, dialogue is where two such clocks are brought into alignment. Conversation is a marvel of temporal coordination. Speakers do not simply take turns; they take them with a precision that leaves, on average, only about a fifth of a second of silence between them—and they do so across every language yet measured, a striking universal beneath modest cultural variation [119]. Since a word takes longer than that to plan and articulate, the listener must be predicting the end of the current turn and preparing her own while the speaker is still talking. The classic account of this machinery, the turn-taking model of Sacks, Schegloff, and Sacks et al. [108], describes a system of rules operating in real time at transition-relevant places—a temporal protocol of remarkable exactness.

The exactness demands anticipation. Because a reply cannot be planned in the fifth of a second that separates turns, the listener must project the end of the current turn before it arrives, using syntactic, prosodic, and semantic cues to forecast the landing point and to launch her own turn to meet it. Controlled experiments confirm that hearers predict turn-ends with considerable accuracy, and that syntax carries most of the predictive load [42]. Dialogue is thus not a relay of discrete signals but a continuously coupled system in which each participant runs a real-time model of the other—two clocks locked not after the fact but in anticipation.

Turn-taking is only the most visible synchrony. Over the course of a conversation, interlocutors progressively *align*: they come to reuse one another's sounds, words, and grammatical constructions, converging on a shared way of speaking. Pickering and Pickering and Garrod [94] model this interactive alignment as automatic priming propagating simultaneously at the phonetic, phonological, lexical, syntactic, semantic, and situational levels—six coupled channels of coordination. Over longer spans the same tendency appears as *accommodation*: speakers adjust their accent and style toward (or, to mark distance,

away from) their interlocutors [57]. All of these are forms of *entrainment*: the synchronisation of one speaker's linguistic rhythms and choices to another's. Entrainment is the mechanism by which the private oscillators of individuals are coupled into the shared cycles of a community—and thus the first bridge from the micro-scale of the utterance to the larger scales above it.

4.4 Nesting: Cycles within Cycles

The examples accumulate into a general principle: linguistic cycles are *nested*. We saw it already within speech rhythm, where the syllable nests in the foot and the foot in the phrase. It holds equally across the wider spectrum. A turn at talk is built from utterances, which are built from prosodic phrases, which are built from words and syllables; and the turn is itself a unit within the larger cycle of the conversation, which is one episode within the speaker's biography, which is one life within the history of the language. At every step the faster cycle runs to completion many times inside a single period of the slower one. A speaker produces on the order of a hundred million syllables in a lifetime; a language passes through dozens of lifetimes in a single sound change.

This nesting is not merely a tidy way of listing durations. It is the structural precondition for everything that follows, because it is what allows events at one scale to *aggregate* into events at the next. A bias that shows up once per syllable, invisibly small, has a hundred million chances to express itself in one speaker and countless more across a community. Nesting turns the repetition of fast cycles into a statistical pressure on slow ones. To see how that pressure actually reshapes a language, we must turn to the coupling between the scales.

4.5 Frequency, Repetition, and Entrenchment

Nesting explains how fast cycles *could* aggregate into slow ones; but the aggregation is not merely arithmetic, and the mechanism that converts repetition into structure deserves a section of its own, for it is the true engine of the meso-scale. Usage-based linguistics has shown that the sheer *frequency* with which a form is used—how often its cycle turns—directly shapes its fate in the language. Joan Bybee [24, 25] documents two complementary effects, both consequences of repetition acting on memory. High token frequency *entrenches* a form: often-used words and phrases are stored and accessed as wholes, becoming resistant to the analogical pressures that regularise their rarer neighbours—which is why the commonest verbs of a language (*be*, *go*, *have*) are so often the irregular ones, fossils preserved by frequency. At the same time, high frequency drives phonetic *reduction*: the more a sequence is repeated, the more it is compressed and

automated, so that *going* to becomes *gonna* and, over centuries, a full verb erodes into a grammatical marker.

Here, then, is the missing link between the scales. Repetition at the micro-scale does not simply accumulate passively; it actively sculpts, through entrenchment and reduction, the very units that the slower cycles transmit. *Chunking*—the binding of frequently co-occurring elements into a single processing unit—is the cognitive hinge on which this turns, and it operates continuously, every time a cycle repeats. Grammar itself, on this view, is the slow residue of fast repetition: the sedimentation of micro-scale usage into meso- and macro-scale structure. The architecture of time cycles is not a static filing system for durations but a productive machine, and frequency is the rate at which it runs.

4.6 Coupling and the Actuation Problem

The scales are not sealed. They are *coupled*, and the coupling runs both ways (Figure 4.2). The upward path is *accumulation*: small, systematic biases in fast-scale performance, summed over enough utterances and enough speakers, become slow-scale change. The best-understood instance is phonetic. John Ohala [91] showed how the ordinary mechanics of speech production and perception—a listener occasionally mis-parsing a coarticulated sound, a speaker slightly reducing an unstressed vowel—supply, at the millisecond scale, exactly the raw variation from which centuries-long sound changes are built. Each micro-event is trivial; their accumulation is a vowel shift. The downward path is *constraint*: the macro-scale facts of a language—its established norms, its prestige forms, its standardised spelling—reach back down to shape the micro-scale choices of every speaker, who adjusts her pronunciation in a job interview precisely because a slow-scale norm bears on a fast-scale act.

It is along the upward arrow that the deepest unsolved problem in the study of language change is located. Weinreich, Labov, and Weinreich et al. [127] posed it as the *actuation problem*: if the phonetic and cognitive biases that drive a change are present in all speakers of a language at all times, why does the change actually begin in one particular community at one particular moment, and not in countless others where the same conditions hold? The raw material is always present; the ignition is rare and localised. In the vocabulary of this book, the actuation problem is precisely the problem of how a perpetually available micro-scale variation occasionally succeeds in propagating up the coupling to become a macro-scale fact. We do not solve it here—no one has—but the architecture makes clear what kind of problem it is: a question about the transmission of a signal across coupled scales, of the sort the temporal framework is designed to pose sharply even where it cannot yet answer.



Figure 4.2. The two-way coupling between the linguistic timescales. Fast-scale biases accumulate upward into slow-scale change; slow-scale norms constrain fast-scale choices downward. The *actuation problem*—why a given change begins where and when it does—is the still-unsolved question of what happens along the upward arrow.

Once a change is actuated, its passage up the coupling has a characteristic temporal shape. Sociolinguists observing change in progress find that new forms spread through a community not linearly but along an *S-curve*: slow at first, then accelerating as the innovation gains social momentum, then levelling off as it approaches completion [75]. The same logistic profile recurs whether the changing item is a pronunciation, a word, or a syntactic construction—and remarkably, when a single change advances across several grammatical contexts at once, it often does so at one constant underlying rate, as though a single meso-scale process were driving all its surface reflexes together. The S-curve is the visible signature, at the macro-scale, of a micro-scale variation that has successfully propagated; it is what actuation looks like once it has happened.

4.7 Summary

Language runs on many clocks. They form a spectrum, from the tens of milliseconds of a phoneme to the millennia of a language family, that the LTC framework

partitions for convenience into micro-, meso-, and macro-cycles. At the fast end, speech rhythm behaves as an entrainable oscillation and prosodic structure is already a nested hierarchy of cycles; in dialogue, the private clocks of speakers are entrained to one another through turn-taking, alignment, and accommodation. The cycles are nested, so that fast periods repeat many times within slow ones, and coupled, so that fast biases accumulate upward into change while slow norms constrain downward—with the actuation problem sitting, unsolved, on the upward path. This is the architecture; the rest of the book explores its floors. We begin at the bottom, with the micro-cycles of cognition and the circadian day, to which Part II now turns.

Coda: Linguistics and the Brahma Kumaris

The architecture of this chapter—fast cycles nested within slow, the whole forming a single coupled system—has a cosmological counterpart in the Brahma Kumaris *kalpa*, itself divided into four nested ages that repeat as one turning wheel [7, 128]. The tradition's ideal of the *trikaldarshi*, the “seer of the three aspects of time” who apprehends past, present, and future at once, is in effect the aspiration to hold the whole nested structure in a single view—the very stance this chapter asks of the linguist who would keep the millisecond and the millennium together. Where the Brahma Kumaris cultivate that vision contemplatively, the temporal linguist cultivates it methodologically; both proceed from the conviction that the moment is legible only against the cycle that contains it.

Part II

The Micro-Cycles: Cognition, Rhythm, and the Day

Circadian and Diurnal Variation in Language

The same speaker is not the same speaker at dawn and at midnight; the grammar is constant, but its voice keeps the hours.

Part I built the framework; Part II begins to inhabit it, starting at the fastest of the recurring scales that the framework recognises—the cycle of the day. Chapter 2 argued that language, like the genome, separates a stable code from its regulated expression, and that this separation is what makes a temporal linguistics possible: if competence is constant but its realisation is gated by context and by time, then there is something for time to modulate. The most regular of all temporal contexts is the twenty-four-hour cycle of waking and sleeping, alertness and fatigue. This chapter asks a simple question with surprisingly deep ramifications: *does language change across the day?* The answer, on the evidence assembled below, is that it does—systematically, measurably, and in ways the LTC framework leads us to expect.

5.1 The Day as a Linguistic Cycle

That cognition varies across the day is among the best-established facts in psychology. What is far less appreciated is that *language* rides on those variations. A speaker's competence does not fluctuate between breakfast and bedtime—she does not forget the rules of her grammar at nightfall—but the resources that turn competence into performance certainly do. Attention, working memory, processing speed, and mood all follow a daily rhythm, and every one of them is implicated in the production and comprehension of speech. If those resources ebb and flow on a circadian schedule, then the fluency, complexity, diversity, and affective colouring of language should ebb and flow with them. The day becomes, in effect, a linguistic cycle: a recurring modulation of expression laid over an unchanging code, precisely the regulated-expression architecture of

Chapter 2 realised at the scale of hours.

5.2 The Chronobiology of Cognition

To see why, we need the barest sketch of how the daily rhythm of cognition is generated. The dominant account is Borbély's *two-process model*, on which alertness at any moment is the resultant of two interacting forces [18]. *Process C* is the circadian signal proper, an endogenous oscillation driven by the body clock that promotes wakefulness on a roughly sinusoidal daily schedule regardless of prior sleep. *Process S* is homeostatic *sleep pressure*, which accumulates monotonically across the waking day—the longer one is awake, the greater it grows—and dissipates during sleep. Moment-to-moment cognitive capacity reflects the difference between the two: high in the late morning when the circadian signal is strong and sleep pressure still low, sagging in the mid-afternoon, and falling in the late evening as sleep pressure mounts.

Two refinements matter for language. First, the phase of Process C is not the same for everyone. An individual's *chronotype*—the timing of her internal day—varies widely across a population, from decided “larks” whose rhythms run early to equally decided “owls” whose peak comes late, with most people somewhere between [102]. Second, and consequently, cognitive performance obeys a *synchrony effect*: people perform best at the time of day that matches their chronotype, and worst at their off-peak hours [83]. Any account of language across the day must therefore reckon not with one universal curve but with a family of curves, shifted along the clock by chronotype—a complication that, as we shall see, is also an experimental opportunity.

5.3 Time-of-Day Effects on Language Processing

Does this circadian machinery actually reach language? Until recently the question was barely posed; higher cognitive functions were assumed to be immune to so “low-level” an influence. The assumption has not survived testing. In a study whose title captures the thesis exactly—“The Tick-Tock of Language”—Rosenberg, Pusch, and Rosenberg et al. [103] probed minimal syntactic processing across conditions of circadian phase and elevated sleep pressure and found that even elementary grammatical operations are sensitive to the time of day and to accumulated time awake. Language processing, it turns out, keeps the clock.

The neural correlate has since been imaged directly. Rosenberg, Reske, Warbrick, and Rosenberg et al. [104] scanned early, intermediate, and late chronotypes during a semantic task and found that chronotype modulated activation

in a network of language-related cortical regions, with early and late types differing measurably in both reaction time and the strength of the haemodynamic response. The conclusion the authors draw is the one this chapter needs: *even basic language processes are associated with chronotype-specific neuronal mechanisms*. The read-out of linguistic competence is, at the level of the brain, a time-of-day-dependent affair.

5.4 Diurnal Patterns in Language Production

If comprehension keeps the clock, so should production—and here the digital record of everyday language provides evidence at a scale earlier researchers could only dream of. In a landmark study, Golder and Golder and Macy [59] applied automated text analysis to some five hundred million public messages from millions of users across dozens of countries, tracking the emotional content of word choice hour by hour. The diurnal signal was unmistakable and cross-culturally robust: positive affect in people’s language is highest shortly after waking, declines across the working day, and shows a modest evening recovery—exactly the profile the two-process model predicts, and a rhythm written not in physiology but in *words*. Figure 5.1 sketches this pattern together with the complementary curve of syntactic resources and the rising tide of sleep pressure that helps drive both.

Beyond affect, the framework predicts a family of further effects, set out in Table 5.1. Syntactic complexity and fluency should be greatest when processing resources peak and should decline into the fatigued evening; lexical diversity should track available working memory; figurative and metaphorical language, bound up with emotional arousal, should follow the affective rhythm; and disfluencies and speech errors should rise as sleep pressure mounts. Some of these predictions are already supported, others remain to be tested with the time-stamped corpora and wearable-sensor methods that Part V will describe. What matters here is that they are genuine, quantitative predictions, generated directly by the LTC picture of language as regulated expression.

The production of *errors* deserves particular attention, for slips of the tongue offer an unusually direct window onto the machinery of expression. The systematic study of such slips has long used them to infer the stages of language production—the level at which sounds, words, and syntactic frames are assembled [38]. If, as the two-process model implies, the resources that keep production on track are depleted as sleep pressure mounts, then the rate and the *type* of slips should shift predictably across the day: more errors overall in the fatigued evening, and a changing balance among them as different production stages lose precision at different rates. The diurnal patterning of speech errors




Figure 5.1. Schematic diurnal profiles of three quantities relevant to language. Positive affect in word choice peaks after waking and declines through the day, with an evening rebound (after 59); syntactic complexity and available resources peak in the late morning; homeostatic sleep pressure rises monotonically across waking hours. The curves are illustrative, not fitted data.

is thus not merely a curiosity but a potential assay of the linguistic expression clock, reading off the state of the production system from the mistakes it makes at each hour.

5.5 Chronotype and the Synchrony Effect

The synchrony effect turns chronotype from a nuisance variable into a powerful experimental lever. Because a lark and an owl reach their cognitive peaks at different clock hours (Figure 5.2), the same linguistic task administered at, say, nine in the morning falls at the lark’s best hour and the owl’s worst. Any purely time-of-day account would predict identical performance; the synchrony account

Table 5.1. Predicted diurnal patterns for a range of linguistic measures, with the proposed circadian or homeostatic mechanism and the current evidential status within the LTC research programme.

Measure	Predicted daily pattern	Mechanism	Status
Positive affect	high on waking, declines, evening rebound	circadian sleep pressure	+ observed
Syntactic complexity	peaks late morning, declines by evening	processing resources	suggestive
Lexical diversity	peaks near midday	working-memory capacity	suggestive
Figurative language	tracks emotional arousal	affective rhythm	hypothesised
Disfluency / errors	rise with time awake	homeostatic pressure	suggestive
Pragmatic care	lower when fatigued	self-regulation	hypothesised

predicts a crossover, and it is the crossover that is observed [83, 104]. This is methodologically precious, for it lets the researcher dissociate the effect of *clock time* from the effect of *internal circadian phase*—the linguistic analogue of the constant-routine and forced-desynchrony protocols by which chronobiologists isolate the body clock. A speaker’s chronotype, in short, is a second hand on the linguistic clock, and attending to it converts an uncontrolled source of variance into a designed contrast.

The synchrony effect also carries a subtler and more surprising lesson for language. Off-peak hours are not simply degraded versions of peak ones; they are qualitatively different, because the reduced inhibitory control that comes with circadian downtime widens the spread of activation in memory [83]. The counterintuitive consequence, sometimes called the *inspiration paradox*, is that tasks requiring broad associative reach can be performed *better* at one’s non-optimal time [129]. Translated into language, this predicts a daily oscillation not merely in the *amount* of linguistic resource but in its *character*: tighter, more conventional, more precisely controlled usage at peak hours, and looser, more associative, more metaphor-prone and error-prone usage off-peak. The fatigued speaker is not only a worse speaker; she is, in a definite sense, a more *creative* and less guarded one—a prediction the framework makes sharp and that time-stamped corpora could readily test.




Figure 5.2. The synchrony effect for language performance. Morning types (larks) reach peak linguistic performance in the late morning, evening types (owls) in the early evening. The same task administered at a fixed clock hour therefore falls at different points on the two curves, producing a crossover that distinguishes true circadian modulation from mere time-of-day artefact.

5.6 Seasonal and Longer Cycles

The day is the fastest of the circadian family, but not the only recurring astronomical cycle that language answers to. The same large-scale study that revealed the diurnal rhythm of affect also uncovered a *seasonal* one: across cultures and hemispheres, the baseline positivity of people’s word choice shifts with the turning of the year, tracking not the absolute length of the day but the relative rate at which daylight is waxing or waning [59]. Language, in other words, keeps not only the hours but the seasons. This is what the LTC framework expects, for the same regulated-expression logic that makes the day a linguistic cycle makes any sufficiently regular environmental period a candidate modulator of expression.

Seasonal linguistic rhythms have both a physiological and a cultural face. On the physiological side, the annual modulation of mood—in its clinical extreme, seasonal affective disorder—leaves its trace in the emotional vocabulary

people reach for at different times of year. On the cultural side, the calendar of festivals, harvests, and liturgical seasons imposes its own recurring genres of discourse: the vocabulary and register of a community are not the same in a season of mourning as in a season of celebration, and these too return with annual regularity. Longer cycles still—the generational and the historical—lie beyond the circadian family and belong to the meso- and macro-scales of later parts. The point to carry forward is that the circadian day is simply the tightest loop in a whole nested series of environmental cycles, each of which modulates the expression of the one unchanging linguistic code.

5.7 Toward a Linguistic Expression Clock

Gathering these strands, the LTC framework proposes a specific construct: a *linguistic expression clock*, the linguistic counterpart of the molecular circadian oscillator that gates gene expression across the day. The proposal is not that language has its own dedicated pacemaker, but that the general circadian and homeostatic machinery of the brain gates the expression of linguistic competence just as it gates other regulated outputs, so that each linguistic subsystem—phonology, lexicon, syntax, pragmatics—has a characteristic daily activation profile. On this view the curves of Figure 5.1 are not epiphenomena but the visible face of a real regulatory rhythm, and the research programme is to chart, for each linguistic module, its phase and amplitude across the day and across chronotypes.

Such a programme has both basic and applied promise. It would refine our theories of the competence–performance distinction by specifying, quantitatively, how much of “performance” is clock-driven. And it points toward time-aware language technologies—predictive text, tutoring systems, and clinical tools that adapt to the user’s circadian state—explored further in Part V. One question, however, this chapter has deferred. Sleep pressure, the rising dashed curve of Figure 5.1, is dissipated only by sleep; and sleep does far more for language than relieve fatigue. It is during sleep that the day’s linguistic experience is consolidated into lasting knowledge. That nightly half of the cycle is the subject of the next chapter.

5.8 Summary

Language keeps the hours of the day. The circadian and homeostatic machinery described by the two-process model gates not only general cognition but the expression of linguistic competence: syntactic processing is sensitive to circadian phase and sleep pressure, chronotype modulates activity in language regions of the brain, and the emotional colouring of word choice traces a robust diurnal curve across hundreds of millions of messages. Chronotype shifts these

patterns along the clock and, through the synchrony effect, offers a lever for disentangling clock time from circadian phase. The LTC framework gathers these findings under the construct of a linguistic expression clock and turns them into a programme of quantitative, testable predictions. What the waking day builds, the sleeping night consolidates—the theme to which we now turn.

Coda: Linguistics and the Brahma Kumaris

No point of contact between this book and the Brahma Kumaris teaching is sharper than the one this chapter reaches. The tradition singles out a particular hour of the circadian day, *Amrit Vela*—the “nectar hour” before dawn, around four in the morning—as the time when the mind is stillest and remembrance deepest, and it builds its central daily practice upon that window [23, 128]. This is, in the framework’s terms, a claim about the linguistic and cognitive expression clock: that there is an optimal circadian phase for the mind’s finest work. Where chronobiology arrives at time-of-day effects through melatonin and the two-process model, the Brahma Kumaris arrived, by contemplative observation, at a practical schedule keyed to the same daily rhythm. That a spiritual tradition and a laboratory science should both privilege particular hours is, for a temporal linguistics, a convergence worth noticing.

Sleep, Consolidation, and the Linguistic Clock

We learn a language twice over: once by day, in the hearing of it, and again by night, when sleep quietly files away what the day let in.

The previous chapter left one curve unexplained. Homeostatic sleep pressure, the rising dashed line of Figure 5.1, is discharged only by sleep—and sleep, it turns out, does far more for language than relieve the day’s fatigue. The hours we spend unconscious are not a blank in the linguistic cycle but its most productive phase, the time when the raw material gathered by day is sorted, stabilised, and woven into the standing fabric of what a speaker knows. If the waking day is when language is *expressed*, the sleeping night is when it is *consolidated*. This chapter follows the linguistic cycle across the threshold of sleep and argues that the nightly reorganisation of memory is an indispensable turn of the LTC wheel—the mechanism by which the fast events of one day become the slow knowledge of a lifetime.

6.1 The Architecture of a Night’s Sleep

Sleep is not a uniform state but a structured progression through distinct stages, cycling several times across a night (Figure 6.1). After drifting down through light stages N1 and N2, the sleeper descends into deep *slow-wave sleep* (SWS, stage N3), marked by large, synchronised cortical oscillations; later comes *rapid-eye-movement* (REM) sleep, in which the brain is active and dreaming is vivid. The crucial fact is that these stages are not evenly distributed: slow-wave sleep dominates the first half of the night, while REM periods lengthen toward morning. Decades of research have established that this architecture is functional, not incidental. The modern consensus, that sleep plays an *active* role in consolidating memory rather than merely protecting it from interference, assigns particular importance to slow-wave sleep for the stabilisation of newly acquired declar-




Figure 6.1. A typical hypnogram across a night’s sleep. Slow-wave sleep (stage N3) is concentrated in the first half of the night, when word-forms are stabilised and integrated into the lexicon; REM periods (highlighted) lengthen toward morning and are associated with abstraction and grammatical generalisation. The nightly cycle thus has its own internal division of linguistic labour.

ative knowledge [98]. Since a language is, in large part, an enormous body of declarative knowledge—tens of thousands of words and their properties—this at once makes sleep a linguistic concern.

6.2 Sleep Consolidates the Words of the Day

The clearest linguistic evidence comes from the learning of new words, and it is worth stating carefully because it isolates sleep from the mere passage of time. A newly heard word-form is learned almost instantly—one can repeat it at once—but *integrating* it into the mental lexicon, so that it competes with similar-sounding words during recognition the way established words do, takes longer, and crucially it takes *sleep*. In an elegant demonstration, Dumay and Dumay and Gaskell [47] taught novel words such as *cathedruke* and showed that they began to inhibit recognition of their real neighbours (*cathedral*) only after

Table 6.1. What sleep does for different kinds of linguistic material, the sleep stage most implicated, and representative evidence. The pattern suggests a division of labour between slow-wave and REM sleep that parallels the lexicon–grammar division in language itself.

Material	Effect of sleep	Stage	Evidence
New word-forms	lexical integration; competition emerges	SWS	Dumay & Gaskell
L2 vocabulary	better retention; cueing boosts it	non-REM	Schreiner & Rasch
Grammatical rules	preferential consolidation of regularities	SWS	Mirković & Gaskell
Chunked L2 input	associative learning enhanced	—	Monaghan et al.
Declarative memory	active system consolidation	SWS	Rasch & Born

a night’s sleep, not after an equivalent interval of daytime waking. The new word does not merely persist across the night; it is *knitted into* the lexical network during it. Sleep is the moment the lexicon is updated.

That the effect is caused by sleep, and not merely correlated with it, has been shown by intervening in sleep directly. Schreiner and Schreiner and Rasch [112] had learners study foreign vocabulary and then, while they slept, softly replayed some of the learned words during non-REM sleep. Memory for the cued words was measurably better afterwards than for uncued ones—a benefit that appeared only when the words were replayed *during sleep*, not during waking. Such *targeted memory reactivation* shows that the sleeping brain is actively rehearsing the day’s linguistic material, and that this rehearsal can even be nudged from outside. Table 6.1 summarises these and related findings.

6.3 Why Language Needs Two Memory Systems

Why should integration require sleep at all? Why not simply store each new word immediately and completely? The answer, developed in the theory of *complementary learning systems*, is one of the most illuminating ideas in cognitive science for a temporal linguistics [84]. The brain, on this account, cannot use a single memory store, because the two things it must do are in tension. To learn a specific new item in one exposure requires a fast store that can change rapidly—the hippocampus. But to extract the statistical regularities of a domain, and to fold new knowledge into old without catastrophic interference, requires a slow store that changes only gradually—the neocortex. The resolution is to have




Figure 6.2. The complementary learning systems account of overnight lexical integration. A new word is captured quickly by the hippocampal fast store during the day; in sleep it is replayed and transferred into the neocortical slow store, where it is interwoven with existing knowledge and begins to compete with its lexical neighbours.

both: the hippocampus captures the day's episodes quickly, and then, offline during sleep, gradually trains the neocortex by replaying them, interleaving the new with the old (Figure 6.2).

For language this architecture is not optional but necessary. A speaker must be able to learn a word from a single encounter, yet must also maintain a lexicon in which tens of thousands of words are smoothly interrelated by sound and meaning. The two demands are exactly those the complementary systems evolved to reconcile. The overnight lexical integration observed by Dumay and Gaskell is precisely what the theory predicts: the fast hippocampal trace of the new word is, during sleep, transferred into and interwoven with the slow neocortical lexicon, at which point it begins to behave like a native resident of the language. Sleep is the nightly session in which the day's words are naturalised.

6.4 Does Sleep Improve Your Grammar?

Words are one thing; rules are another, and here sleep appears to do something subtler than mere stabilisation. When learners are exposed to a novel miniature language and then tested after sleep, what improves is not uniform: sleep

preferentially consolidates certain components of the new system. In a study whose title poses our question directly, Mirković and Mirković and Gaskell [87] found that sleep favoured the retention of the more arbitrary, item-based aspects of a newly learned grammar. More broadly, sleep is repeatedly found to promote *abstraction*: the extraction of the general pattern latent in a set of examples, so that sleepers awaken having grasped a regularity they could not state the night before. Monaghan and Monaghan et al. [88] similarly show that sleep aids the chunking and associative learning on which second-language acquisition depends.

The effect is visible at the very dawn of language. In a striking demonstration with fifteen-month-old infants, Gómez, Bootzin, and Gómez et al. [61] exposed babies to an artificial language and then let some nap and kept others awake. The infants who stayed awake remembered the specific word sequences they had heard; the infants who napped did something more sophisticated—they abstracted the underlying *pattern*, generalising to novel strings that obeyed the same rule. A single nap converted rote memory into an abstract generalisation. That the machinery of sleep-dependent abstraction is already running in the second year of life, precisely when grammar is being acquired, is a powerful hint that sleep is woven into language acquisition from the start.

This hints at a striking parallel. The nightly cycle appears to divide its linguistic labour much as language itself is divided: slow-wave sleep, early in the night, stabilising the particulars—the word-forms, the arbitrary items—and REM sleep, later, favouring the abstraction of general structure, the raw material of grammar. The lexicon–grammar distinction that organises linguistic theory may have, in the architecture of a single night’s sleep, a temporal echo. The claim remains tentative, but it is exactly the kind of cross-scale resonance the LTC framework is built to notice.

6.5 Sleep Loss and the Fraying of Language

If sleep builds language, sleep’s absence should visibly erode it—and it does. The effects of sleep deprivation fall hardest on exactly the linguistic functions that depend on the prefrontal cortex, the region most vulnerable to sleep loss. Harrison and Harrison and Horne [66] found that a night without sleep significantly impaired novel, non-routine language tasks—generating words to a constraint, suppressing a habitual response in favour of an apt one—while leaving well-practised, automatic language relatively intact. The sleep-deprived speaker does not lose her grammar, but her language grows more clichéd, more perseverative, less flexible and less apt: the creative, effortful edge of language dulls first.

This is the mirror image of consolidation, and it completes the case that language is genuinely coupled to the sleep–wake cycle rather than merely coincident with it. The same nightly process that, when it runs, integrates new words and abstracts new rules is the process whose absence leaves expression frayed and formulaic. Both directions—the gains from sleep and the losses from its denial—point to the same conclusion: the offline hours are not linguistically inert, and a full account of language in time must give the night its due.

6.6 The Nightly Reorganisation and the Expression Clock

We can now complete the picture of the daily linguistic cycle begun in Chapter 5. The waking day, gated by the circadian and homeostatic rhythms of the two-process model, is the phase of *expression*: competence is read out into speech, variably and by the hour, and new linguistic material is taken in and held provisionally by the fast store. The sleeping night is the phase of *consolidation and reorganisation*: sleep pressure discharges, the day’s episodes are replayed, new words are integrated into the lexicon, and general patterns are abstracted toward grammar. The two phases are not separate systems but the two strokes of a single engine—the linguistic counterpart of the way a cell alternates between the expression of its genes and the repair and reorganisation of its genome across the circadian cycle.

This is what the framework means by proposing that sleep functions as a linguistic “reorganisation phase,” and it gives the construct of a linguistic expression clock its second hand. The clock does not merely modulate output across the waking day; it also schedules, in the offline hours, the updating of the very competence that will be expressed tomorrow. The practical corollary is significant: because the reorganisation can be influenced from outside—by the timing of study relative to sleep, by napping, and even by targeted reactivation during sleep itself [19, 112]—the nightly phase of the cycle is a point of leverage for language learning and rehabilitation, a theme Part V will develop.

6.7 Summary

Sleep is the other half of the daily linguistic cycle, and the more constructive half. Its structured stages do real linguistic work: slow-wave sleep stabilises new word-forms and integrates them into the lexicon, as shown by the emergence of lexical competition only after sleep, while later sleep favours the abstraction of grammatical regularities. The complementary-learning-systems architecture explains why this offline consolidation is necessary—a fast hippocampal store for one-shot learning, a slow neocortical store trained during sleep for smooth integration—and why language, needing both, depends on the nightly transfer

between them. Read within the LTC framework, waking expression and sleeping consolidation are the two strokes of one engine, and sleep is the reorganisation phase of the linguistic expression clock. Having traced the micro-cycle of the day in full, we turn in the next chapter to the fast structure of language in time—to recursion and repetition as a temporal syntax.

Coda: Linguistics and the Brahma Kumaris

This chapter cast sleep as the mind’s nightly reorganisation. The Brahma Kumaris daily discipline frames the same boundary of the day in complementary terms: the practice of reviewing the day’s thoughts and actions before rest, and of rising at *Amrit Vela* to reorder the mind before the day begins, treats the transition into and out of sleep as a time for settling and re-sorting the *sanskars* [23, 128]. Where cognitive science describes an offline consolidation that integrates the day’s words and abstracts its patterns, the tradition prescribes a deliberate offline hygiene aimed at the same boundary—the hours when the conscious self loosens its grip and the deeper impressions settle. The convergence is not on mechanism but on *timing*: both identify the edges of sleep as the moments when what the day deposited is quietly filed away.

Recursion and Repetition as Temporal Syntax

To nest is to fold time inward; to repeat is to lay it out in a line. Language does both, and in the doing keeps its rhythm.

The two preceding chapters followed language through the outer cycle of the day—its waking expression and its sleeping consolidation. This chapter turns inward, to the fast temporal structure of language itself, and to two devices by which that structure is organised in time. The first is *recursion*, the nesting of a structure inside another of the same kind, which folds time inward into layers. The second is *repetition*, the return of the same form at intervals, which lays time out along a line. They look like opposites—depth against periodicity—but both are, at root, ways of using time to build meaning, and both connect the fast oscillations of Chapter 4 to the larger cycles of the book. Together they constitute what we may call a *temporal syntax*.

7.1 Two Faces of Temporal Structure

Consider two famous sentences. Caesar’s “I came, I saw, I conquered” places three clauses of identical shape side by side; its power is periodic, a beat repeated. A center-embedded sentence such as “the rat the cat the dog worried chased fled” places three clauses *inside* one another; its difficulty is one of depth, of holding an outer structure open while an inner one is completed. Figure 7.1 sets the two against each other. Recursion runs vertically, nesting structure within structure; repetition runs horizontally, returning the same structure again and again. The remainder of the chapter treats each in turn and then shows how, at the fast scale, they meet—in meter, in mantra, and in the priming that makes speakers echo themselves without knowing it. Table 7.1 previews the devices to be discussed and the temporal character of each.




Figure 7.1. The two faces of temporal syntax. *Recursion* (left) nests a structure inside another of the same kind, folding time inward into layers of dependency. *Repetition* (right) returns the same form at fixed intervals, laying time out as a periodic chain. The one builds depth; the other builds rhythm.

7.2 Recursion: Structure Nested in Time

Recursion has a strong claim to be the defining formal property of human language. On the influential proposal of Hauser, Chomsky, and Hauser et al. [67], the uniquely human core of the language faculty may reduce to precisely this: the capacity to embed structures within structures without limit, generating an infinite range of expressions from finite means. A relative clause may contain a relative clause, a sentence may be quoted inside a sentence, and in principle the nesting never has to stop. In temporal terms, recursion is a linguistic *feedback loop*: the output of a rule becomes the input to the same rule, so that the structure re-enters itself. It is the syntactic image of a cycle that contains a smaller copy of itself.

This self-embedding gives language a distinctive expressive power: the ability to represent time *about* time, reasoning nested within reasoning. A sentence

Table 7.1. Devices of recursion and repetition across levels of linguistic structure, with their temporal character and characteristic function.

Device	Level	Temporal type	Function
Reduplication	morphology	local repetition	intensity, plurality
Center-embedding	syntax	nesting (recursion)	clause within clause
Anaphora	rhetoric	periodic return	emphasis, cohesion
Parallelism	poetics	periodic	the poetic function
Meter	phonology	oscillation	the clock of verse
Refrain	discourse	periodic return	memory, unity
Mantra	ritual	slow oscillation	entrainment
Structural priming	processing	short-term recurrence	persistence

such as “every time I try, I fail because I try too hard” folds a temporal generalisation inside a causal claim inside a habitual frame; recursion is what lets a speaker hold such nested temporalities in a single utterance. The device that builds grammatical depth is thus also the device that lets language model the layered structure of time itself—a small but telling instance of the book’s larger theme that linguistic form and temporal structure mirror one another.

The claim that recursion is universal has been vigorously contested—most prominently by Everett’s argument that the Amazonian language Pirahã manages without it [53]—and the debate remains open. But whatever its status as a universal, recursion is unquestionably *available* to human language, and its temporal signature is what concerns us: it is the device by which language escapes the tyranny of the line, allowing a speaker to suspend one thought inside another and return to it. That return, however, is not free. It must be paid for in time and in memory.

7.3 The Real-Time Limits of Recursion

Recursion is unbounded in the grammar but sharply bounded in performance, and the gap between the two is one of the clearest cases in all of linguistics of the competence–performance distinction playing out *in time*. A singly center-embedded sentence is easy; a doubly embedded one is hard; a triply embedded one is very nearly unintelligible, though the grammar sanctions it perfectly. Gibson’s *dependency locality theory* explains why in explicitly temporal terms:

holding an unfinished dependency open across intervening material carries a *storage cost*, and completing it across a distance carries an *integration cost*, both of which grow as the nesting deepens and the relevant elements are pushed further apart in time [56]. Center-embedding fails not because it is ungrammatical but because it overruns the temporal capacity of working memory. Recursion, in short, is a competence that unfolds against a clock, and the clock imposes a depth limit the grammar itself does not. This is the coupling of Chapter 4 seen from the inside: an unbounded structural possibility, gated by the real-time resources of the micro-scale.

7.4 Repetition as Temporal Marker

If recursion folds time, repetition measures it. Far from being the flaw that prescriptivists deplore, repetition is a pervasive and functional organiser of discourse. Deborah Tannen [122] showed that ordinary conversation is saturated with repetition—of sounds, words, and syntactic frames—and that this repetition does essential work: it binds discourse into cohesive wholes, signals engagement, buys planning time, and creates the interpersonal rhythm within which understanding happens. At the level of rhetoric, *anaphora*—the repetition of an opening phrase across successive clauses, as in “we shall fight on the beaches, we shall fight on the landing grounds”—turns repetition into force. At the level of poetics, Roman Jakobson [72] made repetition constitutive: the *poetic function* of language, he argued, projects equivalence onto the chain of speech, so that parallelism—the recurrence of matched structures—is the very principle of verse.

Repetition is also the linguistic substance of ritual and identity. Chants, slogans, creeds, and mantras all recruit the mnemonic and affective power of the returning phrase; the form that comes back is the form that is remembered, and the form that is remembered binds a community across time. Here the horizontal periodicity of repetition reaches up from the micro-scale of the utterance to the macro-scale of cultural continuity—a first hint of the coupling that Part IV will develop in full.

7.5 Meter and Mantra: Language as Oscillator

Nowhere is the temporal character of repetition sharper than in metered language, where it becomes a genuine *oscillation*. A line of iambic pentameter is a sequence of five weak–strong feet, a periodic alternation of prominence that the ear tracks as a beat (Figure 7.2). Meter is, quite literally, a linguistic clock: it imposes a fixed period on the stream of syllables and synchronises producer and

listener to it, exactly the entrainable oscillator that Chapter 4 found at the base of speech rhythm, here regularised and made aesthetic.


The oscillation is not merely metaphorical, and it couples to the body. In a remarkable study, Bernardi and colleagues found that reciting a rosary or a yoga mantra slows the breath to almost exactly six cycles per minute, and that at this rate the recitation drives powerful, synchronised oscillations in heart rate and blood pressure, entraining the cardiovascular system to the rhythm of the words [12]. Ritual repetition, in other words, is a physiological entrainment device: the returning phrase pulls the body's own slow rhythms into lock with it. The mantra is language used deliberately as an oscillator, and its effect is measured in the pulse. In the vocabulary of this book, meter and mantra are micro-scale linguistic cycles engineered to entrain the still slower cycles of the body—rhythm reaching across scales.

Meter and rhyme are also, and not coincidentally, memory technologies. Long before writing, epics and genealogies were carried across generations by word of mouth, and they survived because their form constrained their content: a line that must scan and rhyme cannot drift far without breaking, so the metrical and repetitive scaffolding of oral verse actively corrects transmission errors and holds the text stable through time [106]. Repetition here does at the scale of centuries what priming does at the scale of seconds—it makes the near future resemble the recent past—and it is one of the mechanisms by which a purely spoken tradition achieves the fidelity that Part IV will require of linguistic transmission. The rhythm of the mantra and the meter of the epic are, in this light, not ornament but the very means by which language endures without a script.

7.6 Structural Priming: Repetition Below Awareness

Repetition is not always chosen. Some of the most revealing repetition in language happens beneath awareness, in the phenomenon of *structural priming*. Bock's classic experiments showed that having just produced a sentence in one syntactic frame—a passive, say, or a double-object dative—a speaker is measurably more likely to reuse that same frame for an unrelated message moments later, even when she attends neither to the structure nor to the connection [17]. Syntactic structure *persists*: it echoes forward in time within a single mind. This is the intrapersonal counterpart of the interactive alignment of Chapter 4, in which structure echoes *between* minds [94], and it reveals repetition as a basic temporal property of the production system itself—a short-horizon memory that biases the near future toward the recent past. The tendency of language to repeat itself, that is, is built into the machinery, not merely into the rhetoric.

METER AS A LINGUISTIC OSCILLATOR (iambic pentameter)



one weak-strong foot = one period; five feet per line

Figure 7.2. A line of iambic pentameter as an oscillation. Each weak–strong foot is one period of a periodic prominence contour; the stressed syllables (S) fall on the peaks, the weak syllables (w) in the troughs. Metered verse imposes a fixed temporal period on speech and entrains speaker and listener to it.

7.7 Computational Models: Recursion versus Iteration

The distinction between the two faces of temporal syntax is made exact by formal language theory. In the Chomsky hierarchy, mere *repetition*—a fixed pattern looped any number of times—is within the power of a *finite-state machine*, an automaton that cycles among a fixed set of states with no memory of how it got there. True *recursion*, by contrast, with its unbounded nesting of matched dependencies, demands at least a *push-down automaton*, a machine equipped with a stack on which to store the open structures awaiting completion. The stack is memory laid out in time, and its finite depth in any real system is the formal echo of the center-embedding limit of Section 7.3.

Modern connectionist models make the same contrast in a different guise. Simple recurrent networks, which feed their own hidden state back as input, were introduced precisely to let a network “find structure in time” [51]; their descendants, the long short-term memory networks of Hochreiter and Hochreiter and Schmidhuber [69], added a gated memory able to hold information across longer temporal gaps—an engineered solution to exactly the storage problem that center-embedding poses to human speakers. Whether symbolic or neural, the formal lesson is the same: repetition is cheap, requiring only a loop, while

recursion is expensive, requiring memory that persists through time. The two faces of temporal syntax are, at bottom, two different demands on the same scarce resource—time-bound memory.

7.8 Summary

Recursion and repetition are the two ways language structures itself in time. Recursion folds time inward, nesting structures within structures; it is the formal signature of human syntax, but it runs against a real-time memory limit that caps how deep the nesting can go. Repetition lays time out in a line, returning the same form at intervals; it organises conversation, powers rhetoric and poetry, binds ritual communities, and, in meter and mantra, becomes a genuine oscillator that can entrain the rhythms of the body. Beneath awareness, structural priming shows the production system repeating itself as a matter of course, and formal language theory locates the whole contrast in a single trade-off between looping and memory. Both devices connect the fast oscillations of speech to the larger cycles of culture and cognition—temporal syntax as the hinge between the scales.

Coda: Linguistics and the Brahma Kumaris

This chapter's two devices both find a home in the Brahma Kumaris practice, and in its cosmology. Repetition is the very texture of spiritual discipline: the daily return to the *Murli*, the continual practice of remembrance (*yaad*), and the use of short repeated affirmations such as *Om Shanti* all work by the mnemonic and entraining power of the returning phrase—the same power that Bernardi measured in the slowing pulse of the mantra reciter [12, 23]. The tradition, that is, uses repetition deliberately as a temporal technology of the mind, tuning the inner rhythm by the outer word. Recursion, in turn, has a cosmological echo: the *kalpa* is a cycle that contains, in its predestined Drama, a perfect smaller copy of itself each time it turns—a world that re-enters itself exactly, the recursive figure written at the largest scale [128]. Where the linguist sees a clause nested in a clause and a beat returning in a line, the tradition sees a cycle nested in eternity and a Drama returning without end; the temporal syntax of language and the temporal structure of the cosmos are, on this reading, the same two figures—nesting and return—seen at incommensurable scales.

Part III

The Meso-Cycles: Lifespan and Society

Social Time and Register Shifts

We do not speak one language but many, and we choose among them by the clock of our lives and the calendar of our society.

Part II followed language through the micro-cycles of the day and the fast structure of the utterance. Part III now steps up to the meso-scale—the middle band of the spectrum, measured in years and decades rather than milliseconds and hours. Its governing idea is *social time*: the position a speaker occupies within her own lifespan, her generational cohort, and the recurring occasions of her society. This chapter argues that much of what looks like free stylistic variation is in fact a systematic response to social-temporal position, and that the mechanism by which speakers make that response—the *register shift*—is the meso-scale’s characteristic act, the way language adapts to *when*, in the social sense, it is being spoken.

8.1 Register as Adaptation to the Social Moment

No competent speaker uses one invariant style. The same person speaks differently to a child, a colleague, and a judge; differently at a funeral and at a football match; differently when relaxed and when on guard. This variation is a *register shift*—a coordinated adjustment of vocabulary, pronunciation, syntax, and formality to suit the situation. Early variationist work treated style as a function of the *attention* a speaker pays to her own speech, more standard forms emerging under greater self-monitoring. Allan Bell [11] reframed the matter as *audience design*: speakers shape their style above all for their listeners, adjusting most for those they address directly and less for those who merely overhear. Halliday’s account of register as the language appropriate to a configuration of field, tenor, and mode gave the notion its functional grounding [64], and Eckert’s later work recast style as active social *practice*—not a passive reflex of the situation but a resource speakers deploy to construct who they are [50]. Table 8.1 sets out the principal registers this chapter treats, each keyed to a characteristic temporal frame.

Table 8.1. Registers and the social-temporal frames to which they are adapted, with their characteristic linguistic form.

Register	Temporal frame	Linguistic character
Child-directed speech	early lifespan	simplified syntax, exaggerated prosody
Adolescent vernacular	adolescence	slang, identity marking
Professional standard	adult worklife	formal, careful, prestige forms
Elder speech	late lifespan	conservative, sometimes archaic
Ritual register	ritual time	elevated, archaic, formulaic
Crisis communication	crisis time	simplified, repetitive, deictic urgency
Digital / viral	“viral time”	compressed, hybrid, rapidly renewed

8.2 Language Across the Lifespan

The slowest register cycle is the lifespan itself, over which a speaker’s language changes in patterned ways. The developmental arc is by now well charted, and one feature of it is among the most robust findings in all of sociolinguistics: the *adolescent peak*. Use of vernacular, non-standard forms rises through childhood to a maximum in the teenage years—when peer identity is paramount and the standard language least valued—and then falls sharply as young adults enter the workforce and the pressures of what Eckert calls the “linguistic marketplace,” before often rising again in retirement as marketplace pressures relax [49]. Figure 8.1 sketches this curvilinear pattern. It is a genuinely *cyclic* phenomenon in the book’s sense: each speaker traverses the same curve, so the pattern recurs, lifespan after lifespan, as a stable feature of the community even though no individual stays fixed.

This lifespan curve poses a famous methodological puzzle that is really a question about time. If we survey a community once and find that older speakers use a form less than younger ones, are we seeing *age-grading*—a difference that each cohort will repeat as it ages, so the community is stable—or a *change in progress*, in which the young are carrying a genuine innovation that will persist as they grow old? The *apparent-time* construct assumes the latter, reading the distribution across ages as a snapshot of change through time. Testing that assumption requires real-time data, returning to the same community decades later. Sankoff and Sankoff and Blondeau [110], restudying the pronunciation of




Figure 8.1. The curvilinear age pattern in the use of vernacular, non-standard forms. Use peaks in adolescence, declines to a minimum in the working years under the pressures of the “linguistic marketplace,” and often rises again after retirement. Each speaker traverses the same curve, so the pattern recurs across generations (after 49).

/r/ in Montreal French across three decades, largely vindicated the apparent-time method: individuals were mostly stable across their lifespans, and when they did change, they moved in the direction the community was moving. The meso-scale, in short, can be read—with care—from a single synchronic slice, because the lifespan cycle and the historical change are systematically related.

8.3 Ritual and Institutional Temporalities

Beyond the lifespan, society structures time into recurring occasions, and many of these summon registers of their own. *Ritual time*—the wedding, the funeral, the court session, the religious service—characteristically calls forth elevated, archaic, and formulaic language: the “thee” and “thou” of the liturgy, the fixed verbal formulae of the marriage vow, the ritualised phrasing of the courtroom. This linguistic conservatism is not accidental. Formulaic and archaic language reinforces collective identity and continuity, and its very resistance to change

is a source of authority: to speak the old words in the old way is to invoke a tradition that stands outside ordinary time. Institutions preserve these registers deliberately, and the registers in turn preserve the institutions. Ritual language also has a peculiar temporal logic of its own: much of it does not describe the world but *acts* upon it. Austin's analysis of the *performative* utterance turns precisely on such ritual formulae—"I now pronounce you married," "I hereby name this ship"—which bring about the state of affairs they name, provided they are spoken by the right person, in the right place, at the right moment [6]. The efficacy depends on exact repetition: the words must be the received words, or the act misfires. This is why ritual registers are so resistant to change; unlike ordinary speech, where paraphrase is harmless, ritual speech binds its power to a fixed form that must recur unaltered occasion after occasion—a temporal anchor holding a fixed verbal form steady across the generations a community passes through.

Figure 8.2 places such ritual and institutional registers, along with the digital ones to come, on a landscape defined by formality and by temporal tempo.

8.4 Digital Chronotopes and Register Fluidity

At the opposite temporal extreme lies the register ecology of digital communication, where the meso-scale runs unusually fast. Online, speakers code-switch fluidly and often within minutes—from a formal email to a casual text to the compressed, playful idiom of a meme—so that a single hour may span registers that once belonged to widely separated occasions. New time-sensitive registers arise and spread with remarkable speed: the "Zoom speak" and the epidemiological vocabulary of a pandemic, the ever-renewing lexicon of a social platform. Far from the "linguistic ruin" that alarmists predicted, close study of these registers reveals them to be structured and rule-governed: Tagliamonte and Tagliamonte and Denis [121], examining teenagers' instant messaging, found it to be a distinctive hybrid of speech and writing with its own systematic patterns, neither spoken nor written but a new register keyed to a new temporality [36]. Digital communication compresses the meso-scale, layering registers that coexist in real time and turning register-switching from an occasional adjustment into a continuous, moment-to-moment practice.

8.5 Linguistic Capital and the Marketplace

Behind the lifespan curve and the choice of register lies an economy. Pierre Bourdieu argued that language is never merely communicative but always also a medium of *symbolic power*: particular ways of speaking carry *linguistic capital*, a value assigned by a "linguistic market" that rewards the legitimate,




Figure 8.2. A landscape of registers, arranged by formality (vertical) and by temporal tempo (horizontal). Ritual and institutional registers cluster in the slow, formal quadrant—their power lies in endurance; digital and viral registers occupy the fast, casual quadrant—their life is measured in hours or days. Register choice is, among other things, a choice of temporal stance.

prestige variety and penalises the stigmatised one [21]. This is the machinery beneath Eckert’s “linguistic marketplace.” The sharp decline of vernacular use when young adults enter the workforce (Figure 8.1) is not a change of taste but a response to a market: the standard variety is the currency of professional advancement, and speakers invest in it accordingly. The relaxation of the curve in retirement is the same market ceasing to press.

The marketplace gives the meso-scale its motive force. Register shifts are not free choices but positioned moves within a field of unequal linguistic value, and the value itself is set socially and shifts over historical time—today’s prestige form was yesterday’s innovation and may be tomorrow’s cliché. The market thus couples the scales: it translates the macro-scale facts of social structure and linguistic history into the micro- and meso-scale choices of individual speakers, who feel it as the simple sense that some ways of speaking “sound right” for

some occasions and wrong for others. It is when that sense of value shifts suddenly, across a whole community at once, that the meso-scale erupts into a phase transition.

8.6 Sociolinguistic Phase Transitions

Occasionally the slow accumulation of register change gives way to something abrupt. Social upheavals—revolutions, wars, pandemics, the arrival of a new communicative technology—can trigger rapid, wholesale transformations of the register system: the emergence of new public styles and protest languages, the sudden breakdown of formality in media and politics, bursts of lexical borrowing and coinage to meet new conditions. These resemble *phase transitions* in physical systems, or population bottlenecks in genetics: a system that has drifted slowly for a long time reorganises quickly once a threshold is crossed. The mathematics of Appendix C makes the analogy precise. In the Kuramoto model of coupled oscillators, a population of speakers shifts from incoherence to collective synchrony not gradually but suddenly, once the coupling exceeds a critical value. A sociolinguistic phase transition is the linguistic image of that bifurcation: the moment when a critical mass of speakers locks onto a new norm and the register landscape is redrawn. Here the meso-scale hands upward to the macro-scale, for a register revolution, once complete, is the beginning of a lasting change—the subject of Part IV.

8.7 Summary

Language at the meso-scale is governed by social time. Speakers command not one style but a repertoire of registers, and they shift among them in systematic response to audience, situation, and social-temporal position—most slowly across the lifespan, whose curvilinear age pattern recurs generation after generation and whose relation to historical change underwrites the apparent-time method; more sharply across the recurring occasions of ritual and institution, whose conservative registers trade on endurance; and most rapidly of all in digital communication, where registers are layered and renewed by the hour. When social change is abrupt, register change can be abrupt too, in sociolinguistic phase transitions that carry the meso-scale up into the macro. Register shifting is thus the meso-scale's signature act—the continual retuning of language to the social moment.

Coda: Linguistics and the Brahma Kumaris

The Brahma Kumaris community offers a vivid illustration of register as social-temporal practice, for entry into the tradition is marked by a pronounced register shift. To become a *Brahmin*—in the tradition’s terms, a child of the Confluence Age—is to adopt a whole new linguistic repertoire: the elevated vocabulary catalogued in Appendix E, the forms of address (*Baba, Dadi*), the greeting *Om Shanti* that reaffirms spiritual identity at every meeting, and the distinctive directive register of *Shrimat*, God’s elevated instructions, set explicitly against the *manmat* of one’s own mind [23, 128]. Sociologists of the movement have noted precisely this constructive use of language in the making of a new self [126]—register as the instrument of identity, exactly as Eckert describes. There is a temporal resonance, too, in the tradition’s own framing. It locates its community at the *Sangam Yug*, the Confluence Age—the hinge between the old cycle and the new—and understands the turn of the cycle as the most radical of all transitions. In the vocabulary of this chapter, the Confluence Age is a sociolinguistic phase transition raised to a cosmic key: the moment when a community locks onto a wholly new register because it understands itself to stand at the turning of time.

Acquisition, Transmission, and Linguistic Heredity

A language is never inherited whole; it is rebuilt, from scratch, in every child—and rebuilt, each time, a little changed.

Chapter 8 traced how a speaker’s language shifts across her lifespan. This chapter asks the deeper meso-scale question: how does a language pass from one lifespan to the next at all? In the genetic analogy of Chapter 2, this is the domain of *heredity*—but it is precisely here, we shall see, that the analogy is most instructively strained. A genome is copied and handed on; a language is not. It is *reconstructed*, anew, by each generation of learners, on the evidence of the utterances they happen to hear. That difference—transmission by reconstruction rather than by copying—turns out to be not a weakness of language but the very engine of its structure, and the source of some of the most remarkable phenomena in the human sciences: the spontaneous birth of grammar where none existed before.

9.1 The Timetable of Acquisition

Every normally developing child acquires a first language on a schedule so regular that it can be plotted (Figure 9.1). Cooing gives way to babbling around six months, first words appear near the first birthday, a vocabulary “spurt” follows around eighteen months, two-word combinations around two years, and the machinery of morphology and complex syntax is largely in place by four or five—all on the basis of input that is fragmentary, unlabelled, and never systematically corrected. This regularity is the meso-scale’s developmental clock, and it runs within a window. Lenneberg’s *critical period* hypothesis held that the capacity for effortless first-language acquisition is greatest in early childhood and declines through puberty [81], and the general finding has held up: those deprived of language until after this window never fully attain native




Figure 9.1. The developmental timetable of first-language acquisition, set against the critical-period window. Milestones from babbling to complex syntax cluster in the first five years; the underlying capacity to acquire a first language (shaded) is high in early childhood and declines through puberty. Acquisition is a clock running inside a window.

competence, and second-language learners who begin later rarely reach the fluency of those who begin early [90].

Newport observed a striking twist in this maturational story, the “less is more” effect: the very cognitive *limitations* of the young child—her smaller memory and narrower attention—may *help* acquisition, by forcing her to break the input into small, analysable units rather than storing it whole [90]. The immaturity of the learner is not only an obstacle overcome but, in part, a condition of success—a first hint that the properties of the *learner*, and not only of the input, shape what a language becomes.

9.2 Learning from the Stream

Before a child can learn words she must find them, for speech arrives as an unbroken stream with no reliable pauses between words. How the infant solves

this segmentation problem reveals the temporal character of acquisition. In a landmark experiment, Saffran, Aslin, and Saffran et al. [109] exposed eight-month-old infants to two minutes of a continuous, monotone stream of nonsense syllables in which the only cue to word boundaries was *transitional probability*—syllables within a “word” followed one another reliably, while those spanning a boundary did not. The infants detected the difference, segmenting the stream on statistical evidence alone. This *statistical learning* is a temporal computation par excellence: the infant is tracking, across time, which sounds predict which, and using the troughs in predictability to cut the stream into units. The raw material of acquisition is the sequential, probabilistic texture of experience—time turned into structure.

9.3 Iterated Learning: Transmission as a Structuring Force

We can now state the chapter’s central idea. Because each generation learns the language afresh from a limited sample of the previous generation’s speech, transmission passes through a *bottleneck*: no learner ever hears the whole of her language, only a finite and partial sample of it. To use the language beyond that sample—to say new things—the learner must *generalise*, and generalisation favours whatever is regular, patterned, and compositional, because regular systems are precisely the ones that can be inferred from a part and extended to the whole. Repeated over generations, this pressure accumulates: the language reorganises itself, without anyone intending it, into an ever more structured and learnable form.

This is not merely a computer model. Kirby, Cornish, and Kirby et al. [74] ran the process in the laboratory, having human participants learn a small artificial language and then using their output as the input for the next participant, in a chain of “generations.” Exactly as predicted, the languages evolved, over the chain, from unstructured jumbles into compositional systems—systems in which parts of words came to signal parts of meanings—becoming both more structured and easier to learn, purely as a consequence of being transmitted [73]. Structure is thus not only stored in the head and imposed on the language; it also *emerges from* the act of transmission itself. Table 9.1 sets this mode of transmission alongside the others treated in this chapter.

Why should generalisation yield *compositional* structure in particular? The iterated-learning experiments give a subtle answer: structure emerges under the joint pressure of two competing demands. A language must be *learnable*—compressible enough to pass the bottleneck—but it must also be *expressive*, able to convey many distinct meanings. A perfectly learnable language (one word for everything) is useless; a perfectly expressive one (a unique, unanalysable signal




Figure 9.2. The iterated learning model. Each generation learns from a partial sample of the previous generation’s output—the transmission bottleneck. Because a partial sample can be reproduced in full only by generalising, structure and learnability increase over generations, even with no intention to design a better language.

for every meaning) cannot be learned from a partial sample. Compositionality is the resolution: by building meanings from reusable parts, a language becomes both learnable and expressive at once. Structure, on this account, is what a communication system settles into when it must be squeezed through the bottleneck of learning while still doing the work of communication [74].

9.4 Language Born from Broken Transmission

If transmission itself creates structure, then the most dramatic evidence should come from cases where transmission is broken and a new generation must build a language almost from nothing. Such cases exist, and they are extraordinary. When adults speaking different languages are thrown together and improvise a rudimentary *pidgin*—a makeshift jargon with little consistent grammar—the children who grow up hearing it do not merely reproduce it: they *creolize* it, spontaneously supplying the grammatical structure the pidgin lacked, and

Table 9.1. Modes of linguistic transmission, with their direction, characteristic fidelity, and typical outcome. Only the first is broadly “vertical”; all involve reconstruction rather than copying.

Mode	Direction	Fidelity	Outcome
L1 acquisition	community → child	high, reconstructive	native competence
L2 acquisition	horizontal	variable, partial	interlanguage, accent
Borrowing	horizontal	selective	loanwords, contact features
Iterated learning	generational chain	lossy (bottleneck)	emergent structure
Creolization	pidgin → child	child-enriched	new full grammar
Sign emergence	cohort → cohort	child-driven	new language

within a generation a full language with consistent syntax exists where none did before. Bickerton took this as evidence for an innate “language bioprogram” filling the gap [13].

Sign languages have furnished an even cleaner demonstration. When deaf children in Nicaragua were first brought together in schools around 1980, they had no shared language; from their improvised gestures a new sign language emerged—and, crucially, each successive cohort of young children made it more systematic than the cohort before, introducing hierarchical and grammatical structure that their older models had lacked [113]. The children were not copying a language; they were *creating* one, and the youngest learners drove the change. The same regularizing impulse appears when a single deaf child receives inconsistent input: the child imposes a consistency the input never had, surpassing her own models [115]. In every case the lesson is the one iterated learning predicts—the learner is not a passive vessel but an active regularizer, and transmission through the child is a creative act.

The creative impulse does not even require a community. Deaf children born to hearing parents who do not sign, and who are exposed to no sign language, nonetheless invent gestural systems of their own—*homesign*—that exhibit basic linguistic properties the surrounding gestures lack: stable word-like units, consistent ordering, and rudimentary morphology [60]. A single child, with no linguistic model at all, reconstructs from within herself the beginnings of grammar. Homesign is the purest demonstration of this chapter’s thesis: the structure of language comes not only from the input but from the learner, and

the drive to build a patterned system is so robust that it surfaces even in the total absence of a language to copy.

9.5 The Reconstruction Each Generation

We can now see precisely where the genetic analogy breaks, and why the break is fruitful. In biological heredity the code itself is copied: the genome is physically transmitted, and the offspring receives the instructions directly. In linguistic “heredity” nothing of the kind occurs. The child never receives her parents’ grammar; she receives only their *behaviour*—a stream of utterances—from which she must reconstruct a grammar of her own. Transmission is indirect, passing through the narrow channel of overt speech and back up into a freshly built competence, exactly the feedback from expression to code that the central dogma forbids in biology (Chapter 2). This is why linguistic transmission is at once so faithful and so creative: faithful, because the reconstruction is tightly constrained by the input and by the shared biases of human learners; creative, because reconstruction is not copying, and every act of it is an opportunity for regularization, generalization, and change. The meso-scale of transmission is therefore not a conveyor belt but a workshop. What passes down the generations is remade at every step, and it is in the remaking—the nightly reconstruction of Chapter 6 writ large across the lifespan—that the slow macro-scale changes of Part IV are ultimately born.

9.6 Summary

Linguistic heredity is transmission by reconstruction, not by copying. Each child rebuilds her language on a regular developmental timetable, within a maturational critical period, using powerful statistical learning to extract structure from the stream of speech. Because every learner sees only a partial sample, transmission passes through a bottleneck that forces generalization—and iterated learning shows that this bottleneck, repeated over generations, is itself a source of linguistic structure, as confirmed both in the laboratory and in the spontaneous birth of creoles and new sign languages, where children supply grammar the input never contained. The deep disanalogy with genetic heredity—that language is reconstructed from behaviour rather than copied as code—is exactly what makes linguistic transmission creative, and what carries the meso-scale upward into the historical changes of the macro-scale, to which Part IV now turns.

Coda: Linguistics and the Brahma Kumaris

The theme of this chapter—heredity as the carrying-forward of something across the generations—meets the Brahma Kumaris teaching at two points. The first returns us to Chapter 3. Alongside the cultural transmission of language, the tradition posits a second and more intimate channel of inheritance: the *sanskars*, the impressions borne by the soul itself from birth to birth [7, 23]. Where the linguist describes a grammar rebuilt from the community’s speech, the tradition describes a self that arrives already bearing its own accumulated tendencies—a soul-borne heredity running parallel to, and beneath, the behavioural one, of which the reported cases of xenoglossy would be the linguistic trace. The second point concerns fidelity. The iterated-learning model shows how transmission by imperfect reconstruction lets a language drift and change; the Brahma Kumaris *Drama*, by contrast, is the limiting case of *perfect* transmission—a world cycle that repeats identically, without loss, each *kalpa* [128]. Between these two poles—the lossy, creative reconstruction that makes human languages change, and the lossless repetition of the cosmic cycle—lies the whole tension of a temporal linguistics: how much of what returns is the same, and how much is made anew.

Part IV

The Macro-Cycles: History and Culture

Diachronic Genealogy Trees and Computational Phylogenetics

As species descend with modification from common ancestors, so do languages; and the tree that draws the one will, with care, draw the other.

Part IV enters the macro-scale—the slow cycles of centuries and millennia over which whole languages are born, diverge, and die. It is here that the genetic analogy of Chapter 2 is not merely suggestive but has already been *cashed out*, more thoroughly than in any other domain of linguistics. For more than a century, historical linguists have reconstructed the branching descent of languages; over the last two decades they have done so with the very mathematics that biologists use to reconstruct the tree of life. This chapter tells that story: how the family tree became a computational object, how a “clock” can be read in the vocabulary, how the tree meets the network where languages borrow, and how linguistic trees, ancient DNA, and archaeology are now being braided together to reconstruct the deep human past. It is the chapter in which the macro-cycle of language becomes quantitative.

10.1 Language Families as Lineages

The observation that languages fall into families descended from common ancestors is older than Darwin, and indeed helped to inspire him. When Sir William Jones observed in 1786 that Sanskrit, Greek, and Latin bore to one another a resemblance “stronger than could possibly have been produced by accident,” he concluded that they had “sprung from some common source, which, perhaps, no longer exists.” That common source—Proto-Indo-European—is the ancestor of a family that today includes most of the languages of Europe and much of South Asia, from English and Russian to Hindi and Persian. Within a few decades the philologist August Schleicher, explicitly borrowing Darwin’s imagery, was drawing languages as a branching *Stammbaum*, a family tree. The parallel that

Darwin himself remarked—that the formation of languages and of species is “curiously parallel” [40]—was thus a two-way exchange from the beginning, and the family tree has been the master image of historical linguistics ever since.

The unit of relatedness is the *cognate*: a word in two languages descended, by unbroken transmission, from a single word in their common ancestor—English *father*, German *Vater*, Latin *pater*, Sanskrit *pitr*. Cognates are the linguistic homologues of shared genes, and like homologous genes they are the evidence from which descent is inferred.

10.2 The Comparative Method

How does one tell a true cognate from a chance resemblance or a borrowing? The answer is the *comparative method*, the central achievement of nineteenth century linguistics and still the gold standard. Its key insight is that sound change, though it transforms words over time, is overwhelmingly *regular*: a given sound, in a given environment, tends to change in the same way throughout the vocabulary of a language. The most famous instance is *Grimm’s Law*, the systematic shift by which the Proto-Indo-European voiceless stops became fricatives in the Germanic branch—so that Latin *p* corresponds to English *f* (*pater* / *father*, *pes* / *foot*, *piscis* / *fish*) not sporadically but as a rule. Regular correspondences of this kind let linguists distinguish inheritance from accident, and even to *reconstruct* features of the unattested ancestor by working the changes backwards. The reconstructed proto-language is a hypothesis, marked with an asterisk, but a disciplined one: it is what the regular correspondences, run in reverse, require.

The comparative method thus does something with no exact parallel in biology: it reconstructs the actual ancestral “sequences,” the proto-forms, and not merely the shape of the tree. It is the linguistic time machine, and everything that follows depends on it.

10.3 The Lexical Clock and Glottochronology

If languages change at a roughly steady rate, then the *amount* of change between two related languages should measure the *time* since they parted—a linguistic version of the molecular clock. This was the hope of *glottochronology*, developed by Morris Swadesh [120]. Swadesh assembled a list of basic, culturally neutral vocabulary—body parts, low numbers, natural phenomena—that every language must have and that is resistant to borrowing, and proposed that such core vocabulary is replaced at an approximately constant rate, so that the proportion of cognates two languages still share decays exponentially with time (Figure 10.1). From an empirically estimated retention rate of roughly 86 per




Figure 10.1. The lexical clock. The proportion of cognates that two related languages still share decays with time depth. Core, high-frequency vocabulary (upper curve) is replaced more slowly than peripheral, low-frequency vocabulary (lower curve), so the “clock” ticks at a rate that depends on which words one counts—one reason the simple constant-rate assumption of classical glottochronology fails.

cent per millennium for the basic list, one could, in principle, read a date of divergence directly off the shared vocabulary.

Classical glottochronology was, in its strong form, a failure, and it is important to say so. The assumption of a single universal, constant rate does not hold: rates vary across languages, across meanings, and across time, and borrowing contaminates even the basic list. Critics were right to reject the promise of a simple linguistic date-stamp [26]. Yet the underlying intuition—that vocabulary carries a temporal signal—was not wrong, only too crudely formalised. What the field needed was not to abandon the clock but to replace the single fixed rate with a proper statistical model in which rates could vary and be estimated from the data. That is exactly what computational phylogenetics would provide.

10.4 Frequency and the Conservation of Words

Before turning to those methods, one result deserves its own section, for it ties the macro-scale directly back to the micro-scale mechanisms of Chapter 4. Pagel, Atkinson, and Pagel et al. [93] asked *which* words change fastest, and found a strikingly simple answer: the *frequency* with which a word is used predicts how slowly it is replaced. Across two hundred meanings and the whole recorded history of Indo-European, frequently used words—the commonest verbs, pronouns, and numerals—have been replaced far more slowly than rare ones, some persisting recognisably for ten thousand years or more. This is the upper curve of Figure 10.1, and it is the exact linguistic counterpart of a well-known biological fact: functionally important, heavily “used” genes are the most evolutionarily conserved. The connection to this book’s framework is direct. The entrenchment that Chapter 4 identified at the scale of the individual—high token frequency storing a form as a resistant whole—appears here, integrated over the whole community and ten millennia, as macro-scale evolutionary conservation. The micro-scale mechanism and the macro-scale pattern are one and the same, observed at opposite ends of the spectrum of linguistic time.

10.5 Computational Phylogenetics

The decisive step was to import, wholesale, the model-based methods that transformed evolutionary biology. Instead of counting shared cognates and converting to a date by a fixed rule, one specifies a probabilistic model of how cognates are gained and lost along the branches of a tree, and then uses *Bayesian inference* to find the trees—and the divergence dates—that best account for the observed vocabulary across many languages at once, complete with measures of uncertainty. Gray and Gray and Atkinson [62] first applied this apparatus to a matrix of 87 Indo-European languages scored for 2,449 cognate sets, and Bouckaert and Bouckaert et al. [20] extended it to an explicitly geographic model over 103 languages. The output is a dated phylogenetic tree of the kind shown in Figure 10.2: the branching descent of the family, with Anatolian and Tocharian splitting earliest, the great radiation of the core branches, and the modern languages at the tips.

The methods transfer because the problems are formally identical, as Table 10.1 sets out. A cognate set is a character, like a nucleotide site; cognancy is homology; borrowing and chance resemblance are the linguistic forms of homoplasy, the convergence that misleads; and the same likelihood and Bayesian machinery, the same clocks, the same outgroup rooting and fossil—or, here, ancient-text—calibration apply on both sides. The genetic analogy, which for most of this book has been a source of hypotheses, is at the macro-scale a working,

THE INDO-EUROPEAN FAMILY TREE




Figure 10.2. A schematic phylogeny of the Indo-European family, with time running from the Proto-Indo-European root (left) to the present (right). The earliest-splitting branches, Anatolian and Tocharian, are extinct (†); the core branches radiate and diversify into the modern languages. Trees of this form are now inferred computationally, with dates and uncertainties, from cognate data.

industrial-strength method.

10.6 Trees, Networks, and the Problem of Borrowing

The tree model has a built-in limit, and it is the disanalogy of Chapter 2 in a new guise. A tree represents pure vertical descent, branches diverging and never rejoining. But languages, unlike species, borrow freely from their neighbours, and borrowing creates resemblances that no tree can capture—horizontal links across the branches. When borrowing is heavy, a single bifurcating tree is not merely incomplete but misleading. Linguists have responded, again in step with biologists, by adopting *network* methods—NeighborNet and related techniques—

Table 10.1. The correspondence of methods between molecular and linguistic phylogenetics. The formal identity of the two problems is what allows the biologist’s tools to be applied, essentially unchanged, to language.

Concept	Molecular biology	Historical linguistics
Character	nucleotide site	cognate set
Homology	shared descent of a site	cognancy of words
Homoplasy	convergent mutation	borrowing, chance resemblance
Distance	genetic distance	% shared cognates
Inference	likelihood / Bayesian	likelihood / Bayesian
Clock	molecular clock	lexical clock
Rooting	outgroup	outgroup / attested ancestor
Calibration	dated fossils	dated ancient texts

that display the conflicting signals of descent and contact as a reticulated web rather than forcing them into a tree. The nineteenth-century debate between the *Stammbaum* (tree) model and the *Wellentheorie* (wave) model, on which changes spread across a dialect continuum irrespective of genealogy, is thus alive and quantified: real linguistic history is a tree overgrown with a network, and modern phylogenetics increasingly models both at once.

10.7 Archaeogenetics: Aligning Languages, Genes, and Pots

The richest payoff comes when linguistic trees are set beside two other records of the past: the archaeological record of material culture, and—newly and dramatically—the genetic record recovered from ancient skeletons. The test case is the century-old puzzle of the Indo-European homeland. Two hypotheses long competed. The *Anatolian* hypothesis places the origin in what is now Turkey some 8,000 to 9,500 years ago, spreading with farming; the *Steppe* (or *Kurgan*) hypothesis places it on the Pontic-Caspian grasslands some 6,000 years ago, spreading with horses and wheeled vehicles. Strikingly, the two leading kinds of evidence at first disagreed. The Bayesian phylogenetic analyses of linguistic vocabulary favoured the earlier, Anatolian date [20]; but the ancient-DNA revolution then delivered a spectacular result on the other side. Haak and colleagues, sequencing the genomes of dozens of prehistoric Europeans, found evidence of a massive migration from the steppe into Europe about 4,500 years ago, carrying an ancestry still present in Europeans today and fitting the Steppe

hypothesis [63, 99].

The apparent conflict has since begun to resolve into synthesis. A large 2023 analysis using “sampled ancestor” phylogenetic methods on 161 languages proposed a *hybrid* model: an ultimate origin south of the Caucasus around 8,000 years ago—reconciling the deeper linguistic date—followed by a secondary dispersal through the steppe, reconciling the genetic evidence [68]. Whatever its eventual verdict, the episode is a vivid demonstration of the framework’s central promise at the macro-scale: linguistic descent, genetic descent, and material culture are three coupled records of the same human history, and reading them together—each correcting and calibrating the others—is now among the most productive enterprises in the study of the deep past. The correlation of linguistic and genetic lineages, anticipated in Chapter 2, has become a science [92].

10.8 Applications and Visualisations

The maturing of linguistic phylogenetics has practical issue. Dated language trees can be draped over geography and time in GIS systems, animating the spread of families across continents and millennia. Comparative databases now function as something like a “linguistic genome browser,” letting a researcher inspect the phonological, lexical, and grammatical divergence of hundreds of languages at once. And the evolutionary framing has proved a powerful pedagogy, teaching the history of languages through the same tree-thinking that students meet in biology. Whatever one thinks of particular dates or homelands, the macro-scale of language is now, for the first time, a fully quantitative and visual science.

10.9 Summary

At the macro-scale the genetic analogy stops being a metaphor and becomes a method. Languages descend with modification from common ancestors; the comparative method reconstructs both the shape of the family tree and the proto-forms at its root; and the vocabulary carries a temporal signal—a lexical clock—that classical glottochronology formalised too crudely but that computational phylogenetics has rehabilitated with proper statistical models. Frequently used words are conserved just as heavily used genes are, tying the macro-scale pattern to the micro-scale mechanism of entrenchment. The tree must be supplemented by the network wherever languages borrow, and, most excitingly, linguistic trees are now being read together with ancient DNA and archaeology to reconstruct the deep human past, as in the still-unfolding story of the Indo-European homeland. The slow macro-cycle of linguistic descent has become

measurable; the next chapter turns from the tree of languages to the rhythms of discourse that play out upon it.

Coda: Linguistics and the Brahma Kumaris

The comparative method reconstructs a lost common source—a single ancestral tongue from which the branching family descends—and in doing so it recovers, by rigorous inference, an image the Brahma Kumaris teaching holds on other grounds: that of an original linguistic unity, the single connected tongue of the Golden Age from which the diversity of Kaliyug has since unfolded (Chapter 3). The tree of Figure 10.2, read to its root and beyond, points toward the same convergence that the *Kalpa Vriksha*, the world tree, depicts [23, 128]. Here, though, intellectual honesty requires marking a real divergence rather than a resonance. The timescales do not agree: computational phylogenetics dates Proto-Indo-European alone to some eight thousand years, and the deeper human linguistic past to tens of thousands more, whereas the Brahma Kumaris *kalpa* is a cycle of five thousand years. This is the sharpest point of tension in the book between its mainstream science and its cosmological frame, and we do not resolve it here. The reader who holds the longer chronology will read the tradition's cycle as symbol; the reader who holds the cycle will read the radiometric and phylogenetic dates as an artefact of assumptions about constant rates—the very assumptions that, as this chapter showed, classical glottochronology already had to abandon once. The framework's task is not to pretend the tension away but to state it plainly, and to let the two readings stand in honest dialogue, each aware of the other.

Chronotopic Pragmatics and Discourse

Every utterance carries a clock. It says not only what but when—and by saying when, it says where in time the speaker stands, and where the listener is asked to stand with him.

Chapter 10 traced the slow descent of languages down the macro-scale. This chapter stays at the macro-scale but shifts from the *system* to its *use*—to discourse, the actual weaving of language into texts and talk—and asks how time is built into that weaving. The claim is that temporality is not an occasional topic of discourse but a constant structuring dimension of it: every stretch of language is anchored in time, orders time, and situates its speaker and hearer within overlapping temporal frames that reach from the passing second to the historical epoch. The organising concept, borrowed from literary theory, is the *chronotope*; the mechanisms are the grammar of tense and deixis, the architecture of narrative, and the indexical linking of the small utterance to the large scale.

11.1 The Chronotope

The term *chronotope*—literally “time-space”—was coined by Mikhail Bakhtin [8] to name the way a genre of discourse fixes a characteristic configuration of time and space, an inseparable weld of the two that gives the genre its shape. The adventure romance, the biography, the folktale, the news report—each carries its own chronotope, its own implicit answer to *when and where does this kind of language happen*. Bakhtin’s insight, developed for the novel, generalises: all discourse is chronotopic, situated in a particular time-space and legible only against it. To adopt the LTC perspective is to take this seriously as a claim about temporal structure. A genre is, among other things, a temporal template—a habitual way of positioning events, speaker, and audience in time—and the recurrence of genres across a culture is one more of the macro-scale cycles this

book tracks. The remainder of the chapter unpacks the linguistic machinery through which discourse achieves its temporal situatedness.

11.2 Deixis and the Origo

The most basic temporal work language does is *deixis*: the anchoring of utterance to its context of production. Words like *now*, *then*, *today*, *yesterday*, *ago*, and *soon* have no fixed reference; they point, and what they point to depends entirely on *when* they are spoken. Their common anchor is the *origo*, the “here-now-I” of the speaking moment from which all deictic reckoning proceeds [54, 82]. Figure 11.1 lays out this deictic field: the *now* of speech at the centre, the past to one side and the future to the other, with the temporal deictics arrayed around the origo they depend on.

Tense is deixis grammaticalised. To mark a verb as past, present, or future is to locate its event relative to the speaking now, and the traditional analysis due to Reichenbach [100] shows how finely this is done: tenses coordinate not two times but three—the speech time, the event time, and a *reference* time from which the event is viewed—so that the past perfect (“had left”) places the event before a reference point that is itself before now. *Aspect*, by contrast, concerns not the location of an event in time but its internal temporal shape—whether it is viewed as complete or ongoing, punctual or extended [32]. Between them, tense and aspect give every finite clause a precise temporal signature, computed automatically from the origo. Language cannot help telling time.

11.3 The Grammar of Time Across Languages

If every language tells time, they do not all tell it the same way, and the cross-linguistic diversity is itself instructive. Some languages, such as Mandarin Chinese or Yucatec Maya, are effectively *tenseless*: they do not obligatorily mark past, present, and future on the verb at all, relying instead on aspect, temporal adverbs, and context to locate events [32, 39]. This does not mean their speakers are indifferent to time—temporal reference is recovered as readily as in tensed languages—but it shows that grammaticalised tense is one option among several for encoding a universal concern. Other languages grammaticalise temporal-epistemic distinctions that European languages leave to context: systems of *evidentiality* obligatorily mark whether the speaker witnessed an event, inferred it, or heard of it, folding the *time and manner of the speaker’s own knowing* into the grammar of the reported event [2]. Aspect, meanwhile, varies richly—the perfective and imperfective, the perfect, the habitual, the progressive—each carving the internal time of an event differently.




Figure 11.1. The field of temporal deixis. All temporal deictics—*now*, *yesterday*, *tomorrow*—and the tense system reckon from the *origo*, the “now” of the speaking moment (top). Discourse can relocate this deictic centre: a narrator may shift the “now” into the past and narrate from there, as in the historical present (bottom).

The lesson for the LTC framework is that temporal encoding is a universal function realised by language-particular means. Every language must let its speakers locate, shape, and source events in time; *how* it does so is one of the parameters that Chapter 3 saw set differently in each branch of the family tree. Table 11.1 summarises the principal devices, cutting across these typological differences, by which discourse encodes time.

11.4 Narrative and the Ordering of Time

Discourse does more than anchor events to the now; it *orders* them, and narrative is the art of that ordering. Here a fundamental duality appears, long recognised in narratology: the distinction between the time of the events told (the *story time*, or *fabula*) and the time of the telling (the *discourse time*, or *sjuzhet*). The two

Table 11.1. The principal linguistic devices by which discourse encodes and manipulates time, with their form and temporal function.

Device	Form	Temporal function
Tense	past / present / future marking	locate event relative to the now
Aspect	perfective / imperfective	internal temporal shape of event
Evidentiality	witnessed / inferred / reported	source and time of knowing
Temporal deixis	<i>now, then, yesterday</i>	anchor to the origo
Discourse markers	<i>meanwhile, since, already</i>	relate events in time
Historical present	present tense for past events	vividness, immediacy
Framing formulae	"once upon a time," "breaking now"	set genre and stance
Proximization	"imminent," "the time is now"	collapse distance, mobilise

need not coincide, and the systematic study of their divergence is due above all to Gérard Genette [55], who distinguished three dimensions along which they can pull apart (Figure 11.2). *Order* concerns sequence: a narrative may depart from chronology through *anachronies*—the flashback that tells a later event first, the flash-forward that anticipates. *Duration* concerns pace: a narrative may compress years into a sentence (summary) or dilate a moment across pages (scene). *Frequency* concerns repetition: an event that happened once may be told many times, or a habitual event told once ("every morning she would...").

Everyday narrative exploits the same resources. Labov and Labov and Waletzky [78], studying oral stories of personal experience, showed that even unrehearsed narratives have a regular temporal architecture—orientation, complicating action, evaluation, resolution—and that their backbone is a sequence of clauses whose order mirrors the order of events, a *temporal juncture* that listeners reconstruct automatically. Narrative is thus a machine for manipulating the relation between two timelines, and much of its power—suspense, revelation, irony—lives in the controlled mismatch between them.

11.5 Discourse Across Sociopolitical Cycles

Zoom out from the single narrative to the discourse of whole societies, and further temporal cycles appear. Genres of public discourse are keyed to the recurring phases of political and economic life. Wartime rhetoric recurs with

NARRATIVE TIME




Figure 11.2. The two clocks of narrative. *Story time* (bottom) is the chronological order of events; *discourse time* (top) is the order in which they are told. Their divergence—here a flashback that narrates event E3 first—is the resource narrative exploits, along dimensions of order, duration, and frequency [55].

its appeals to unity, sacrifice, and urgency; crisis communication recurs with its repetition, simplification, and deictic pressure of the immediate; post-crisis discourse recurs with its narratives of reconstruction, blame, and healing. These are chronotopes at the scale of history: each returns when its phase of the cycle returns, bringing its characteristic vocabulary and stance.

One temporal device deserves particular note, because it has been studied closely in political discourse. *Proximization*, analysed by Cap [27], is the rhetorical strategy of construing distant threats—temporally or spatially remote—as approaching the audience’s here-and-now, collapsing the distance in order to mobilise. To say that a danger is *imminent*, that the time to act is *now*, that history is *at a turning point*, is to manipulate the temporal deixis of Section 11.2 for suasive ends. Temporal framing, in such discourse, is not decoration but the very mechanism of persuasion, and its recurrence across the political cycle is one of the more consequential of the macro-scale’s rhythms.

11.6 Chronotopes of Digital Media

Digital communication has generated new chronotopes at a startling rate, and they are unlike the old ones in their temporal structure. There is a *viral time* in which a meme or a hashtag rises and decays within hours, a compressed and frantic tempo with no pre-digital equivalent. There is an *archival time* in which the same utterance, once posted, persists indefinitely and searchably, so that words spoken in one moment can be retrieved and recontextualised years later. And these coexist: a single platform layers the ephemeral and the permanent, the live feed and the searchable archive, so that discourse participants inhabit several timescales at once. The meme itself is a chronotopic unit—a compressed, shareable packet of time-space that carries its moment with it. Temporal anchoring, in this environment, becomes fragmented and multiple, several clocks running in the same stream of talk.

This layering is a new intensity of an old condition. Bakhtin held that discourse is inherently *heteroglot*, carrying many voices—and with them many times—within a single text: the reported speech of another moment, the cliché that arrives with its whole history attached, the quotation that reaches back centuries. Digital media merely make this polytemporality vivid and continuous, so that a single feed splices the live, the recent, the archived, and the ancient into one scrolling present. To read such discourse is to move constantly among timescales, and the competent reader does so without effort—further evidence that holding many times at once is not an exotic feat but the ordinary condition of language in use.

11.7 Orders of Indexicality: The Small Utterance and the Large Scale

We can now draw the chapter's threads together with a concept that speaks directly to the architecture of this book. When a speaker uses a form—a word, an accent, a temporal framing—that form does more than convey its literal content; it *indexes*, points to, the larger contexts in which it has value. Silverstein's notion of *indexical order* and Blommaert's of *orders of indexicality* describe how a single feature of talk is simultaneously legible at several levels: as a move in this conversation, as a marker of a social identity, as a token of a genre, as an emblem of a historical moment [15, 114].

Blommaert's development of this idea is especially apt here, for he frames it explicitly in terms of *scale*. Sociolinguistic phenomena, he argues, are not merely spread horizontally across space but layered vertically across scales, from the local and momentary to the translocal and enduring, and a single utterance operates on several scales at once, indexing the here-and-now while also invoking the timeless and the historical [16]. This is, in the vocabulary of

this book, precisely the coupling of the micro to the macro: a small discourse event, produced in a passing second, reaches up through nested scales to index the epoch. The chronotope, the deictic origo, the two clocks of narrative, and the orders of indexicality are all, in the end, ways that a single utterance holds many timescales together—the book’s central image, realised in the fabric of discourse itself.

11.8 Summary

Discourse is temporal through and through. Genres are chronotopes, weld configurations of time-space that recur across a culture; deixis and tense anchor every utterance to the origo of the speaking now, and can relocate that origo for narrative effect; narrative runs two clocks at once, story time and discourse time, and lives in the controlled divergence between them; public discourse cycles through recurring genres keyed to the phases of political and economic life, with temporal proximization as a potent suasive device; and digital media layer viral, archival, and historical times in a single stream. Beneath all of these, the orders of indexicality show how one small utterance indexes many scales at once, coupling the passing second to the historical epoch. The macro-scale of discourse, like the macro-scale of descent in Chapter 10, turns out to be built from the same nested, coupled cycles that organise the whole of language in time.

Coda: Linguistics and the Brahma Kumaris

If any discourse is chronotopic to its core, it is the Brahma Kumaris *Murli*. Spoken daily, and ideally received at the pre-dawn *Amrit Vela*, the *Murli* continually situates its listener at a precise cosmic time-space: not the ordinary now, but the *Sangam Yuga*, the Confluence Age, the hinge between the old cycle and the new [23, 128]. Its characteristic discourse performs exactly the deictic relocation of Figure 11.1, shifting the deictic centre from the mundane present to “this most auspicious confluence,” and it performs, too, the proximization of Section 11.5, construing the turning of the cycle as imminent and the moment for transformation as *now*. The ideal of the *trikaldarshi*, the seer of the three times, is the aspiration to hold past, present, and future in one view—the origo widened until it encompasses the whole Drama. In the terms of this chapter, the tradition’s discourse is a sustained exercise in orders of indexicality: every ordinary word of the *Murli* is made to index the largest scale of all, the cycle of time itself, so that the smallest utterance and the whole *kalpa* are held together in a single breath. That is the chronotopic ambition of the framework, stated in its own devotional key.

Language and Cultural Evolution

Language is the fastest-evolving thing our species carries, and the one that most shapes us in return.

This chapter completes Part IV and, with it, the descent through the three scales. Chapter 10 gave the macro-scale its tree; Chapter 11 gave it its discourse; this chapter gives it its deepest frame. For language is not only a biological endowment and a cognitive achievement—it is a *cultural* system, subject to the same Darwinian logic of variation, selection, and inheritance that governs all of culture, yet running faster and cutting deeper than any other cultural form. To see language as a cultural evolutionary system is to see the whole book’s temporal architecture converge: the micro-scale biases of Part II, the meso-scale transmission of Part III, and the macro-scale descent of Chapter 10 are, in the end, the machinery of a single evolutionary process operating on a cultural replicator. This chapter draws that process out, follows it to the co-evolution of language and the human brain, and ends by gathering the scales into one picture.

12.1 Language as a Cultural System

Culture, like biology, is a system of inheritance. It has *variation* (individuals differ in what they know and do), *selection* (some variants spread while others vanish), and *transmission* (variants pass from person to person, and generation to generation). Wherever these three are present, Darwinian dynamics follow, and a mature science of *cultural evolution* has grown up to study them, with the quantitative models of Cavalli-Sforza and Cavalli-Sforza and Feldman [28], the dual-inheritance theory of Boyd and Boyd and Richerson [22], Richerson and Boyd [101], and the synthesis of Mesoudi [86]. Human beings, on this view, are the products of two coupled inheritance streams—genes and culture—that have shaped each other over evolutionary time [123].

Language is the paradigm case of the cultural stream, and in some ways its purest. It is discretely structured, richly variable, densely transmitted, and—crucially for this book—traceable through time as no other cultural system is, as

Chapter 10's phylogenies showed. If culture is a Darwinian system, language is the place to watch it evolve.

12.2 The Meme and the Lingueme

The idea that culture contains its own replicators is due to Dawkins [41], who coined the *meme* as the cultural analogue of the gene—a unit of imitation, copied from mind to mind, competing for the limited resources of attention and memory. Blackmore pushed the idea furthest, arguing that humans are essentially “meme machines” and that language itself arose as a system for spreading memes with high fidelity [14]. We met the difficulties of this proposal in Chapter 2: the meme, as a general unit, is hard to delimit and count. But we also met its linguistic rescue—Croft's *lingueme*, the concrete token of linguistic structure in an utterance, which supplies the countable replicator that memetics in general lacks [34]. Language, precisely because its replicators are relatively discrete and recorded, is where the memetic programme comes closest to being a measurable science. The words, constructions, and sounds of a language are cultural replicators whose frequencies rise and fall over time under selection—and, as Chapter 10 showed, that rise and fall can now be tracked with the same phylogenetic tools used for genes.

12.3 Linguistic Ecology and Transmission Biases

If linguemes are replicators competing in an ecology of minds, what determines which of them spread? Cultural evolution has identified a small set of *transmission biases*—systematic tendencies in what people copy—and each has a direct linguistic manifestation, drawing together threads from across this book (Table 12.1). A *content bias* favours variants that are intrinsically easier to learn, remember, or use—the pressure toward learnable, compositional structure that iterated learning revealed in Chapter 9. A *prestige* or model-based bias favours variants used by high-status individuals—the engine of borrowing from prestige varieties and of the linguistic marketplace of Chapter 8. A *conformity bias* favours whatever the majority does—the force behind regularization, koineization, and standardization. A *frequency-dependent bias* means that how often a form is encountered affects whether it is copied—the entrenchment that Chapter 4 found in the individual and Chapter 10 found conserved across ten millennia. And beneath all of these runs *drift*, the random sampling that changes frequencies by chance, especially in small populations [9].

The point of the table is not merely to classify but to unify. The macro-scale process of cultural evolution is not a new mechanism bolted on above the others; it *is* the others, integrated over populations and centuries. The same

Table 12.1. The principal transmission biases of cultural evolution and their manifestations in language, with the chapters where each appears. The macro-scale process of cultural evolution is built from mechanisms already met at the smaller scales.

Bias	What it is	Linguistic effect
Content bias	easier-to-learn variants favoured	compositional, learnable forms (Ch. 9)
Prestige bias	copy high-status models	borrowing, prestige spread (Ch. 8)
Conformity bias	copy the majority	regularization, standardization
Frequency bias	frequency affects copying	entrenchment, conservation (Ch. 4, 10)
Drift	random sampling in copying	chance change, small populations (Ch. 2)

entrenchment that resists a slip of the tongue at breakfast conserves a word for ten thousand years; the same prestige that shifts a register in a job interview drives a sound change across a nation. Cultural evolution is what the micro- and meso-scale biases of Parts II and III look like from the top of the spectrum of linguistic time.

12.4 Chronocultural Co-Evolution: Technology and Language

Language does not evolve in a vacuum; it co-evolves with the material culture of communication, and each transformation of communicative technology has reshaped linguistic form (Figure 12.1). The invention of *writing* some five thousand years ago gave language permanence and made possible law, scripture, and history—and a new, more deliberate register freed from the limits of memory. The *printing press* drove standardization and the rise of national vernaculars. The *telegraph* rewarded compression, breeding a terse “telegraphese.” Broadcast media propagated spoken standards; the *internet* produced “written speech” and a wave of informalization; social media compressed the tempo into viral time and the meme. Each epoch of communicative technology has its own chronotope and its own selective pressures on linguistic form—a *chronocultural* co-evolution in which the tools of communication and the language they carry remake each other, epoch by epoch.

CHRONOCULTURAL CO-EVOLUTION: MEDIA EPOCHS & LANGUAGE




Figure 12.1. The co-evolution of communicative technology and linguistic form. Each epoch of media—from writing to the printing press to the internet and beyond—brings its own selective pressures and its own consequences for language, from standardization to informalization to viral compression. (Time is not drawn to scale.)

12.5 Brain–Language Co-Evolution

The deepest co-evolution is with the human brain itself, and here cultural evolution returns us, transformed, to the innateness question of Chapter 3. The traditional view held that language is the way it is because the brain contains an innate Universal Grammar that language must obey. Christiansen and Christiansen and Chater [31] argued for something close to the reverse. Because language change is far faster than genetic change—the point dramatised in Figure 12.2—language is a “moving target” that the genome could never have tracked; a genetically fixed Universal Grammar could not have co-evolved with so fast-moving a partner. The better explanation, they propose, is that *language has adapted to the brain*, not the brain to language: over cultural evolutionary time, languages have been shaped, by the transmission biases of the previous section, into forms that human brains find easy to learn and process. Grammar is thus, in large part, the accumulated residue of what generations of learners could most easily acquire—the iterated-learning result of Chapter 9 raised to the scale of the species.

This does not abolish biology’s role but redistributes it. The brain still sup-




Figure 12.2. The asymmetry of brain–language co-evolution. Because language evolves orders of magnitude faster than genes, the dominant influence runs downward: the brain’s learning biases shape language (strong arrow), while language shapes the genome only weakly and slowly, through the *Baldwin effect* and *niche construction* (dashed arrow). Language is the moving target.

plies the biases, and language still exerts a slow return pressure on the genome—through the *Baldwin effect*, by which learned abilities can create selection for their own easier acquisition, and through *niche construction*, by which a language-using species remakes the very environment its descendants must adapt to [44]. Deacon’s “symbolic species” is the name for the result: a creature and its language that have made each other, each the other’s cause across different scales of time. The nature-versus-nurture debate over language dissolves, in this frame, into a question of *coupled evolutionary tempos*—exactly the kind of question the LTC framework is built to pose.

12.6 Temporal Bottlenecks: Death, Creolization, and Revival

Cultural evolution, like biological evolution, has its catastrophes and its interventions. Languages can die: when transmission to children ceases—through conquest, assimilation, or the crushing prestige of a dominant tongue—a lan-

guage passes out of the world, and on current estimates a large fraction of the world's several thousand languages are so threatened [35]. Such *temporal bottlenecks*, the linguistic counterpart of the population bottlenecks of genetics, are moments when the transmission chain of Chapter 9 is severed. Sometimes life bursts through the break, as in the creolization and new sign languages already met. And sometimes humans intervene deliberately in linguistic time—the one thing no gene can do. The revival of Hebrew from a liturgical language into a living vernacular, and the sustained efforts to reverse the decline of Welsh, Māori, and other languages, are acts of conscious, *directed* cultural evolution, the Lamarckian and intentional character of linguistic change (Chapter 2) exercised on the largest scale. Language is the one inheritance system whose bearers can choose to save it.

12.7 The Convergence of the Scales

We can now close the arch of the book's argument. Language is at once a biological phenomenon, gated by the circadian and developmental clocks of Parts I and II; a cognitive one, structured by the fast processes of memory and rhythm; a social one, tuned by the register cycles of Part III; and a cultural one, evolving by selection and drift across the historical time of Part IV. These are not four different objects but one object seen at four bands of a single spectrum of time. The entrenchment of a frequent word is a micro-scale fact, a meso-scale transmission bias, and a macro-scale evolutionary conservation—the same phenomenon, thrice named. Cultural evolution is the level at which this unity becomes explicit, because it is the level that integrates all the others: it takes the biases generated in the individual mind, transmits them through the community, and accumulates them into the historical descent of languages. Language, in the fullest sense, is a multi-scale temporal artifact, and the cultural-evolutionary frame is where its scales are finally seen to be one.

12.8 Summary

Language is a cultural evolutionary system—perhaps the clearest one we have. Its replicators, Croft's *linguemes*, spread and decline under the transmission biases of cultural evolution: content, prestige, conformity, frequency, and drift, each of which is a mechanism already met at the smaller scales, now integrated over populations and centuries. Language co-evolves with the material culture of communication, epoch by epoch, and—most profoundly—with the human brain itself, which it has shaped and been shaped by across vastly different tempos, so that the innateness question becomes a question of coupled evolutionary clocks. Its transmission can be broken in temporal bottlenecks, and, uniquely,

it can be deliberately revived. In cultural evolution the three scales of the book converge: language is one temporal artifact, evolving at every scale at once. With the descent of language thus completed, Part V turns from what the framework claims to how it might be tested and built.

Coda: Linguistics and the Brahma Kumaris

Cultural evolution describes how a tradition reproduces itself in time, and the Brahma Kumaris are, among other things, a case study in exactly that: a community that sustains a distinctive linguistic culture—the Murli, the specialised vocabulary of Appendix E, the practices of remembrance—through deliberate daily transmission [126, 128]. Two deeper resonances are worth marking. First, the direction of the story. Mainstream cultural evolution generally narrates language as ascending from simple to complex, structure accumulating over time; the Brahma Kumaris narrate the reverse—an original perfection and unity in the Golden Age, *declining* through the cycle into the fragmentation of the present [23]. The two need not be reconciled to be set instructively side by side: one is a story of cultural construction, the other of cultural descent, and both are narratives of language moving through time. Second, the theme of deliberate revival with which this chapter ended finds its largest possible expression in the tradition's own project: the conscious effort, at the Confluence Age, to restore the lost purity of thought, word, and deed—a directed transformation of culture and of language undertaken in the belief that the cycle is about to turn. Where the linguist sees communities choosing to save a dying tongue, the tradition sees souls choosing to remake themselves before the wheel of time completes its round.

Part V

Methods, Models, and Futures

Experimental LTC Studies

A framework earns its keep not by the pictures it paints but by the experiments it makes it natural to run.

Parts I through IV built the LTC framework and traced it down the three scales. Part V asks the harder question: is it *true*? A temporal linguistics that generated no testable predictions would be, as Chapter 1 warned, literature rather than science. This chapter therefore turns from theory to method, and surveys how the framework's claims can be put to empirical test. The task is unusual, because LTC predictions span timescales from the millisecond to the millennium, and no single method reaches across them. What is required is a coordinated *programme*, borrowing the experimental apparatus of chronobiology, psycholinguistics, neuroscience, and variationist sociolinguistics, and pointing it at the specifically temporal hypotheses the framework generates. This chapter sets out that programme, scale by scale, and ends with a table that ties each of the book's central predictions to a design that could confirm or refute it.

13.1 Isolating the Linguistic Clock

The micro-scale prediction of Part II—that language performance varies with circadian phase—faces an immediate methodological obstacle. Performance at a given clock hour is confounded: it reflects not only the endogenous circadian rhythm (Process C) but also accumulated sleep pressure (Process S), the effects of meals and activity, and the demands of whatever task fills that hour. To isolate the clock, one must strip these confounds away, and chronobiology has developed exactly the protocols for doing so (Figure 13.1). In the *constant routine*, participants remain awake in constant dim light, temperature, and posture, taking identical small meals at regular intervals, so that any remaining rhythm in their performance must be endogenous [46]. In *forced desynchrony*, participants live on an artificial day length that the circadian clock cannot follow, so that circadian phase and time-awake gradually separate and their effects can be statistically disentangled.




Figure 13.1. Isolating the endogenous circadian component of language performance. Raw performance at a given clock hour is confounded by sleep pressure, chronotype, meals, and activity. Chronobiological protocols—the constant routine and forced desynchrony—hold these masking factors constant, revealing the underlying circadian rhythm of the linguistic measure [46].

Applied to language, these designs turn the informal observations of Chapter 5 into rigorous tests. One would sample spontaneous or elicited speech at many points across such a protocol, measuring syntactic complexity, lexical diversity, disfluency, and figurative-language use, and ask whether they trace an endogenous circadian curve. *Chronotype* provides an additional lever, exactly as in the synchrony-effect studies of Chapter 5: because larks and owls differ in circadian phase, comparing them at fixed clock hours dissociates clock time from internal time [83, 103, 111]. Circadian biomarkers—melatonin onset, the cortisol rhythm, core body temperature—provide the independent clock against which linguistic performance can be phase-aligned.

13.2 Measuring Language Over Time

Testing temporal predictions requires temporal data, and a quiet revolution in measurement has made such data abundant. Three approaches complement one another. *Naturalistic recording* captures spontaneous language in the wild: the Electronically Activated Recorder samples ambient speech at intervals across the day [85], and dense home recording can track language development continuously. The most remarkable such corpus, the Human Speechome Project, recorded almost the entire first three years of one child's life, allowing the very moment a word entered the child's vocabulary to be located and its triggers analysed [105]—a meso-scale acquisition study of a fidelity previously unimaginable. *Digital corpora* supply the same density at population scale: the time-stamped output of social media constitutes a natural experiment in diurnal and seasonal language, as the mood study of Chapter 5 demonstrated [59]. And *diachronic corpora*—collections of dated texts spanning centuries—are the macro-scale's instrument, letting the lexical clock and the frequency–conservation law of Chapter 10 be measured directly. Across all three, the quantitative measures are shared: mean length of utterance, type–token ratio and other indices of lexical diversity, measures of syntactic depth, and rates of disfluency and error.

13.3 Neural Oscillations and the Linguistic Hierarchy

The framework's boldest micro-scale claim—that language is organised by nested oscillations (Chapter 4, Appendix C)—can be tested directly in the brain, and the evidence is already striking. Using magnetoencephalography, Ding, Melloni, and Ding et al. [45] showed that cortical activity *tracks the linguistic hierarchy*: when listeners hear connected speech, distinct cortical rhythms lock onto distinct levels of structure at once—a fast rhythm to the syllables, a slower one to words, and a slower one still to phrases—even when those levels have no acoustic cue and must be built internally by the grammar (Figure 13.2). This is the nesting of Appendix C made visible in neural tissue: the brain represents the linguistic hierarchy as a hierarchy of coupled oscillations.

The broader research programme, developed by Giraud and Giraud and Poeppel [58], holds that delta, theta, and gamma oscillations “package” the incoming speech stream into units of the appropriate temporal grain, and that their cross-frequency coupling is foundational to comprehension. For the LTC framework this is a direct vindication at the fastest scale: the entrainable oscillators posited on linguistic grounds in Chapter 4, and formalised in Appendix C, correspond to measurable brain rhythms, and their nesting—gamma within theta within delta—is the neural signature of the nested time cycles the whole book describes. Time-resolved fMRI and the study of oscillatory dynamics dur-

NEURAL OSCILLATIONS TRACK NESTED LINGUISTIC UNITS




Figure 13.2. Cortical oscillations track the nested units of linguistic structure. Slow delta rhythms lock onto phrases, theta rhythms onto syllables, and fast gamma rhythms—nested within the theta cycles—onto phonemic detail. The brain’s cross-frequency nesting mirrors the nested time cycles of the LTC framework [45, 58].

ing spontaneous versus rehearsed speech extend the same programme to the coordination of the temporal language network as a whole.

13.4 Sleep and Second-Language Learning

The consolidation claims of Chapter 6 are unusually testable, because sleep can be manipulated. The canonical design contrasts learning followed by sleep with learning followed by an equal interval of waking, holding total elapsed time constant, so that any advantage must be attributed to sleep rather than to the mere passage of time—the logic that established overnight lexical integration [47] and sleep-dependent grammatical abstraction. *Nap studies* compress the design into a single day; *targeted memory reactivation*, in which cues associated with newly learned words are replayed during sleep, moves from observation to intervention, showing that consolidation can be actively steered [19, 88, 112].

Such designs have obvious application to language pedagogy: they specify *when*, relative to sleep, study is most effective, and how the offline hours might be recruited to accelerate second-language learning.

13.5 Longitudinal and Panel Designs

The meso- and macro-scales demand designs that themselves extend over time. The *apparent-time* method reads change from a single synchronic sample by comparing age groups, but it rests on an assumption—that speakers are stable across their lives—that only *real-time* data can test. Two designs do so. A *trend study* re-samples a community at intervals with fresh participants; a *panel study* re-records the very same individuals as they age. Sankoff and Sankoff and Blondeau [110] combined both to validate the apparent-time construct for Montreal French, and such longitudinal work remains the gold standard for the meso-scale [77]. At the macro-scale, the diachronic corpus is the natural experiment that no researcher could stage, preserving centuries of change for retrospective measurement. Together these designs let the lifespan curve of Chapter 8 and the descent trees of Chapter 10 be tested against the actual flow of time rather than inferred from a single slice of it.

13.6 A Programme of LTC Experiments

The preceding sections can be gathered into a single research programme, set out in Table 13.1. Its virtue is coverage: it attaches to each of the book's central predictions, at each scale, a concrete design and a measurable outcome, so that the framework as a whole becomes falsifiable rather than merely suggestive. No one laboratory could execute all of it—the timescales are too disparate—but that is precisely the point. A temporal linguistics is inherently a collaborative enterprise, pooling chronobiology, neuroscience, psycholinguistics, and sociolinguistics, because language lives on all their timescales at once. The unifying contribution of LTC is not a single experiment but the recognition that these scattered literatures are studying one phenomenon at different magnifications.

13.7 Summary

The LTC framework is testable, and the instruments already exist. At the micro-scale, chronobiological protocols can isolate the endogenous circadian rhythm of language performance, and neuroimaging has already shown the brain tracking the linguistic hierarchy as a hierarchy of nested oscillations—a direct confirmation of the framework's central image. Naturalistic, digital, and diachronic corpora supply temporal data at every scale, from the moment a child learns

Table 13.1. A programme of experiments testing LTC predictions across the scales. Each row pairs a central prediction of the book with a design that could confirm or refute it and the measure it would use.

Scale	Prediction	Design	Measure
Micro	syntax simplifies toward evening	constant routine, time-of-day sampling	MLU, clause depth
Micro	oscillations track linguistic units	MEG/EEG cortical tracking	delta/theta/gamma coherence
Micro	sleep integrates new words	nap / pre-post, TMR	lexical competition, recall
Meso	adolescent peak, age-grading	apparent + real-time panel	vernacular variant rate
Macro	frequency predicts conservation	diachronic corpus phylogenetics	replacement rate vs frequency

a word to the millennial descent of a family. Sleep can be manipulated to test consolidation, and longitudinal panel and trend studies can test the lifespan and historical predictions against real time. Assembled into a single programme, these designs make the framework falsifiable across the whole spectrum of linguistic time. What remains is to *model* the phenomena thus measured—to build systems that simulate and reproduce the temporal dynamics of language—which is the task of the next chapter.

Coda: Linguistics and the Brahma Kumaris

The experimental logic of this chapter has an unexpected counterpart in Brahma Kumaris practice. The constant routine of Figure 13.1 works by stripping away masking influences so that an endogenous rhythm can be observed in isolation—and this is, in a different register, precisely the aim of the pre-dawn meditation of *Amrit Vela*: to still the outer noise of activity and sensation so that the soul's own state can be observed directly [23, 128]. The tradition proposes a disciplined, first-person method of inner observation, and it treats the transformation of *sanskars* over months and years as something to be tracked and measured by sustained self-study—a longitudinal design turned inward. We do not suggest that contemplative introspection and controlled experiment are interchangeable; they answer to different standards of evidence. But they share a common intuition, which is also the methodological heart of this chapter: that to find a rhythm one must first quiet what obscures it. Third-person science isolates the

clock by controlling the environment; the contemplative isolates it by quieting the mind. Both are, in their way, experiments in time.

Computational Modeling and AI Systems

To model a thing in time is to ask whether the mechanisms we imagine would really produce the history we observe.

Theory proposes and experiment tests; between them stands a third mode of inquiry, *computational modeling*, which asks whether a proposed mechanism, set running in a simulation, actually generates the phenomenon it is meant to explain. For a framework organised around time, this mode is indispensable, because temporal processes are exactly the ones whose consequences are hardest to intuit: small biases compounded over millions of utterances, structure emerging from transmission, meaning drifting across centuries. This chapter surveys how the LTC phenomena can be modelled and simulated—from agent-based models of language change to the neural sequence models that now dominate artificial intelligence—and it closes by turning the lens on those very models, which are themselves, this chapter argues, artifacts frozen in linguistic time.

14.1 Agent-Based Models of Language Change

The macro-scale processes of Part IV are, at bottom, the aggregate of countless local interactions between speakers, and *agent-based models* make that aggregation explicit. One populates a simulation with agents, each carrying some linguistic state, lets them interact by simple rules, and watches what emerges at the population level over simulated generations. Luc Steels [116] pioneered the approach with “language games”: in the *naming game*, agents that begin with no shared vocabulary converge, through repeated pairwise interactions and feedback, on a common lexicon, with no central coordination—a vivid demonstration that shared convention can self-organise from local negotiation. The utterance selection model of Baxter and Baxter et al. [9], met in Chapter 2, gives the same idea a population-genetic form, deriving the S-curves and actuation dynamics of Chapter 4 from the differential replication of linguistic variants.

Agent-based models are the natural computational home of the macro-cycle: they let one ask whether a proposed micro-scale bias—a prestige preference, a learnability pressure—really would, compounded over a community and a history, produce the macro-scale change observed.


14.2 Neural Sequence Models and the Representation of Time

If agent-based models simulate the community, neural network models simulate the individual processor of sequential language—and their history is, revealingly, a history of changing ways to represent *time* (Figure 14.1). The first breakthrough was the *recurrent neural network*, introduced to let a network “find structure in time” by feeding its own hidden state back as input at each step, so that a trace of the past is carried forward through the sequence [51]. Its limitation—that information decays across long gaps—was addressed by the *long short-term memory* network, whose gated memory can hold information across many steps [69], an engineered answer to exactly the storage problem that center-embedding poses to human speakers (Chapter 7).

Then came a conceptual reversal. The *Transformer* architecture that underlies today’s large language models *abolished recurrence altogether* [125]. Instead of processing a sentence word by word in time, it processes all words in parallel, using an *attention* mechanism in which every word can directly consult every other. But a model with no sequential processing has no intrinsic sense of order, so order must be supplied artificially, as an explicit *positional encoding* added to each word. This is a striking move from the standpoint of this book: where the RNN treated time as *process*—the sequential unfolding of computation—the Transformer treats time as *position*, a coordinate stamped onto each token and handled like any other feature. The extraordinary success of the Transformer shows that much of what language requires can be captured by position without process; whether the loss of genuinely sequential dynamics matters for the deepest temporal phenomena of language remains an open and intriguing question.

14.3 Iterated Learning and the Evolution of Structure in silico

The transmission processes of Chapter 9 were, in fact, first established through simulation. Kirby’s *iterated learning models* implement a chain of artificial learners, each acquiring a language from the output of the previous one, and show that compositional structure emerges over the generations purely from the pressure to pass a language through the bottleneck of learning [73, 74]. Modern versions replace the hand-built learners with neural networks, letting deep-learning



DIACHRONIC SEMANTIC DRIFT (word-embedding space)




Figure 14.2. Semantic change made measurable. In a word-embedding space trained on successive periods, a word’s vector migrates as its meaning shifts: “gay” moves from the region of *merry, cheerful* toward that of *homosexual*, and “broadcast” from *sow, scatter, seed* toward *transmit, media*. Diachronic embeddings turn semantic drift into a trajectory that can be measured and its laws discovered [65].

scale. Hamilton, Leskovec, and Hamilton et al. [65] used this method to uncover *statistical laws* of semantic change—for instance, that frequent words change meaning more slowly, and that polysemous words change faster—laws that echo, in the domain of meaning, the frequency–conservation result that Chapter 10 found for word forms [93]. Figure 14.2 illustrates the phenomenon: a word’s vector migrates across the decades as its sense shifts, so that “gay” drifts from the neighbourhood of *merry* toward that of *homosexual*, and “broadcast” from *scatter seed* toward *transmit signal*. The abstract semantic drift of the seed paper’s “awful” becomes, in this framework, a measurable trajectory through a vector space.

Table 14.1 gathers these computational approaches, each keyed to the scale it models and the phenomenon it captures, from the agent-based simulation of change to the neural modelling of the individual sequence processor to the

Table 14.1. Computational approaches to the LTC phenomena, with the scale each addresses, what it captures, and a representative method or result.

Approach	Scale	Captures	Example
Agent-based models	macro	change from inter-action	naming game (Steels)
Iterated learning	meso	emergence of structure	Kirby et al.
RNN / LSTM	micro	sequential syntax	Elman; Hochreiter
Transformer	micro	order as position	Vaswani et al.
Diachronic embeddings	macro	semantic drift	Hamilton et al.
Time-aware LMs	all	temporal adaptation	Lazaridou et al.

corpus measurement of drift.

14.5 Large Language Models as Temporal Artifacts

The most consequential computational objects of the present moment—the large language models now woven into daily life—are, from the LTC standpoint, profoundly *temporal* objects, and in a way their designers did not primarily intend. A large language model is trained on a corpus gathered up to a fixed *cutoff* date, and thereafter frozen. It is, in the most literal sense, a snapshot of linguistic time: a dense statistical fossil of how a language was used up to a certain moment, preserved unchanging while the living language moves on. This has a measurable consequence. Lazaridou and Lazaridou et al. [80] showed that language models degrade when asked to predict text from *after* their training period—and that the degradation worsens the further into the future one goes. The model ages, in effect, not by changing but by standing still while the world changes around it; its knowledge of language has a definite temporal centre and decays with distance from it. The very knowledge cutoff that any user of such a system encounters is a direct manifestation of the LTC thesis: language is a moving target, and a model that does not move with it falls behind.

This diagnosis also points to a prescription. The lesson of Chapter 12—that language is shaped by its learners over time—is now literally instantiated in these systems, whose form is shaped by the temporal distribution of their training data. A genuinely *time-aware* language technology would treat time not as an accident of the training set but as a modelled variable: systems that update continuously rather than freezing, that represent the multiple timescales

of Chapter 4 explicitly, and that can adapt their output to the temporal context of use—predictive text tuned to the user’s circadian state, tutoring systems that exploit the sleep-consolidation windows of Chapter 6, conversational systems whose register tracks the social time of Chapter 8. The LTC framework does not merely describe the temporality of language; it prescribes what a language technology built in full awareness of that temporality would look like.

14.6 Summary

Computation is the third mode of LTC inquiry, testing whether imagined mechanisms really produce observed histories. Agent-based models show how macro-scale change self-organises from local interaction; neural sequence models instantiate the individual processor of temporal syntax, and their evolution from recurrent networks to Transformers is a shift from representing time as process to representing it as position; iterated-learning simulations confirm that transmission through a bottleneck generates structure; and diachronic embeddings turn semantic drift into a measurable trajectory governed by discoverable laws. Most strikingly, the large language models of the present are themselves temporal artifacts—snapshots of linguistic time that decay as the living language moves past them—and the LTC framework both diagnoses this and prescribes the time-aware systems that would remedy it. With the framework thus extended into simulation and technology, one task remains: to gather the whole argument into a single picture and a research programme, which is the work of the final chapter.

Coda: Linguistics and the Brahma Kumaris

The image with which this chapter ends—the language model as a fossil, dense with knowledge yet frozen at its moment of training and slowly falling out of step with the living language—finds a pointed contrast in the Brahma Kumaris understanding of knowledge. Machine knowledge is stored and static: captured once, it can only decay. The tradition’s *gyan*, by contrast, is understood as living and self-renewing: the *Murli* is not archived and consulted but *spoken afresh each day*, re-given to the community morning by morning, so that knowledge is sustained not by storage but by continual re-transmission [23, 128]. There is a deeper resonance still. This chapter’s prescription was that a language system should *live in time*—update continually, move with the moving target—rather than freeze; and this is precisely the tradition’s own model of how truth is kept alive, through daily renewal rather than mere preservation. And in the cosmology’s largest frame, the *Drama* is the ultimate time-aware system: a knowledge that never needs updating because it is perfectly cyclic, the same

eternal script returning identically each *kalpa*. Between the frozen snapshot of the machine and the eternal recurrence of the cycle lies the living, daily renewal that both the contemplative and the temporal linguist, in their different idioms, hold up as the proper way for knowledge to inhabit time.

Toward a Unified Temporal Grammar

We set out to show that language keeps time. We end by proposing that time is the grammar within which all its other grammars are written.

This book began with a single conjecture—that language is organised by time, on many nested scales at once—and has spent fourteen chapters unfolding it. This final chapter gathers the argument into one picture, assesses what the framework has and has not achieved, states the positive proposal toward which the whole book has been building, and marks honestly the questions it leaves open, including the one on which its scientific and its contemplative readings most sharply diverge. It is a conclusion, and so it looks back; but a temporal linguistics, of all things, should also look forward, and it ends by sketching the research programme it invites.

15.1 The Argument in Brief

The claim was never that language merely *changes* in time—everyone knows that—but that it is *regulated by* time, at every scale on which it exists, and that these scales are coupled into a single system. Part I built the foundations: the temporal turn against the Saussurean snapshot (Chapter 1); the analogy with the regulated expression of the genome (Chapter 2); the question of where the linguistic kernel comes from, and its possible seat in the soul (Chapter 3); and the formalisation of the architecture as nested, coupled cycles (Chapter 4). Parts II, III, and IV then descended that architecture. At the *micro*-scale, language was shown to keep the hours of the circadian day, to be consolidated by sleep, and to run on the fast oscillations of rhythm and recursion (Chapters 5–7). At the *meso*-scale, it was shown to shift register across social time and to be reconstructed—never merely copied—by each generation of learners (Chapters 8–9). At the *macro*-scale, it was shown to descend in branching family trees, to structure discourse through

Table 15.1. The LTC framework at a glance: the three scales with their characteristic cycles, the phenomena treated at each, and the central open problem each poses.

Scale	Cycles	Phenomena (ch.)	Open problem
Micro	circadian, cognitive, rhythmic	expression clock, sleep, rhythm, recursion (4–7)	how fast biases become slow change
Meso	lifespan, generational, social	register, acquisition, transmission (8–9)	how competence is reconstructed and varied
Macro	historical, cultural	phylogeny, discourse, cultural evolution (10–12)	how the scales couple quantitatively

chronotope and deixis, and to evolve as the fastest of all cultural systems (Chapters 10–12). Part V turned to method: how the framework’s predictions can be experimentally tested (Chapter 13) and computationally modelled (Chapter 14). What remains is to see the whole as one.

15.2 The Unity of the Scales

The deepest claim of the book is not that language operates on several timescales—a truism—but that those timescales are *one coupled system*, and that the same phenomena recur, transformed, at each of them. Figure 15.1 draws the framework whole: the three nested scales, the phenomena that populate each, and the two directions of coupling that bind them—*accumulation* upward, by which fast biases compound into slow change, and *constraint* downward, by which slow norms shape fast choices. Table 15.1 sets out the same synthesis in words, pairing each scale with its characteristic cycles, its phenomena, and the central problem it leaves open.

The unity is not merely diagrammatic. Again and again a single phenomenon has proved to be one thing seen at different scales. *Entrenchment*—the storing of a frequent form as a resistant whole—is a fact about an individual’s memory (Chapter 4), a transmission bias in a community (Chapter 12), and an evolutionary conservation law across ten millennia (Chapter 10): the same mechanism, thrice named. *Entrainment* is the locking of a syllable to a beat (Chapter 4), the alignment of two speakers in dialogue (Chapter 4), and the condensation of a community onto a shared norm (Chapter 8). *Reconstruction*—the rebuilding of structure from partial evidence—is what sleep does to the day’s words each




Figure 15.1. The LTC framework as a whole. Language sits at the centre of three nested scales—micro, meso, and macro—each populated by the phenomena of the chapters shown. The scales are bound by two directions of coupling: accumulation (fast biases compounding into slow change) and constraint (slow norms shaping fast choices). The whole rests on the genetic analogy and the innate kernel of Part I.

night (Chapter 6) and what each generation does to its language across a lifetime (Chapter 9). To see these identities is to see why the scales are not four subjects but one.

15.3 The Analogy Assessed

The genetic analogy carried the argument, and in closing it deserves an honest audit. Where did it earn its keep, and where did it break? At the macro-scale it

was more than analogy: the mathematics of phylogenetics transferred wholesale, and linguistic descent is now studied with the biologist's own tools (Chapter 10). At the micro-scale it was genuinely generative, importing the experimental logic of chronobiology to ask how linguistic expression varies across the day, and finding, in the brain's nested oscillations, a striking vindication of the framework's central image (Chapters 5, 13). At the meso-scale it was suggestive but strained, for it was here that the disanalogies bit hardest.

And the disanalogies, it turned out, were among the framework's most useful findings. That linguistic transmission is *horizontal* as well as vertical, that it is *Lamarckian*, that it proceeds by *reconstruction* rather than copying—each of these, far from embarrassing the analogy, marked a precise respect in which language is *more* evolutionarily fluid than the genome, and each pointed to a real mechanism: borrowing, deliberate reform, the creative role of the learner. An analogy that failed informatively is worth more than one that succeeded blandly. The genetic parallel did what a good analogy should: it generated questions, transferred methods, and, where it broke, showed exactly where language is its own kind of thing.

15.4 What a Unified Temporal Grammar Would Be

We can now state the positive proposal toward which the whole book has pointed. A *unified temporal grammar* would be a theory of language whose primitives are not static rules but *coupled cycles*: a hierarchy of oscillators, from the syllabic to the civilizational, each entrained from within by the level below and modulated from without by the level above, with the formal apparatus of Appendix C—phase, entrainment, cross-frequency coupling—as its mathematical core. On such a view, a grammar is not a timeless system that happens to change but a standing pattern in a flow, the momentary cross-section of a set of processes running at every scale at once.

Such a grammar would dissolve two old oppositions. It would dissolve the *synchronic–diachronic* divide with which Chapter 1 began, since on this view the synchronic state simply *is* the diachronic process viewed at an instant. And it would reframe the *universals–diversity* debate that has divided the field [52]: the universals of language would be the properties every coupled system of this kind must share—the constraints of learnability, memory, and entrainment that hold at every scale—while its diversity would be the many stable settings such a system can occupy, the different phases onto which different communities have locked. Universality and diversity would both be consequences of the same temporal architecture. That is the theory the framework gestures toward; building it in full is the work of a field, not a book.

15.5 Open Problems and a Research Programme

That work has a shape, sketched in Figure 15.2. It begins with *measurement*: charting, for each linguistic module, its phase and amplitude across the scales, using the instruments of Chapter 13. It proceeds to *modelling*: building the coupled-oscillator and iterated-learning systems of Chapter 14 that can reproduce the couplings quantitatively. It aims at *unification*: a formal theory of cross-scale coupling that would, among other things, finally engage the *actuation problem*—why a perpetually available micro-scale variation occasionally, and only occasionally, propagates up into macro-scale change—which has stood unsolved since Chapter 4 and remains, in the author’s view, the single deepest open question in the study of language in time. And it issues in *application*: language technology built, at last, in full awareness that its object is a moving target.

A RESEARCH ROADMAP FOR TEMPORAL LINGUISTICS




Figure 15.2. A research programme for a temporal linguistics: from measuring the phase and amplitude of linguistic cycles, through modelling their coupling, to a unified theory of cross-scale interaction and its application in time-aware technology.

Other problems stand alongside it: the *unit problem*—what, exactly, is the replicating lingueme—which the whole book has treated as multi-level and never fully resolved; the quantification of coupling strength between adjacent scales; and the integration of the framework’s biological, cognitive, social, and cultural strands into a single formalism. None of these is small. But a framework

is fruitful not when it answers every question but when it poses the right ones sharply, and the sharpness of these questions is, we submit, the framework's best credential.

15.6 The Two Readings

This book has been written in two registers, and honesty requires acknowledging their relation plainly. Most of it is mainstream science: testable, cited, and answerable to evidence. But at certain points—the seat of the linguistic kernel in the soul (Chapter 3), the reading of divergence against the cycle of time, the coda to each chapter—it has entertained a second, contemplative reading drawn from the Brahma Kumaris understanding from which the framework takes its name. We have tried never to blur the two: to present the science as science and the metaphysics as a clearly-marked interpretive proposal, and, where they conflict—most sharply over the timescales of Chapter 10, where phylogenetic dating far exceeds the five-thousand year cycle—to state the tension rather than dissolve it. The reader is free to take the temporal linguistics without the metaphysics; the framework stands as science on its own. But the book has also wished to show that a disciplined science of language in time and a contemplative vision of time itself can be held in the same mind without either being falsified, each illuminating the other at the edges. Whether that is a strength or a stretch, the reader must judge; we have preferred to make the double vision visible rather than to hide half of it.

15.7 Envoi

Language is the most temporal thing we do. It happens in time, is built of time, and carries time within it at every scale—the syllable's beat, the day's rise and fall, the lifetime's arc, the millennium's slow drift. To study it as if it stood still, as the grammars have largely done, is to study the melody by staring at a single note. This book has tried to hear the whole tune: to show that the note, the phrase, and the symphony are one music, played on nested clocks, and that language is what it is because it moves. If the framework is right, then time is not the medium in which language happens but the grammar in which its grammar is written—and to speak, in the end, is to keep time.

Coda: Linguistics and the Brahma Kumaris

It is fitting that a book named for the time cycle should close with the teaching that gave it that name. Every chapter's coda has traced a single correspondence: that the temporal structure the linguist finds in language—nesting, entrainment,

return—is, in the Brahma Kumaris vision, the temporal structure of reality itself. The syllable returning in its metre, the day returning at *Amrit Vela*, the language returning to its root, the soul returning through its births, the whole world returning through the *kalpa*: these are, on the reading this book has entertained, the same figure of nested recurrence written at every magnitude, from the beat of a word to the turning of the cosmos. The linguist measures the smallest of these wheels; the tradition contemplates the largest; and the wager of the LTC framework has been that they are wheels within one wheel. Whether or not the reader accepts the largest claim, the smallest ones stand, and they point the same way: that to exist in time is to move in cycles, and that language, our most human possession, is the moving image of that moving image—time made audible, and, for a moment, made meaning. *Om Shanti*.

Glossary of Core LTC Terms

This appendix collects, for quick reference, the technical terms introduced across the book, each with a compact definition as it is used within the Linguistics in the Time Cycle (LTC) framework. Fuller discussion appears at the point of first use, indexed at the back of the volume. The glossary will be expanded as later chapters introduce further vocabulary; the entries below anchor the terms established in Part I.

Micro-cycle A fast temporal cycle—on the scale of milliseconds to a day—governing the moment-to-moment expression of language, including circadian and cognitive rhythms.

Meso-cycle An intermediate cycle—years to a lifetime—governing the transmission of language across development and generations.

Macro-cycle A slow cycle—centuries to millennia—governing the historical and cultural transformation of a language.

Coupling The dependence of one temporal scale on another, whereby fast-scale biases accumulate into slow-scale change (upward coupling) and slow-scale facts constrain fast-scale choices (downward coupling).

Entrainment The synchronisation of an internal linguistic rhythm to an external periodic signal, by analogy with the entrainment of the body clock to daylight.

The Full GTC–LTC Correspondence

The genetic analogy that motivates the LTC framework is developed across the main text—introduced in Chapter 1, elaborated term by term in Chapter 2, and formalised as an architecture of coupled cycles in Chapter 4. This appendix gathers the complete correspondence between the Genetic Time Cycle (GTC) and the Linguistic Time Cycle (LTC) into a single reference, so that the parallel can be seen whole. The correspondence is organised into four blocks—the units of the code, its regulated expression, the forces of variation and selection, and the mechanics of transmission—followed by a statement of the systematic *disanalogies*, which are as informative as the parallels themselves. The table is a map, not a claim of literal identity; the cautions of Section 1.7 apply throughout.

The Systematic Disanalogies

An honest correspondence must record where it fails, for each failure marks a place where language is its own kind of evolving system and the biological toolkit must be adapted rather than merely applied (Section 2.5). Three differences are fundamental.

Mode of transmission. Genetic inheritance is overwhelmingly *vertical*, from ancestor to descendant; linguistic transmission is also promiscuously *horizontal*, through borrowing and contact between coexisting languages. A language genealogy is therefore a tree crossed by a network, not a pure tree.

Lamarckian change. Acquired linguistic characteristics are inherited: an innovation coined within a speaker’s lifetime can pass directly into the language, with no analogue of the Weismann barrier that in biology insulates the germ line. Change can even be *directed*, as in deliberate reform or revival.

Porous genotype–phenotype boundary. In the organism the gene and the body are sharply separated; in language, competence (“code”) and utterance (“expression”) are continuously coupled, since the utterances a learner hears are exactly what build her competence. The feedback the central dogma forbids in biology is, in language, the very mechanism of acquisition.

Table B.1. The full correspondence between the Genetic and Linguistic Time Cycles, organised by function. Each genetic element is paired with its linguistic counterpart and the mechanism through which the linguistic version operates.

Genetics (GTC)	Linguistics (LTC)	Mechanism / note
<i>Units and code</i>		
Nucleotide base	Phonological feature	smallest contrastive unit
Codon	Phoneme	minimal combinatorial unit
Gene	Morpheme / construction	minimal meaningful unit
Genome	Grammatical competence	the stored system (<i>langue</i>)
<i>Expression and regulation</i>		
Transcription	Encoding of an utterance	mapping meaning to form
Translation	Articulation	plan realised as speech/text
Protein	Utterance	the overt expressed product
Promoter / enhancer	Syntax	licenses and structures output
Regulation, timing	Pragmatics, circadian phase	context gates expression
Epigenetic marks	Register, style	durable settings on expression
<i>Variation and selection</i>		
Point mutation	Innovation / neologism	error, analogy, reanalysis
Recombination	Blending, borrowing	new combinations of material
Genetic drift	Lexical diffusion	chance change in frequency
Selection pressure	Social prestige, utility	differential replication
Replicator (gene)	Lingueme (utterance token)	the entity that is copied
Interactor (body)	Speaker	the entity that interacts
<i>Transmission</i>		
Heredity	Acquisition	code rebuilt in each learner
Molecular clock	Lexical clock (Swadesh)	change rate dates divergence
Phylogeny	Language family tree	branching descent
Speciation	Language split	loss of mutual intelligibility

These disanalogies are not defects of the framework but signposts within it. They indicate, in each case, that language is *more* evolutionarily fluid than the genome—faster to borrow, quicker to inherit the acquired, and less firmly divided between what is stored and what is expressed.

Mathematical Notes: Coupled Oscillators and Entrainment

Several chapters—most directly Chapter 4—appeal to the language of oscillation, phase, and entrainment to make the notions of rhythm, synchrony, nesting, and coupling precise. The main text deliberately keeps that vocabulary informal, so as to remain readable to linguists. This appendix collects the minimal mathematical apparatus behind it, for the reader who wishes to see the formal skeleton. Nothing here is original; it is the standard theory of coupled phase oscillators, stated compactly and with linguistic interpretation attached.

C.1 The phase of a limit-cycle oscillator

A self-sustained rhythm—a beating heart, a circadian clock, or, we propose, the rhythmic generator underlying speech—is modelled as a *limit-cycle oscillator*: a system that settles onto a closed loop in its state space and circulates around it. On that loop the entire state can be summarised by a single *phase* variable $\theta \in [0, 2\pi)$, which advances at a constant natural rate ω , the oscillator’s angular frequency:

$$\frac{d\theta}{dt} = \omega, \quad \omega = \frac{2\pi}{T}, \quad (\text{C.1})$$

where T is the period. The phase reduction is what makes rhythmic phenomena tractable: however complicated the underlying mechanism, its timing is captured by θ . In the linguistic setting, the syllable-scale oscillator of Section 4.2 has a period T of a couple of hundred milliseconds; the circadian expression clock of Chapter 5 has $T \approx 24$ hours.

C.2 Entrainment by a periodic force

An oscillator rarely runs in isolation; it is driven. Let a periodic input of frequency Ω act on an oscillator of natural frequency ω , coupling to it with strength K . Writing the phase difference $\phi = \theta - \Omega t$, the classic reduced dynamics is the

Adler equation,

$$\frac{d\phi}{dt} = \Delta\omega - K \sin \phi, \quad \Delta\omega = \omega - \Omega. \quad (\text{C.2})$$

A stable fixed point—*phase locking*, or *entrainment*—exists precisely when the detuning is not too large relative to the coupling:

$$|\Delta\omega| \leq K. \quad (\text{C.3})$$

The set of $(\Delta\omega, K)$ for which locking occurs forms the wedge-shaped region known as an *Arnold tongue*. This is the formal content of the claim that a speaker's rhythm can be *entrained*: within the tongue, the internal oscillator abandons its own frequency and adopts that of the driver—daylight entraining the body clock, or a metronomic beat pulling speech onto it. Outside the tongue, the phase slips and entrainment fails.

C.3 Coupled oscillators: the Kuramoto model

Dialogue and community are not one oscillator and a fixed drive but *many oscillators driving one another*. The canonical model of such mutual synchronisation is due to Kuramoto. For N oscillators with phases θ_i and natural frequencies ω_i , coupled all-to-all with strength K ,

$$\frac{d\theta_i}{dt} = \omega_i + \frac{K}{N} \sum_{j=1}^N \sin(\theta_j - \theta_i), \quad i = 1, \dots, N. \quad (\text{C.4})$$

The degree of collective synchrony is measured by the complex *order parameter*

$$r e^{i\psi} = \frac{1}{N} \sum_{j=1}^N e^{i\theta_j}, \quad (\text{C.5})$$

where $r = 0$ denotes incoherence and $r = 1$ perfect synchrony. Kuramoto's central result is a *phase transition*: if the natural frequencies are spread with a distribution of width comparable to $g(0)$, synchrony emerges only once the coupling exceeds a critical value,

$$K_c = \frac{2}{\pi g(0)}, \quad (\text{C.6})$$

with g the density of the frequency distribution about its mean. Below K_c the population drifts incoherently; above it, a macroscopic cluster locks into a common rhythm. This is the mathematical image behind the sociolinguistic claim of Section 4.3 that the private rhythms of individual speakers can, once interaction is strong enough, condense into the shared cycles of a speech community—and behind the “phase transitions” invoked for register change and the actuation of sound change.

C.4 Nesting as a hierarchy of frequencies

Finally, the *nesting* of Section 4.4 is expressed formally as a separation of timescales. Let a fast oscillator of frequency ω_f be modulated by a slow one of frequency ω_s with $\omega_f \gg \omega_s$. Their interaction is legible in the spectrum as *cross-frequency coupling*: the amplitude or phase of the fast rhythm is organised by the phase of the slow one, so that many fast cycles are enveloped within each slow cycle. Writing the ratio

$$n = \frac{\omega_f}{\omega_s} \gg 1, \quad (\text{C.7})$$

n is the number of fast periods per slow period—the syllables per breath group, the utterances per conversation, the diurnal cycles per lifetime of Section 4.1. The whole LTC architecture is, in these terms, a hierarchy of oscillators whose frequency ratios n span the twelve orders of magnitude of Figure 4.1, each level entrained from within by the level below and modulated from without by the level above.

Corpora, Datasets, and Software for Temporal Linguistics

Empirical and computational work on the LTC framework draws on a scattered body of resources: diachronic corpora, time-stamped social-media archives, chronobiological datasets, and software for phylogenetic and agent-based modelling. This appendix catalogues the principal resources referred to in Parts IV and V, with notes on their coverage, access, and suitability for temporal-linguistic analysis.

To be completed alongside Chapters 9, 12, and 13.

Glossary of Brahma Kumaris Terms

The temporal framework of this book intersects, at several points, with the teaching of the Brahma Kumaris on the cycle of time. That teaching has a precise technical vocabulary of its own, much of it Hindi or Sanskrit, and many of its words carry specialised meanings that differ from ordinary usage—meanings fixed above all in the daily discourses known as the *Murlis*, in both their corporeal (*Sakar*) and subtle (*Avyakt*) forms [7, 23, 126, 128]. This appendix glosses the terms used in the book, in a simplified transliteration that omits diacritics. It is offered as an aid to readers unfamiliar with the tradition, not as a doctrinal statement; fuller treatments are given in the scholarly literature cited.

Amrit Vela The “nectar hour”: the pre-dawn period (about 4:00–4:45 a.m.) set aside for the deepest daily remembrance. Of direct interest to the circadian argument of Chapter 5.

Atma The soul: an eternal, conscious point of light, the seat of identity, understood to be located in the centre of the forehead.

Avyakt “Unmanifest,” subtle, incorporeal. The *Avyakt Murlis* are the discourses spoken in subtle form from 1969 onward, as against the earlier corporeal ones.

BapDada The combined form addressed in the discourses—the Supreme Father (Shiv Baba) together with the corporeal founder (Brahma Baba).

Brahma Baba The corporeal founder of the Brahma Kumaris and the medium through whom the Sakar Murlis were spoken (1936–1969).

Brahmin A member of the Brahma Kumaris spiritual community; a “child of Brahma,” living at the Confluence Age.

Dharna The imbibing and sustaining, in daily conduct, of divine knowledge and virtue—the internalisation of what the Murli teaches.

Drama The eternal, predestined “World Drama”: the beginningless and endless play of souls and matter that repeats identically each cycle.

Drishti Spiritual vision; the soul-conscious “seeing” exchanged as a silent blessing.

- Dwapar** The Copper Age, third of the four ages, when soul-connection breaks and divergence—including, on this book’s reading, the divergence of languages—begins.
- Gyan** Divine knowledge; the doctrinal content of the Murli.
- Kalpa** The world cycle of 5,000 years; the *Kalpa Vriksha* is the “world tree” of humanity, an image echoed in the language tree of Chapter 3.
- Kaliyug** The Iron Age, the final and most degraded age, of maximal diversity and separation.
- Karmateet** The stage “beyond karma,” free of the effects of action; the perfected condition of the soul.
- Madhuban** The “forest of honey”: the Brahma Kumaris headquarters at Mount Abu, including Gyan Sarovar and Pandav Bhawan.
- Manmanabhav** The core instruction “merge your mind in Me”—to remember the One with an undivided mind.
- Mansa, Vacha, Karmana** Thought, word, and deed: the three registers through which action, and karma, are performed. *Vacha* (word/speech) is the register of most direct linguistic interest.
- Murli** Literally “flute”; the daily spoken discourse of Shiv Baba, in its *Sakar* (corporeal) and *Avyakt* (subtle) forms. The primary corpus of the tradition’s specialised vocabulary.
- Nirakar** The incorporeal, the formless; an epithet of the Supreme Soul.
- Om Shanti** “I am a peaceful soul”; both the community’s greeting and its fundamental affirmation of spiritual identity.
- Paramdham (Shantidham)** The “supreme abode” or “abode of peace”: the incorporeal world of light and silence, home of the souls and of the Supreme.
- Raja Yoga** “Kingly union”: the Brahma Kumaris practice of meditative remembrance of the Supreme Soul.
- Sakar** Corporeal, manifest; the *Sakar Murlis* are those spoken through Brahma Baba in bodily form.
- Sangam Yug** The Confluence Age: the short, auspicious age between Kaliyug and Satyug in which God is said to descend and the cycle turns. Also *Purushottam Sangam Yug*, the “most elevated” Confluence Age.
- Sanskar** An impression, tendency, or trait carried by the soul and persisting across births—the seat of character. Of direct relevance to the soul-borne linguistic “kernel” and to xenoglossy in Chapter 3.

Satyug The Golden Age, first and most perfect of the ages, of unity and (on this book's reading) a single, undivided tongue.

Shiv Baba The Supreme Soul, the incorporeal Father, conceived as an infinitesimal point of divine light.

Shrimat God's "elevated directions," contrasted with *manmat* (the dictates of one's own mind) and *parmat* (the dictates of others). A threefold contrast of *directive* language.

Sukshma Vatan The subtle region of light—abode of the deity forms Brahma, Vishnu, and Shankar—in which avyakt speech is received.

Treta The Silver Age, the second age, in which souls remain connected.

Trikaldarshi The "seer of the three aspects of time"—past, present, and future—one who knows the whole Drama at once.

Trimurti The three subtle deity forms (Brahma, Vishnu, Shankar) through which the incorporeal Supreme is said to act.

Vinash The destruction at the close of Kaliyug by which the cycle is renewed.

Yaad Remembrance; the continuous linking of the mind to the Supreme.

Bibliography

- [1] David Abercrombie. *Elements of General Phonetics*. Edinburgh University Press, Edinburgh, 1967.
- [2] Alexandra Y. Aikhenvald. *Evidentiality*. Oxford University Press, Oxford, 2004.
- [3] Jean Aitchison. *Language Change: Progress or Decay?* Cambridge University Press, Cambridge, 3rd edition, 2001.
- [4] Quentin D. Atkinson and Russell D. Gray. Curious parallels and curious connections: Phylogenetic thinking in biology and historical linguistics. *Systematic Biology*, 54(4):513–526, 2005.
- [5] Quentin D. Atkinson, Andrew Meade, Chris Venditti, Simon J. Greenhill, and Mark Pagel. Languages evolve in punctuational bursts. *Science*, 319 (5863):588, 2008.
- [6] J. L. Austin. *How to Do Things with Words*. Oxford University Press, Oxford, 1962.
- [7] Lawrence A. Babb. *Redemptive Encounters: Three Modern Styles in the Hindu Tradition*. University of California Press, Berkeley, 1986.
- [8] Mikhail M. Bakhtin. *The Dialogic Imagination: Four Essays*. University of Texas Press, Austin, 1981. Ed. M. Holquist; trans. C. Emerson and M. Holquist.
- [9] Gareth J. Baxter, Richard A. Blythe, William Croft, and Alan J. McKane. Utterance selection model of language change. *Physical Review E*, 73(4): 046118, 2006.
- [10] Clay Beckner, Richard Blythe, Joan Bybee, Morten H. Christiansen, William Croft, Nick C. Ellis, John Holland, Jinyun Ke, Diane Larsen-Freeman, and Tom Schoenemann. Language is a complex adaptive system: Position paper. *Language Learning*, 59(s1):1–26, 2009.
- [11] Allan Bell. Language style as audience design. *Language in Society*, 13(2): 145–204, 1984.

-
- [12] Luciano Bernardi, Peter Sleight, Gabriele Bandinelli, Simone Cencetti, Lamberto Fattorini, Johanna Wdowczyk-Szulc, and Alfonso Lagi. Effect of rosary prayer and yoga mantras on autonomic cardiovascular rhythms: Comparative study. *BMJ*, 323(7327):1446–1449, 2001.
 - [13] Derek Bickerton. The language bioprogram hypothesis. *Behavioral and Brain Sciences*, 7(2):173–221, 1984.
 - [14] Susan Blackmore. *The Meme Machine*. Oxford University Press, Oxford, 1999.
 - [15] Jan Blommaert. *Discourse: A Critical Introduction*. Cambridge University Press, Cambridge, 2005.
 - [16] Jan Blommaert. Sociolinguistic scales. *Intercultural Pragmatics*, 4(1):1–19, 2007.
 - [17] J. Kathryn Bock. Syntactic persistence in language production. *Cognitive Psychology*, 18(3):355–387, 1986.
 - [18] Alexander A. Borbély. A two process model of sleep regulation. *Human Neurobiology*, 1(3):195–204, 1982.
 - [19] Jan Born and Ines Wilhelm. System consolidation of memory during sleep. *Psychological Research*, 76(2):192–203, 2012.
 - [20] Remco Bouckaert, Philippe Lemey, Michael Dunn, Simon J. Greenhill, Alexander V. Alekseyenko, Alexei J. Drummond, Russell D. Gray, Marc A. Suchard, and Quentin D. Atkinson. Mapping the origins and expansion of the indo-european language family. *Science*, 337(6097):957–960, 2012.
 - [21] Pierre Bourdieu. *Language and Symbolic Power*. Harvard University Press, Cambridge, MA, 1991.
 - [22] Robert Boyd and Peter J. Richerson. *Culture and the Evolutionary Process*. University of Chicago Press, Chicago, 1985.
 - [23] Brahma Kumaris World Spiritual University. The murlis: Daily discourses of shiv baba (sakar and avyakt versions), 2020. Prajapita Brahma Kumaris Ishwariya Vishwa Vidyalaya, Mount Abu, India; Sakar Murlis spoken 1936–1969, Avyakt Murlis from 1969.
 - [24] Joan Bybee. From usage to grammar: The mind’s response to repetition. *Language*, 82(4):711–733, 2006.

- [25] Joan Bybee. *Language, Usage and Cognition*. Cambridge University Press, Cambridge, 2010.
- [26] Lyle Campbell. *Historical Linguistics: An Introduction*. MIT Press, Cambridge, MA, 3rd edition, 2013.
- [27] Piotr Cap. *Proximization: The Pragmatics of Political Discourse*. John Benjamins, Amsterdam, 2013.
- [28] Luigi Luca Cavalli-Sforza and Marcus W. Feldman. *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton University Press, Princeton, 1981.
- [29] Noam Chomsky. *Syntactic Structures*. Mouton, The Hague, 1957.
- [30] Noam Chomsky. *Aspects of the Theory of Syntax*. MIT Press, Cambridge, MA, 1965.
- [31] Morten H. Christiansen and Nick Chater. Language as shaped by the brain. *Behavioral and Brain Sciences*, 31(5):489–509, 2008.
- [32] Bernard Comrie. *Tense*. Cambridge University Press, Cambridge, 1985.
- [33] Francis Crick. Central dogma of molecular biology. *Nature*, 227(5258): 561–563, 1970.
- [34] William Croft. *Explaining Language Change: An Evolutionary Approach*. Longman, Harlow, 2000.
- [35] David Crystal. *Language Death*. Cambridge University Press, Cambridge, 2000.
- [36] David Crystal. *Language and the Internet*. Cambridge University Press, Cambridge, 2nd edition, 2006.
- [37] Fred Cummins and Robert Port. Rhythmic constraints on stress timing in english. *Journal of Phonetics*, 26(2):145–171, 1998.
- [38] Anne Cutler. *Slips of the Tongue and Language Production*. Mouton, Berlin, 1982.
- [39] Östen Dahl. *Tense and Aspect Systems*. Blackwell, Oxford, 1985.
- [40] Charles Darwin. *The Descent of Man, and Selection in Relation to Sex*. John Murray, London, 1871.
- [41] Richard Dawkins. *The Selfish Gene*. Oxford University Press, Oxford, 1976.

-
- [42] Jan Peter de Ruiter, Holger Mitterer, and N. J. Enfield. Projecting the end of a speaker's turn: A cognitive cornerstone of conversation. *Language*, 82 (3):515–535, 2006.
 - [43] Ferdinand de Saussure. *Course in General Linguistics*. Payot, Paris, 1916. Ed. C. Bally and A. Sechehaye; trans. W. Baskin, 1959.
 - [44] Terrence W. Deacon. *The Symbolic Species: The Co-Evolution of Language and the Brain*. W. W. Norton, New York, 1997.
 - [45] Nai Ding, Lucia Melloni, Hang Zhang, Xing Tian, and David Poeppel. Cortical tracking of hierarchical linguistic structures in connected speech. *Nature Neuroscience*, 19(1):158–164, 2016.
 - [46] Jeanne F. Duffy and Derk-Jan Dijk. Getting through to circadian oscillators: Why use constant routines? *Journal of Biological Rhythms*, 17(1):4–13, 2002.
 - [47] Nicolas Dumay and M. Gareth Gaskell. Sleep-associated changes in the mental representation of spoken words. *Psychological Science*, 18(1):35–39, 2007.
 - [48] David M. Eberhard, Gary F. Simons, and Charles D. Fennig. *Ethnologue: Languages of the World*. SIL International, Dallas, 26th edition, 2023.
 - [49] Penelope Eckert. Age as a sociolinguistic variable. In Florian Coulmas, editor, *The Handbook of Sociolinguistics*, pages 151–167. Blackwell, Oxford, 1997.
 - [50] Penelope Eckert. *Linguistic Variation as Social Practice*. Blackwell, Oxford, 2000.
 - [51] Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.
 - [52] Nicholas Evans and Stephen C. Levinson. The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, 32(5):429–448, 2009.
 - [53] Daniel L. Everett. Cultural constraints on grammar and cognition in pirahã: Another look at the design features of human language. *Current Anthropology*, 46(4):621–646, 2005.
 - [54] Charles J. Fillmore. *Lectures on Deixis*. CSLI Publications, Stanford, 1997.
 - [55] Gérard Genette. *Narrative Discourse: An Essay in Method*. Cornell University Press, Ithaca, NY, 1980. Trans. J. E. Lewin; French original 1972.

- [56] Edward Gibson. Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68(1):1–76, 1998.
- [57] Howard Giles, Justine Coupland, and Nikolas Coupland. *Contexts of Accommodation: Developments in Applied Sociolinguistics*. Cambridge University Press, Cambridge, 1991.
- [58] Anne-Lise Giraud and David Poeppel. Cortical oscillations and speech processing: Emerging computational principles and operations. *Nature Neuroscience*, 15(4):511–517, 2012.
- [59] Scott A. Golder and Michael W. Macy. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881, 2011.
- [60] Susan Goldin-Meadow. *The Resilience of Language: What Gesture Creation in Deaf Children Can Tell Us About How All Children Learn Language*. Psychology Press, New York, 2003.
- [61] Rebecca L. Gómez, Richard R. Bootzin, and Lynn Nadel. Naps promote abstraction in language-learning infants. *Psychological Science*, 17(8):670–674, 2006.
- [62] Russell D. Gray and Quentin D. Atkinson. Language-tree divergence times support the anatolian theory of indo-european origin. *Nature*, 426(6965):435–439, 2003.
- [63] Wolfgang Haak, Iosif Lazaridis, Nick Patterson, Nadin Rohland, Swapan Mallick, Bastien Llamas, et al. Massive migration from the steppe was a source for indo-european languages in europe. *Nature*, 522(7555):207–211, 2015.
- [64] M. A. K. Halliday. *Language as Social Semiotic: The Social Interpretation of Language and Meaning*. Edward Arnold, London, 1978.
- [65] William L. Hamilton, Jure Leskovec, and Dan Jurafsky. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1489–1501, 2016.
- [66] Yvonne Harrison and James A. Horne. Sleep loss impairs short and novel language tasks having a prefrontal focus. *Journal of Sleep Research*, 7(2):95–100, 1998.

-
- [67] Marc D. Hauser, Noam Chomsky, and W. Tecumseh Fitch. The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298(5598): 1569–1579, 2002.
 - [68] Paul Heggarty, Cormac Anderson, Matthew Scarborough, Benedict King, Remco Bouckaert, et al. Language trees with sampled ancestors support a hybrid model for the origin of indo-european languages. *Science*, 381 (6656):eabg0818, 2023.
 - [69] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
 - [70] David L. Hull. *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. University of Chicago Press, Chicago, 1988.
 - [71] Eva Jablonka and Marion J. Lamb. *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. MIT Press, Cambridge, MA, 2005.
 - [72] Roman Jakobson. Closing statement: Linguistics and poetics. In Thomas A. Sebeok, editor, *Style in Language*, pages 350–377. MIT Press, Cambridge, MA, 1960.
 - [73] Simon Kirby. Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110, 2001.
 - [74] Simon Kirby, Hannah Cornish, and Kenny Smith. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, 105(31):10681–10686, 2008.
 - [75] Anthony S. Kroch. Reflexes of grammar in patterns of language change. *Language Variation and Change*, 1(3):199–244, 1989.
 - [76] William Labov. *Principles of Linguistic Change, Volume I: Internal Factors*. Blackwell, Oxford, 1994.
 - [77] William Labov. *Principles of Linguistic Change, Volume 2: Social Factors*. Blackwell, Oxford, 2001.
 - [78] William Labov and Joshua Waletzky. Narrative analysis: Oral versions of personal experience. In June Helm, editor, *Essays on the Verbal and Visual Arts*, pages 12–44. University of Washington Press, Seattle, 1967.

- [79] Cecilia S. L. Lai, Simon E. Fisher, Jane A. Hurst, Faraneh Vargha-Khadem, and Anthony P. Monaco. A forkhead-domain gene is mutated in a severe speech and language disorder. *Nature*, 413(6855):519–523, 2001.
- [80] Angeliki Lazaridou, Adhiguna Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kocisky, Sebastian Ruder, et al. Mind the gap: Assessing temporal generalization in neural language models. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, 2021.
- [81] Eric H. Lenneberg. *Biological Foundations of Language*. Wiley, New York, 1967.
- [82] Stephen C. Levinson. *Pragmatics*. Cambridge University Press, Cambridge, 1983.
- [83] Cynthia P. May and Lynn Hasher. Synchrony effects in inhibitory control over thought and action. *Journal of Experimental Psychology: Human Perception and Performance*, 25(2):363–379, 1999.
- [84] James L. McClelland, Bruce L. McNaughton, and Randall C. O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, 1995.
- [85] Matthias R. Mehl. The electronically activated recorder (ear): A method for the naturalistic observation of daily social behavior. *Current Directions in Psychological Science*, 26(2):184–190, 2017.
- [86] Alex Mesoudi. *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*. University of Chicago Press, Chicago, 2011.
- [87] Jelena Mirković and M. Gareth Gaskell. Does sleep improve your grammar? preferential consolidation of arbitrary components of new linguistic knowledge. *PLoS ONE*, 6(4):e18001, 2011.
- [88] Padraic Monaghan, Christine Schoetensack, and Patrick Rebuschat. Sleep and second language acquisition: The role of chunking and associative learning. *Language Learning*, 67(S1):192–214, 2017.
- [89] Marina Nespors and Irene Vogel. *Prosodic Phonology*. Foris, Dordrecht, 1986.
- [90] Elissa L. Newport. Maturation constraints on language learning. *Cognitive Science*, 14(1):11–28, 1990.

-
- [91] John J. Ohala. The phonetics of sound change. In Charles Jones, editor, *Historical Linguistics: Problems and Perspectives*, pages 237–278. Longman, London, 1993.
- [92] Mark Pagel. Human language as a culturally transmitted replicator. *Nature Reviews Genetics*, 10(6):405–415, 2009.
- [93] Mark Pagel, Quentin D. Atkinson, and Andrew Meade. Frequency of word-use predicts rates of lexical evolution throughout indo-european history. *Nature*, 449(7163):717–720, 2007.
- [94] Martin J. Pickering and Simon Garrod. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(2):169–226, 2004.
- [95] Kenneth L. Pike. *The Intonation of American English*. University of Michigan Press, Ann Arbor, 1945.
- [96] Steven Pinker. *The Language Instinct: How the Mind Creates Language*. William Morrow, New York, 1994.
- [97] Franck Ramus, Marina Nespor, and Jacques Mehler. Correlates of linguistic rhythm in the speech signal. *Cognition*, 73(3):265–292, 1999.
- [98] Björn Rasch and Jan Born. About sleep’s role in memory. *Physiological Reviews*, 93(2):681–766, 2013.
- [99] David Reich. *Who We Are and How We Got Here: Ancient DNA and the New Science of the Human Past*. Pantheon Books, New York, 2018.
- [100] Hans Reichenbach. *Elements of Symbolic Logic*. Macmillan, New York, 1947.
- [101] Peter J. Richerson and Robert Boyd. *Not by Genes Alone: How Culture Transformed Human Evolution*. University of Chicago Press, Chicago, 2005.
- [102] Till Roenneberg, Tim Kuehnle, Myriam Juda, Thomas Kantermann, Karla Allebrandt, Marijke Gordijn, and Martha Merrow. Epidemiology of the human circadian clock. *Sleep Medicine Reviews*, 11(6):429–438, 2007.
- [103] Jessica Rosenberg, Kathrin Pusch, and Rainer Dietrich. The tick-tock of language: Is language processing sensitive to circadian rhythmicity and elevated sleep pressure? *Chronobiology International*, 26(5):771–796, 2009.
- [104] Jessica Rosenberg, Martina Reske, Tracy Warbrick, and N. Jon Shah. Chronotype modulates language processing-related cerebral activity during functional mri (fmri). *PLoS ONE*, 10(9):e0137197, 2015.

- [105] Brandon C. Roy, Michael C. Frank, Philip DeCamp, Matthew Miller, and Deb Roy. Predicting the birth of a spoken word. *Proceedings of the National Academy of Sciences*, 112(41):12663–12668, 2015.
- [106] David C. Rubin. *Memory in Oral Traditions: The Cognitive Psychology of Epic, Ballads, and Counting-Out Rhymes*. Oxford University Press, New York, 1995.
- [107] Merritt Ruhlen. *The Origin of Language: Tracing the Evolution of the Mother Tongue*. John Wiley, New York, 1994.
- [108] Harvey Sacks, Emanuel A. Schegloff, and Gail Jefferson. A simplest systematics for the organization of turn-taking for conversation. *Language*, 50(4):696–735, 1974.
- [109] Jenny R. Saffran, Richard N. Aslin, and Elissa L. Newport. Statistical learning by 8-month-old infants. *Science*, 274(5294):1926–1928, 1996.
- [110] Gillian Sankoff and Hélène Blondeau. Language change across the lifespan: /r/ in montreal french. *Language*, 83(3):560–588, 2007.
- [111] Christina Schmidt, Fabienne Collette, Christian Cajochen, and Philippe Peigneux. A time to think: Circadian rhythms in human cognition. *Cognitive Neuropsychology*, 24(7):755–789, 2007.
- [112] Thomas Schreiner and Björn Rasch. Boosting vocabulary learning by verbal cueing during sleep. *Cerebral Cortex*, 25(11):4169–4179, 2015.
- [113] Ann Senghas, Sotaro Kita, and Asli Özyürek. Children creating core properties of language: Evidence from an emerging sign language in nicaragua. *Science*, 305(5691):1779–1782, 2004.
- [114] Michael Silverstein. Indexical order and the dialectics of sociolinguistic life. *Language and Communication*, 23(3–4):193–229, 2003.
- [115] Jenny L. Singleton and Elissa L. Newport. When learners surpass their models: The acquisition of american sign language from inconsistent input. *Cognitive Psychology*, 49(4):370–407, 2004.
- [116] Luc Steels. Modeling the cultural evolution of language. *Physics of Life Reviews*, 8(4):339–356, 2011.
- [117] Ian Stevenson. *Xenoglossy: A Review and Report of a Case*. University Press of Virginia, Charlottesville, 1974.

-
- [118] Ian Stevenson. *Unlearned Language: New Studies in Xenoglossy*. University Press of Virginia, Charlottesville, 1984.
 - [119] Tanya Stivers, N. J. Enfield, Penelope Brown, Christina Englert, Makoto Hayashi, Trine Heinemann, Gertie Hoymann, Federico Rossano, Jan Peter de Ruiter, Kyung-Eun Yoon, and Stephen C. Levinson. Universals and cultural variation in turn-taking in conversation. *Proceedings of the National Academy of Sciences*, 106(26):10587–10592, 2009.
 - [120] Morris Swadesh. Lexico-statistic dating of prehistoric ethnic contacts. *Proceedings of the American Philosophical Society*, 96(4):452–463, 1952.
 - [121] Sali A. Tagliamonte and Derek Denis. Linguistic ruin? lol! instant messaging and teen language. *American Speech*, 83(1):3–34, 2008.
 - [122] Deborah Tannen. *Talking Voices: Repetition, Dialogue, and Imagery in Conversational Discourse*. Cambridge University Press, Cambridge, 1989.
 - [123] Michael Tomasello. *The Cultural Origins of Human Cognition*. Harvard University Press, Cambridge, MA, 1999.
 - [124] Michael Tomasello. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Harvard University Press, Cambridge, MA, 2003.
 - [125] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pages 5998–6008, 2017.
 - [126] John Walliss. *The Brahma Kumaris as a 'Reflexive Tradition': Responding to Late Modernity*. Ashgate, Aldershot, 2002.
 - [127] Uriel Weinreich, William Labov, and Marvin I. Herzog. Empirical foundations for a theory of language change. In Winfred P. Lehmann and Yakov Malkiel, editors, *Directions for Historical Linguistics*, pages 95–195. University of Texas Press, Austin, 1968.
 - [128] Frank Whaling. *Understanding the Brahma Kumaris*. Dunedin Academic Press, Edinburgh, 2012.
 - [129] Mareike B. Wieth and Rose T. Zacks. Time of day effects on problem solving: When the non-optimal is optimal. *Thinking & Reasoning*, 17(4): 387–401, 2011.