The Immediate Apprehension of Harmony: Generalizing Root, Dissonance, and Voice-Leading Through Tonalness Theory

Gregory J. Strohman

ORCID: 0009-0008-8077-3842

Abstract

Traditional harmonic analysis describes notation rather than sound. A C-major triad is three noteheads on the page and a dense, continuously evolving band of acoustic energy in performance, and where those two diverge—late chromaticism, microtonal practice, electronic sound-mass—inherited analytical tools have little purchase. This paper applies three measures derived from tonalness theory directly to recorded performance: τ, which renders the continuous landscape of candidate fundamentals a sonority projects and so generalizes the notion of a chord root; η, which measures physiological dissonance as a function of register, timbre, spacing, and tuning rather than interval class; and $\mathbf{\Omega}$, the tonalness progression function, which measures the smoothness of harmonic motion as the transport of acoustic mass between moments. Applied without modification to Strauss, Mozart, Wagner, Liszt, Collier, and Varèse, the measures show that complexity on the page is not dissonance in the ear, that the emancipation of dissonance appears as a rising numerical floor from Mozart to Wagner, and that in performance articulation and orchestration govern harmonic motion as much as voice-leading does. Where notation vanishes altogether, the measures still return values, and still find tension and release. The framework addresses immediate apprehension and complements rather than displaces the synoptic concerns of traditional analysis.

Part I. The Framework

The Rameau problem and the notation illusion

When Jean-Philippe Rameau published *Traité de l'harmonie réduite à ses principes naturels* in 1722,[1] he sought to reduce harmony to its natural principles. His theory, and three centuries of theory after it, was built for the idealized discrete harmonic vocabulary of Common Practice tonal music. It served that vocabulary brilliantly. Friction began to arise as that vocabulary expanded in the later nineteenth and twentieth centuries and these theories were stretched to explain musical languages they were never designed to analyze.

The deeper issue is that traditional analysis describes the notation, not the performance. A C-major triad on the page is a static block of three pitch classes. When sounded in a hall, shaped by the vibrato of a string section, the envelopes of brass, and the inharmonicity of real instruments, it becomes a dynamic and ever-changing acoustic event. Traditional music theories analyze the idealized map and quietly assume it stands in for the acoustic territory. For a J.S. Bach chorale the gap is small enough to ignore. For Richard Wagner's saturated chromaticism, Jacob Collier's microtonal voicings, or Edgard Varèse's sound-masses, the map and the territory part company entirely, and the inherited tools have little to say.

A framework that aims at the sound rather than its symbol must be able to take in any frequency distribution at all and return the same kind of measurement for each. This applies whether the distribution is discrete or continuous, harmonic or inharmonic, consonant or pure noise.

---

[1] Jean-Philippe Rameau, *Traité de l'harmonie réduite à ses principes naturels* (Paris: Ballard, 1722).

Immediate apprehension, and what "tonalness" means

One boundary has to be drawn first. Traditional theory excels at synoptic comprehension. It handles the large-scale formal, cultural, and historical understanding of a complete work, and nothing here competes with that. By contrast, this framework governs only immediate apprehension. It tracks the moment-to-moment, bottom-up processing of the acoustic signal as the ear first meets it. The two are complementary lenses on the same music.

Fourier analysis can decompose any signal into its frequencies, but a raw spectrum is not what we hear, as it ignores the cognition the ear brings to the sound. Tonalness names that cognition. Tonalness is the propensity of the human auditory system to project virtual pitches, which are phantom fundamentals it hears as the common root of a given distribution of spectral energy, onto whatever it is given. Tonalness theory folds three well-established psychoacoustic facts that a plain spectrum misses, namely the critical band, the ear's frequency response, and virtual-pitch perception, back into Fourier analysis. The result is a single time-varying function from which the three analytical tools below are read. The equations themselves are not reproduced here. This paper is concerned with applying the tools to music rather than deriving them. The continuous-case mathematical definitions are given in Appendix A, and their full derivation in the companion paper.[2]

Precedents

The idea that harmony might be explained from the properties of hearing rather than from the conventions of notation is not new, and each of the three tools below has identifiable

[2] Greg Strohman, "Tonalness: A Signal-Domain Transform for Virtual Pitch, with Implications for Musical Harmony" (unpublished manuscript, October 2026), PDF.

ancestors. The psychoacoustic line runs from Helmholtz, who located dissonance in the interference of partials falling within a shared critical band,[3] through Plomp and Levelt, who gave that account its experimental basis,[4] to Terhardt, whose theory of virtual pitch supplied the mechanism by which the ear infers fundamentals not physically present.[5] Parncutt brought this apparatus to musical harmony directly, proposing that a chord projects not a single root but a profile of candidate roots of varying salience.[6] Sethares showed that because dissonance depends on the spectrum of the sounding tones, consonance is a joint property of tuning and timbre rather than of interval class alone.[7] On the theoretical side, Lerdahl's tonal pitch space and Tymoczko's voice-leading geometry both replace categorical rules with continuous quantitative distances,[8] and Cohn's account of parsimonious voice-leading pursues a related economy.[9]

What distinguishes the present framework from all of these is the domain of its input. Each of the models named takes as its object a chord—a list of pitches or of partials, supplied by the analyst from a score. The tools below take a sound pressure signal, and are defined for any

---

[3] Hermann von Helmholtz, *On the Sensations of Tone as a Physiological Basis for the Theory of Music*, trans. Alexander J. Ellis (New York: Dover Publications, 1954), 172–97.

[4] Reinier Plomp and Willem J. M. Levelt, "Tonal Consonance and Critical Bandwidth," *Journal of the Acoustical Society of America* 38, no. 4 (1965): 548–60.

[5] Ernst Terhardt, "Calculating Virtual Pitch," *Hearing Research* 1, no. 2 (1979): 155–82.

[6] Richard Parncutt, "Revision of Terhardt's Psychoacoustical Model of the Root(s) of a Musical Chord," *Music Perception* 6, no. 1 (1988): 65–93; and *Harmony: A Psychoacoustical Approach* (Berlin: Springer-Verlag, 1989).

[7] William A. Sethares, "Local Consonance and the Relationship between Timbre and Scale," *Journal of the Acoustical Society of America* 94, no. 3 (1993): 1218–28; and William A. Sethares, *Tuning, Timbre, Spectrum, Scale*, 2nd ed. (London: Springer, 2005).

[8] Fred Lerdahl, *Tonal Pitch Space* (New York: Oxford University Press, 2001); and Dmitri Tymoczko, *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice* (New York: Oxford University Press, 2011).

[9] Richard Cohn, "Neo-Riemannian Operations, Parsimonious Trichords, and Their *Tonnetz* Representations," *Journal of Music Theory* 41, no. 1 (1997): 1–66.

signal whatever. That is why they can be turned on a recording, on a microtonal performance that no notation captures, and on music that has no score at all.

Three tools of tonalness analysis

τ: The Generalization of the Root

A chord, to the ear, does not have a single, static root. It generates a continuous landscape of candidate fundamentals called virtual pitches, each projected with varying intensity. The Greek letter τ represents the function that renders that landscape. Across the pitch range, it calculates exactly where the ear is trying to organize the sound into a fundamental, and how aggressively it is trying. It is the direct, continuous successor to Rameau's fundamental bass, read off the physical acoustic wave rather than assigned from musical notation.

Two things are read directly from τ. The first is harmonicity. Where the virtual pitches concentrate on a single fundamental, the sonority fuses and sounds like one thing; where they scatter across several competing fundamentals, it disperses. The second is roughness. Where partials and virtual pitches collide within a critical band, τ flutters rather than holding steady, and the sonority beats. These two readings vary separately. A sonority can be tightly fused yet rough, or dispersed yet smooth. The harmonicity reading is why two chords with identical interval content can sound utterly different. A major triad concentrates its virtual pitches on a single fundamental and fuses, whereas a minor triad splits them across two and does not. The roughness reading is the physiological tension that η will later measure. Together they make τ the qualitative tool of the framework. τ is the fingerprint you point to in order to say why a particular voicing of a chord sounds the way it does.

η: The Generalization of Dissonance

In traditional music theory, "dissonance" is a discrete cultural label pinned to interval classes. A tritone is dissonant, a perfect fifth is consonant, etc. In tonalness theory, we set that label aside—not because the cultural framing is unreal,[10] but because η measures something else—and use η to calculate the continuous beating and phase-drift between virtual pitches colliding within the critical bands of the basilar membrane, weighed against the harmonic mass of the sonority. η is an objective measurement of physiological tension, capable of registering the dissonance in a poorly voiced major triad or the comparative consonance of a microtonally tuned jazz voicing.

Two properties make η valuable for analysis. First, it depends on register, timbre, and tuning, not just interval class, so it can distinguish between various voicings, orchestrations, textures, and arpeggiations of a particular chord. Second, because it is normalized, it is comparable across chords of any size: a triad against a seven-note extended jazz voicing against a spectral sound mass. η is the tool we will use for rating harmonic tension.

Ω: The Generalization of Voice-Leading

The species-counterpoint tradition of Common Practice tonal music governs the motion between chords with geometric prohibitions: avoid parallel fifths, resolve the leading tone by step, etc. Ω reframes voice-leading not as a set of rules, but as a measurement of virtual pitch movement. Moving from one sonority to the next requires the physical transport of acoustic mass across the frequency spectrum. Ω, the tonalness progression function (TPF), calculates the total

[10] Peter M. C. Harrison and Marcus T. Pearce, "Simultaneous Consonance in Music Perception and Composition," *Psychological Review* 127, no. 2 (2020): 216–44.

migration of virtual pitches in a sound between two moments in time. A note on naming: the function is named for what it tracks, progression, but it is used to measure smoothness, and the relationship is inverse. A low value means the landscape has barely moved and the progression is smooth. A high value means it has lurched. Headings and figures in what follows label Ω by what it measures, "Smoothness Ω," with low values indicating greater smoothness. The TPF provides a parsimonious metric for smoothness that works equally well for a cadence in a classical sonata or an electronic sound-mass sweeping across the frequency spectrum.

Two properties make Ω valuable for analysis. First, like η it answers to register, timbre, and tuning rather than to notated intervals alone, so it can separate a progression that looks smooth on the page from one that actually lurches in performance. Second, it is the only one of the three tools that measures motion rather than a static snapshot. This makes harmonic progression itself, and not merely the chords at its endpoints, available to analysis. A small Ω marks a small step between harmonic states; a large Ω marks a leap. Ω is the tool for measuring motion.

Advantages of immediate apprehension analysis

Because these three tools measure the nervous system's immediate response to any acoustic signal, they need no separate rulebook for each style or century. The same tools that analyze a harmonic cadence in Mozart also analyze a Varèse sound-mass or Wagner's most evasive chromatic progressions. What follows applies this one analytical kit across functional tonality, late-Romantic chromaticism, microtonal jazz, and pure electronic sound, to show that a single set of acoustic principles is at work beneath all of them.

Part II. From Symbol to Signal: The Tools in Demonstration

Before applying the three tools to repertoire, it is worth seeing what they measure, made concrete in a single sonority and a single perceptual fact.

The notation illusion, shown

Consider again a C-major triad. Notated, it is three noteheads—a static block of three pitch classes, identical on every page it appears (Figure 1). Sounded as the fortissimo "Sunrise" climax of Strauss's *Also sprach Zarathustra*—full orchestra crowned by the pedal and full stops of the organ—it becomes the spectrogram of Figure 2: not three lines, but a dense, seething band of energy—fundamentals and overtones, the difference tones of massed brass, the frequency spread of string vibrato, the attack transients of timpani and organ pipe speech. One harmony, two representations; between them lies everything traditional analysis discards. For a restrained four-voice Bach chorale the notation is a fair proxy for the sound; for Wagner, Collier, and Varèse it is a progressively poorer one, and for Varèse it is absent altogether—which is why the tools that follow read the terrain, not the map.

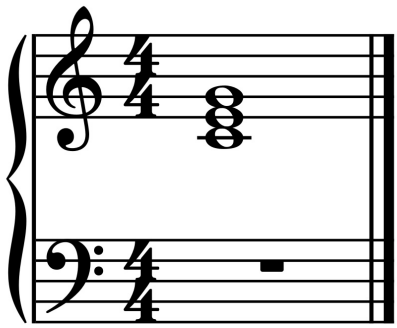


FIGURE 1. A STANDARD NOTATED C-MAJOR TRIAD. A STATIC BLOCK OF THREE PITCH CLASSES.

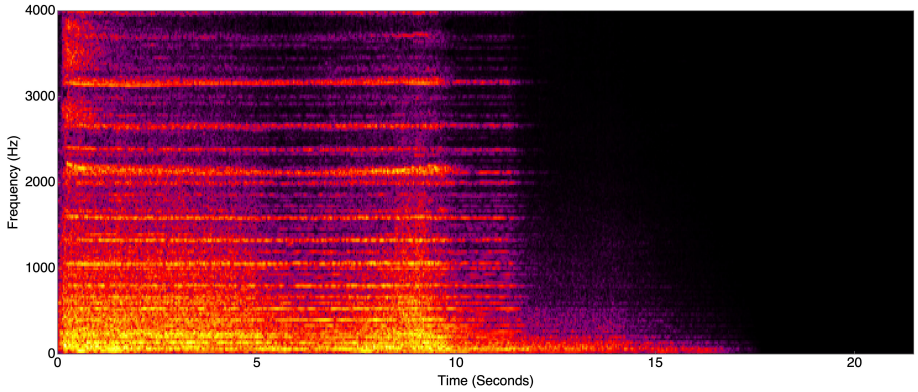


FIGURE 2. SPECTROGRAM OF THE OPENING C-MAJOR CLIMAX OF "SUNRISE," FROM RICHARD STRAUSS'S *ALSO SPRACH ZARATHUSTRA*.[11] THE SAME TRIAD AS A CONTINUOUS, DENSE ACOUSTIC EVENT.

Spectral pitch and virtual pitch

To read the terrain we need one distinction that ordinary spectral analysis does not draw: between the pitches that are physically present in a signal and the pitches the ear infers from them.

A spectral pitch is a frequency actually present in the acoustic signal. It is a component that a microphone registers and a Fourier transform recovers. A virtual pitch is not present in the signal at all. It is a cognitive projection: the auditory system's attempt to deduce a fundamental

[11] Berliner Philharmoniker, "Strauss: Also sprach Zarathustra / Dudamel · Berliner Philharmoniker," YouTube video, 1:55, October 26, 2012, https://www.youtube.com/watch?v=Szdziw4tI9o.

from the pattern of frequencies it receives, whether or not that fundamental is physically sounding.

The clearest demonstration is the missing fundamental. Present the ear with harmonics 3 through 8 of a fundamental at C3 (MIDI 48)—a cluster of upper partials with no energy at the fundamental itself—and it still hears a pitch at C3. Physically, the spectrum contains only those high partials, sitting at their true harmonic frequencies rather than on any tempered pitch (Figure 3); yet the tonalness function projects a strong virtual pitch at the absent fundamental, MIDI 48, together with its octaves (Figure 4). The fundamental is absent from the signal but present in perception. The ear is not measuring the spectrum; it is interpreting it, searching for the harmonic series that best explains the partials it receives and reporting that series' fundamental as a pitch.

The same mechanism, applied to a chord, is what makes harmony cohere. The three notes of a major triad approximate the 4th, 5th, and 6th harmonics of a fundamental two octaves below the root. The auditory system, finding that they fit a single harmonic template, projects a strong virtual pitch there—subordinate in strength to the chord tones themselves, but the only fundamental the ear is offered—and the triad partially fuses into one sonorous object (Figures 5–6). Change the chord to a minor triad and the same notes no longer fit one template cleanly. The ear projects two weaker, competing fundamentals, and the fusion loosens.

These two facts—that the ear infers fundamentals the signal omits, and that it fuses or fails to fuse a chord accordingly—are what τ quantifies. With the distinction in hand, we can turn τ on a set of real sonorities.

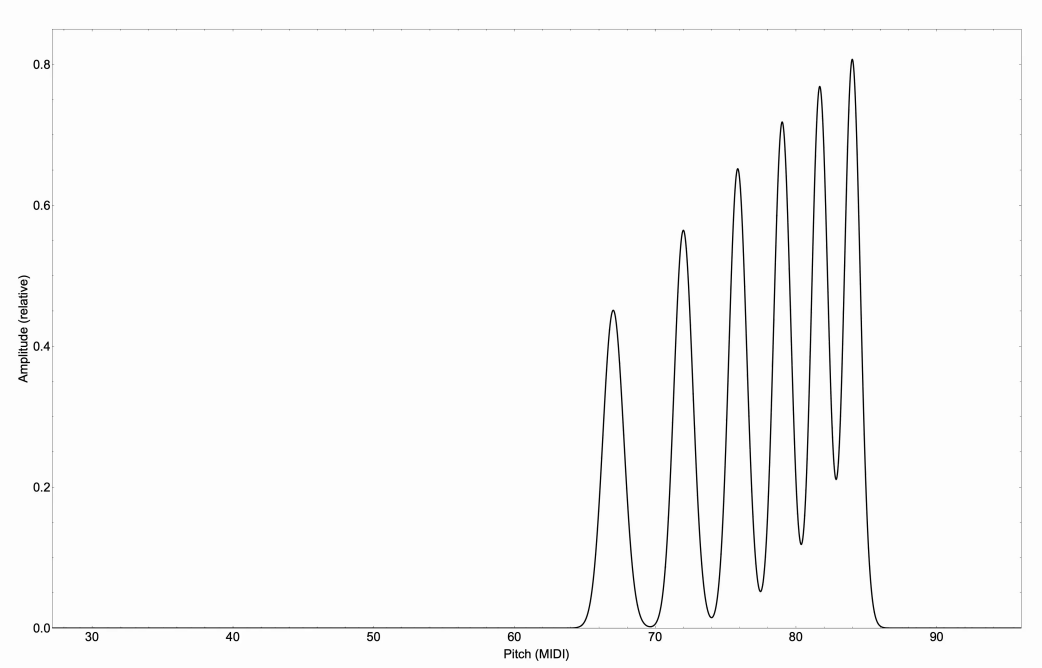


FIGURE 3. THE MISSING FUNDAMENTAL, SPECTRAL VIEW: HARMONICS 3–8 OF C3 (MIDI 48), WITH NO ENERGY AT THE FUNDAMENTAL.

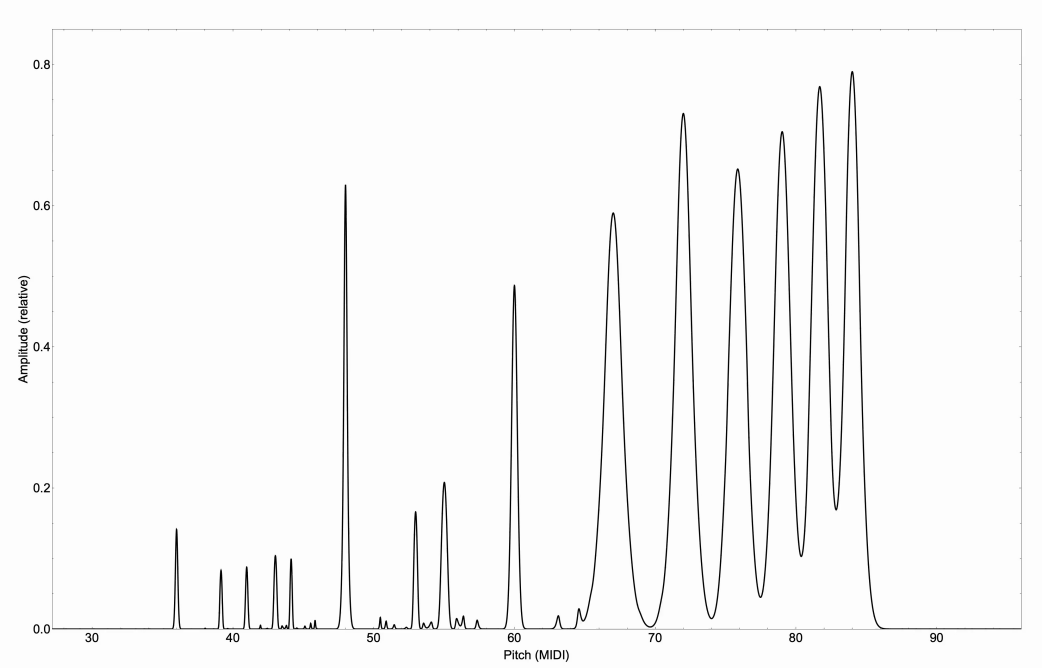


FIGURE 4. THE MISSING FUNDAMENTAL, TONALNESS VIEW: τ PROJECTS A STRONG VIRTUAL PITCH AT THE ABSENT FUNDAMENTAL (MIDI 48) AND ITS OCTAVES.

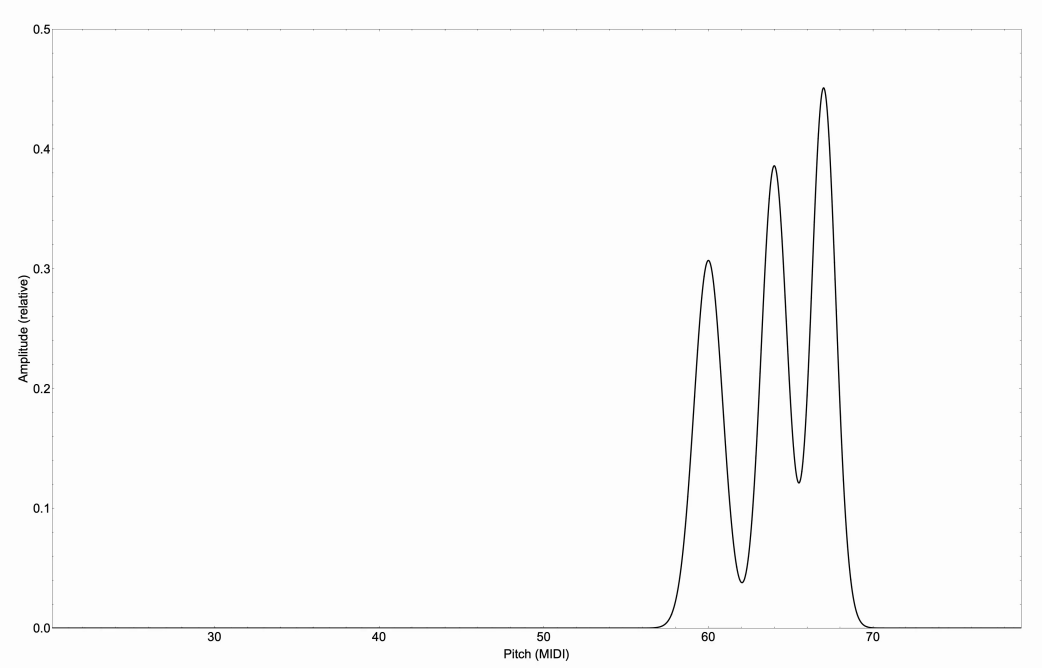


FIGURE 5. C MAJOR TRIAD, SPECTRAL VIEW: THE THREE CHORD TONES ONLY.

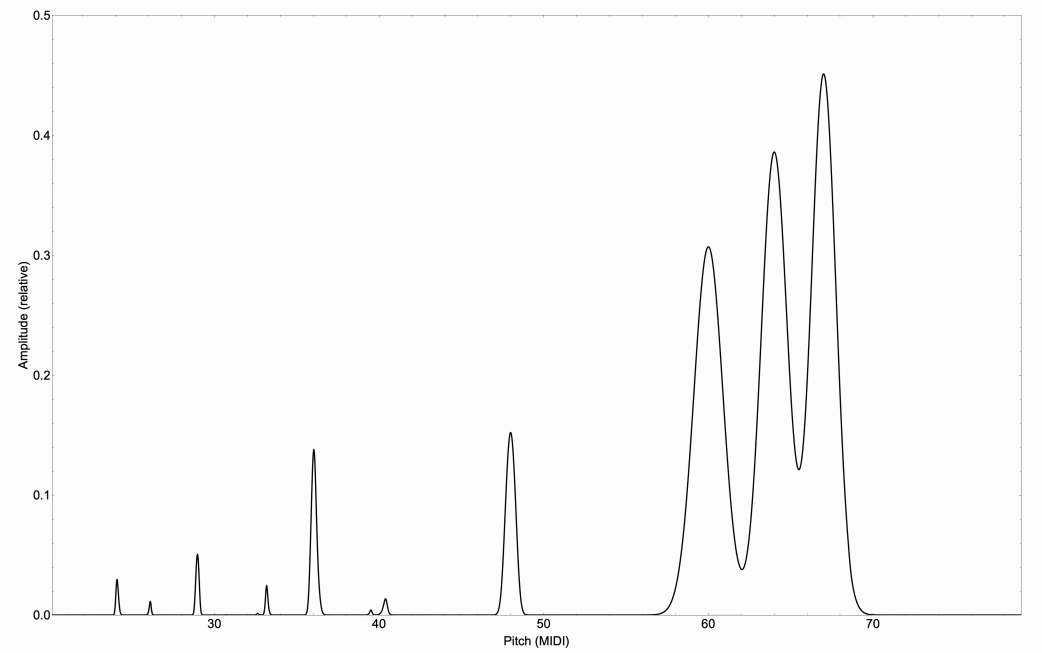


FIGURE 6. C MAJOR TRIAD, TONALNESS VIEW: THE VIRTUAL PITCHES CONVERGE ON A SINGLE LOW FUNDAMENTAL AND ITS OCTAVE—THE TRIAD FUSES.

Part III. The Tonalness Landscape: Analyzing Static Sonorities

A chord does not have a single root. Presented with a distribution of spectral energy, the auditory system projects a continuous landscape of candidate fundamentals, each with some strength; the tonalness function τ renders that landscape across the pitch range. Two things are read from it. Where the projected virtual pitches concentrate on one fundamental, the sonority is harmonic and fuses; where they scatter across several competing fundamentals, it is dispersed. Separately, where components fall close enough to collide within a critical band, τ flutters and the sonority is rough—and it is this roughness, normalized against harmonicity, that the dissonance measure η reports as a single number. Harmonicity and roughness are independent readings, and the eight sonorities below were chosen to pull them apart.

Chord	Pitches (MIDI)	Description	Dissonance η
1	48, 55, 64, 72	C major triad, open spacing	0.0007
8	48, 62, 67, 71	G major triad, over C bass	0.0019
4	29, 36, 43, 50, 57, 64, 71, 78, 85, 92	ten-note stack of perfect fifths	0.0097
6	36, 43, 52, 57, 62, 66, 71	mixed quintal–tertian spread	0.0110
5	36, 47, 58, 69, 80, 91	stack of major sevenths	0.0115
2	41, 48, 57, 63, 68	F dominant seventh, sharp ninth	0.0117
7	43, 48, 55, 61, 66, 73	tritone-laden stack	0.0253
3	60–71 (all twelve)	full chromatic cluster, C4–B4	0.2650

TABLE 1. DISSONANCE η OF EIGHT SONORITIES.

The fused pole

Chords 1 and 8 (Figures 7 and 14) return the lowest dissonance values of the set, and their landscapes show why. Chord 1—a C major triad voiced C–G–E–C across two octaves—concentrates its virtual pitches onto a single low fundamental near C, with clean, well-separated peaks. The ear is offered one unambiguous fundamental and takes it; the sonority fuses. Chord 8, a G major triad over a C bass in similarly open spacing, behaves similarly, though with a weaker low fundamental and with nearly three times the dissonance. Neither chord places two components inside a critical band, so roughness is negligible and η sits at the floor of the set (0.0007 and 0.0019), a clear step below every other sonority measured here. These are the reference sonorities: high harmonicity, low roughness.

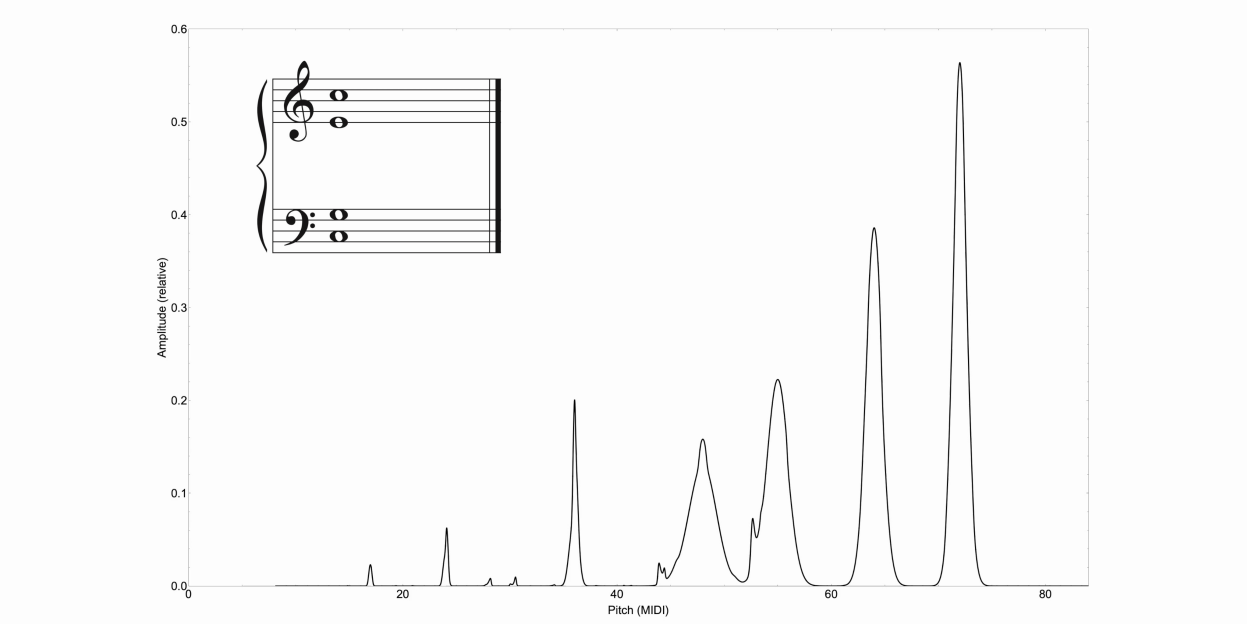


FIGURE 7. CHORD 1 TONALNESS τ: C MAJOR TRIAD, OPEN SPACING.

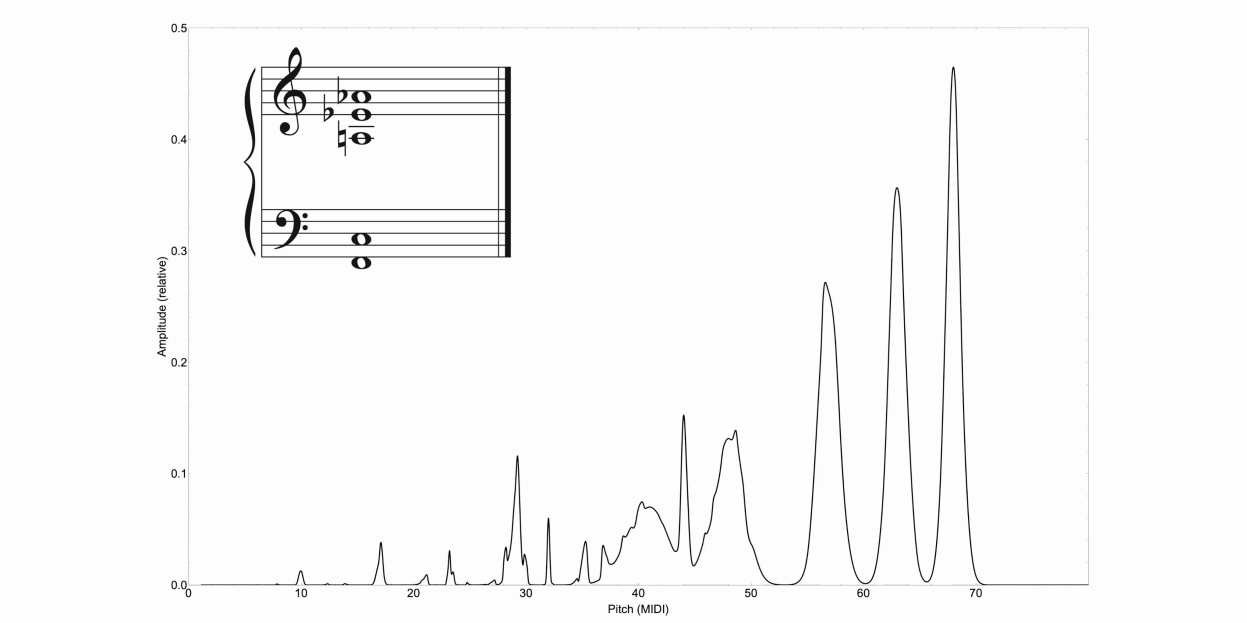


FIGURE 8. CHORD 2 TONALNESS τ: F DOMINANT SEVENTH, SHARP NINTH.

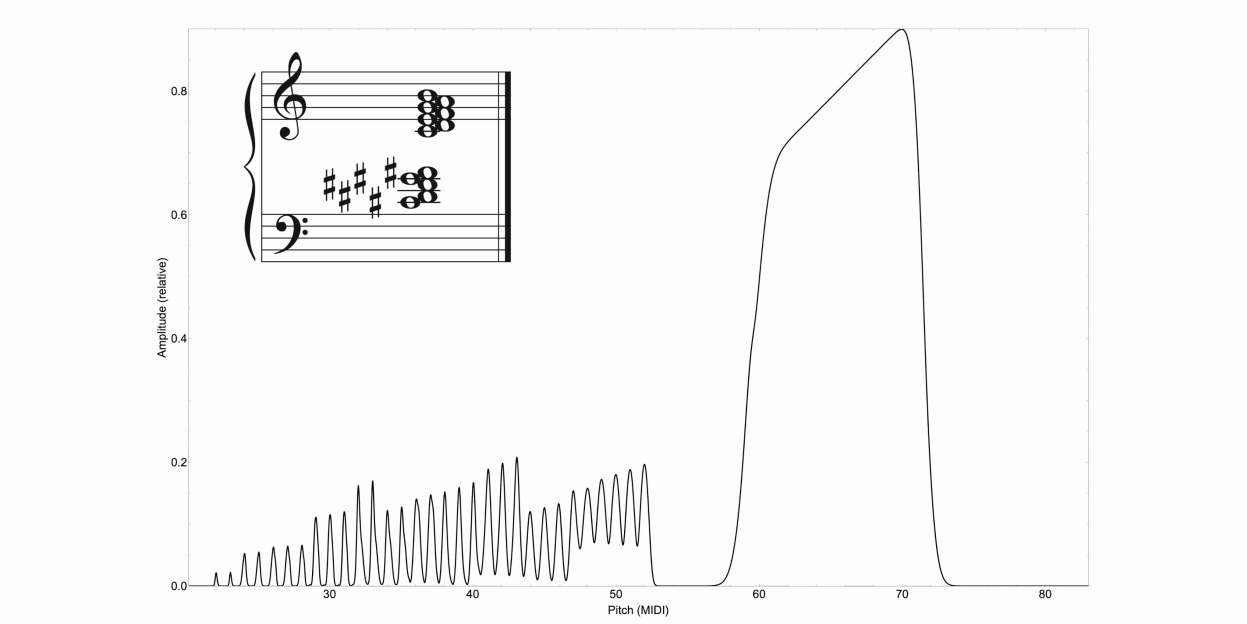


FIGURE 9. CHORD 3 TONALNESS τ: FULL CHROMATIC CLUSTER, C4–B4.

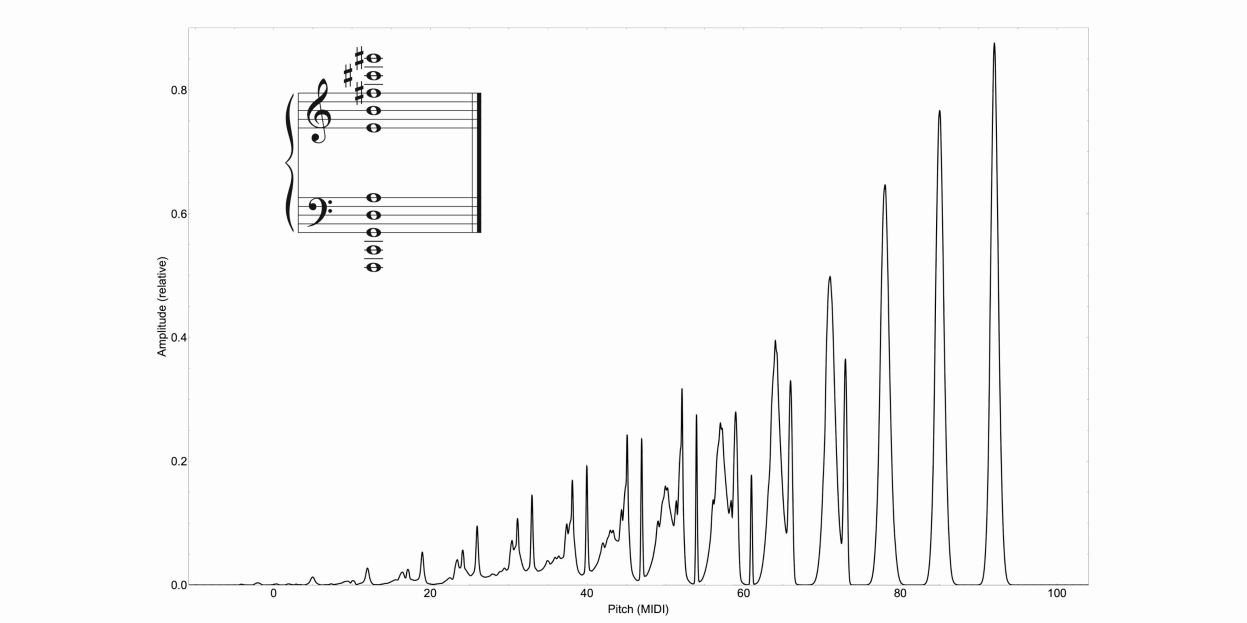


FIGURE 10. CHORD 4 TONALNESS τ: TEN-NOTE STACK OF PERFECT FIFTHS.

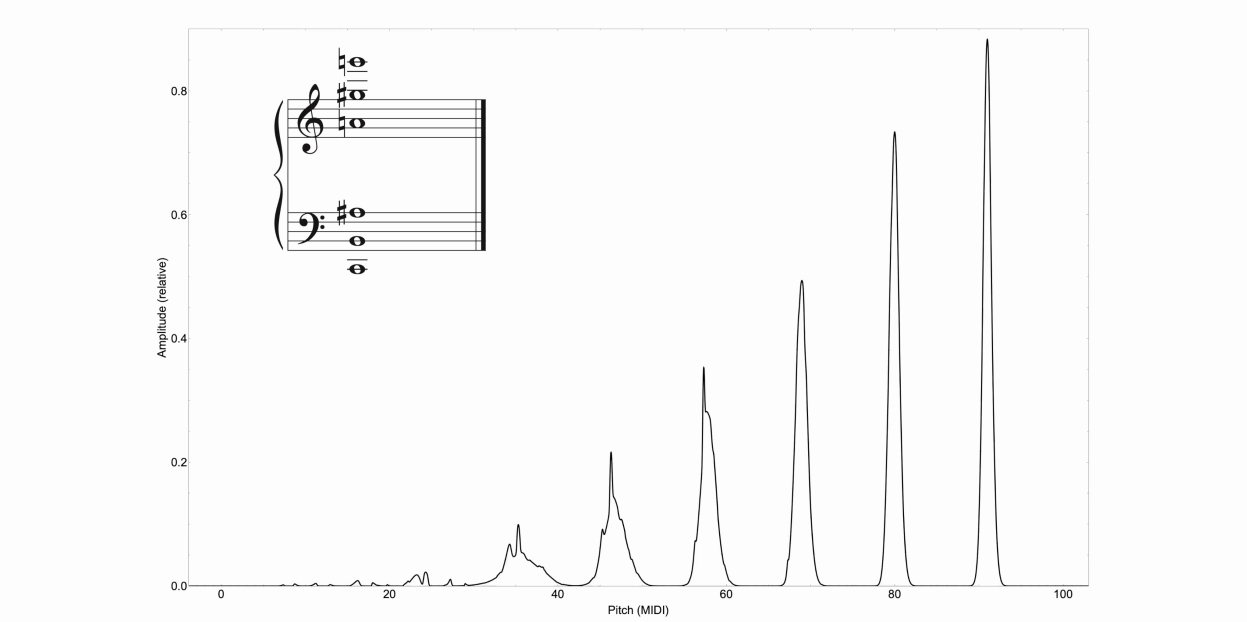


FIGURE 11. CHORD 5 TONALNESS τ: STACK OF MAJOR SEVENTHS.

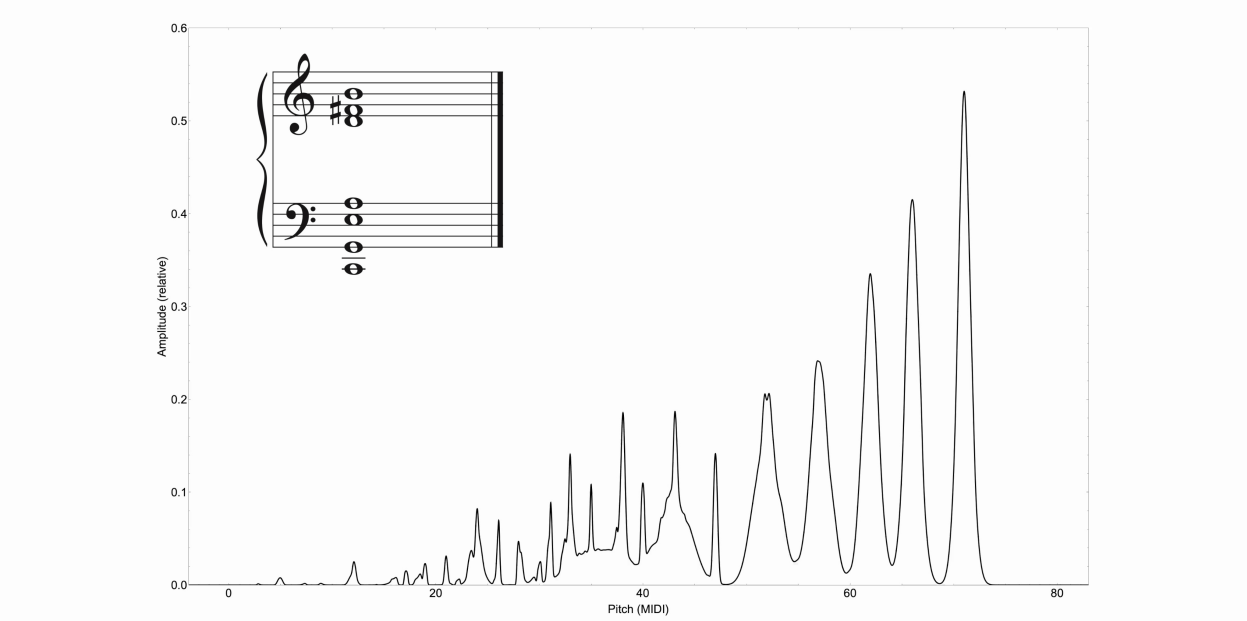


FIGURE 12. CHORD 6 TONALNESS τ: MIXED QUINTAL–TERTIAN SPREAD.

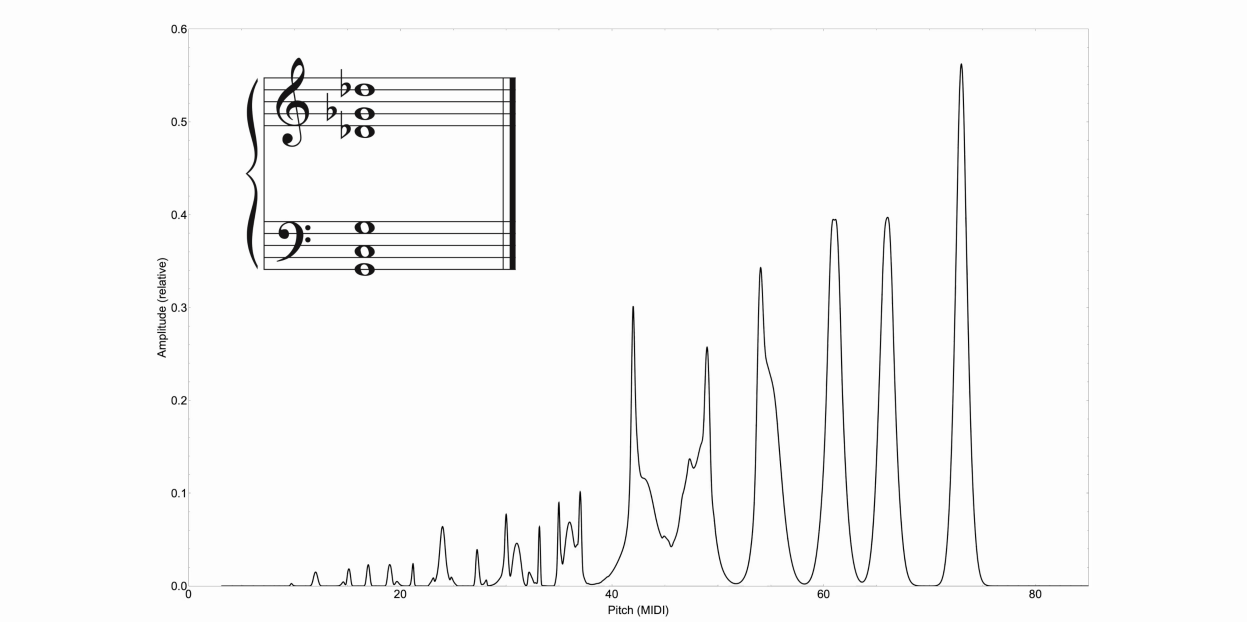


FIGURE 13. CHORD 7 TONALNESS τ: TRITONE-LADEN STACK.

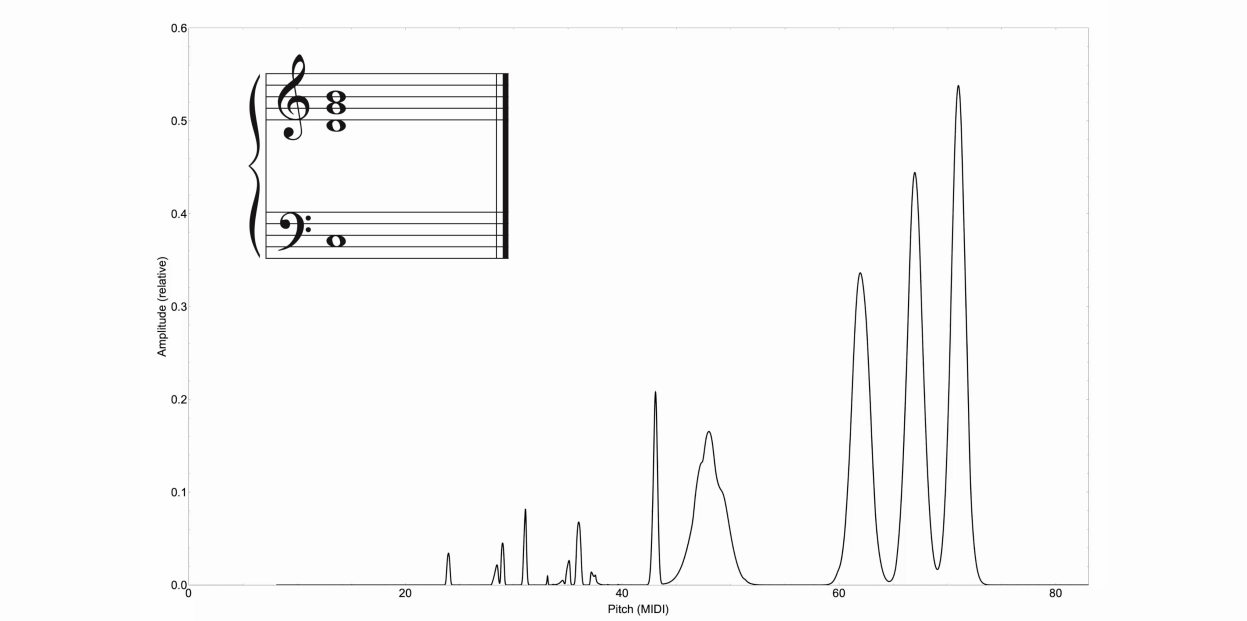


FIGURE 14. CHORD 8 TONALNESS τ: G MAJOR TRIAD, OVER C BASS.

Complexity is not dissonance

The heart of the comparison is the pair of chords that a note count would rank most alike and the ear ranks least alike. Chord 3 is a twelve-note chromatic cluster spanning a single octave, C4 to B4. Chord 4 is a ten-note stack of perfect fifths spanning five and a quarter octaves, F1 to G♯6. Both are dense, many-note sonorities that a traditional analyst would flag as extreme. Their dissonance values differ by a factor of roughly twenty-seven.

Chord 3 packs twelve pitches into the interval of an octave, so every adjacent pair—and most non-adjacent pairs—falls within a critical band. The result is maximal beating: the τ landscape (Figure 9) does not resolve into peaks in the cluster's own register at all, but collapses into a broad, fluttering smear — beneath which the ear scatters a dense comb of weak, competing fundamentals, none of them winning. η reaches 0.2650, an order of magnitude above every other

chord in the set. This is physiological dissonance in its purest form, and it is a property of the spacing, not the note count.

Chord 4 proves the point from the opposite direction. It has nearly as many notes, but they are spread in pure fifths so widely that no two components ever fall inside a critical band. Its landscape (Figure 10) is a clean picket-fence of separated peaks climbing across the register, and its η—0.0097—is a twenty-seventh of the cluster's, for a sonority with nearly as many notes, though it stands well clear of the two triads. Yet chord 4 does not fuse: stacked fifths spiral through ten different pitch classes and offer the ear no single fundamental, so the landscape is thoroughly dispersed. Chord 4 is therefore dispersed and smooth at once—many competing fundamentals, almost no roughness. It is the clearest demonstration available that the harmonicity a listener reads from τ and the roughness that η measures are genuinely separate axes: a sonority can fail to fuse and still be physically consonant.

The middle ground

The four remaining chords occupy the interval between these poles, and what they show is less a graded ranking than a near-tie broken by a single outlier. Chord 6 (0.0110), a mixed quintal–tertian spread; chord 5 (0.0115), a stack of major sevenths; and chord 2 (0.0117), an F dominant seventh with a sharp ninth, fall within about six percent of one another—a spread too small to rank, and a telling one, since the first is a plain quintal–tertian spread and the last the most altered chord in the set. Only chord 7 (0.0253), the tritone-laden stack, separates from them, at better than twice their value, registering the highest roughness short of the chromatic cluster. The ordering that survives tracks how much of a chord's content the voicing forces into critical-

band proximity, not how "advanced" it looks in notation—and where three sonorities carrying three quite different labels converge on the same number, the label predicts nothing at all.

From landscape to metric

These eight readings establish what τ can and cannot do as an analytical tool. It renders the harmonic character of a sonority with a precision no chord symbol approaches, distinguishing fusion from dispersion and both from roughness. But τ is a continuous, high-dimensional function of frequency, and it is also highly dynamic: the landscape of a real sonority shifts continuously as the sound evolves. Mapping it in full across the duration of a piece is not feasible on the static page. For temporal analysis we therefore rely on the two scalar quantities derived from τ—dissonance (η), which collapses the roughness of the landscape at an instant into a single number, and the tonalness progression function (Ω), which measures how far the landscape travels between two instants. The remainder of this section demonstrates both on a single passage—the close of Strauss's "Sunrise"—before the next section turns to the wider repertoire.

The Sunrise in motion: η and Ω on the final progression

The eight chords above were abstract test cases, chosen to isolate what τ measures. We now watch the metrics operate on real music in time. The passage is the close of Strauss's "Sunrise" from *Also sprach Zarathustra*—the cadential gesture that resolves onto the fortissimo C-major arrival—which puts η and Ω to work on a familiar sonority. The notation is deliberately plain: three harmonic events resolving to a sustained C major (Figure 15). The map could hardly

be tidier; the terrain is considerably denser. For each metric the ideal map trace is a MIDI rendering of the notated harmony and the real terrain trace is the acoustic recording.

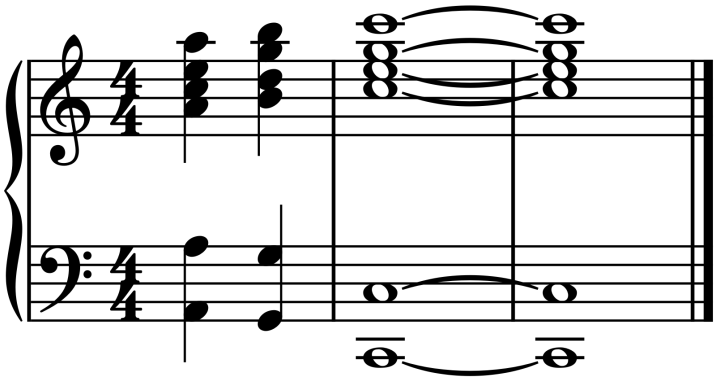


FIGURE 15. NOTATION MAP OF "SUNRISE" FROM ALSO SPRACH ZARATHUSTRA, BY RICHARD STRAUSS.

Dissonance η

On the map, η is a clean descending staircase—three flat plateaus, one per event, stepping downward (approximately 0.0035, then 0.0030, then 0.0017) as the harmony resolves toward the pure triad (Figure 16). The terrain diverges on two independent counts. The first is magnitude: the performed η sits near 0.14 and spikes above 0.20—roughly eighty times the notated values (Figure 17). The interval arithmetic of the map cannot see the broadband attack energy, ensemble vibrato, instrumental inharmonicity, and hall noise that make up most of what the ear receives. The second is shape: where the map is flat, the terrain is volatile, with sharp spikes at each orchestral attack, a turbulent elevated band as the massed chord reaches full cry, sustained to the end of the excerpt. The resolution the map draws as a neat step-down, the terrain accomplishes as a noisy climb that does not come back down.

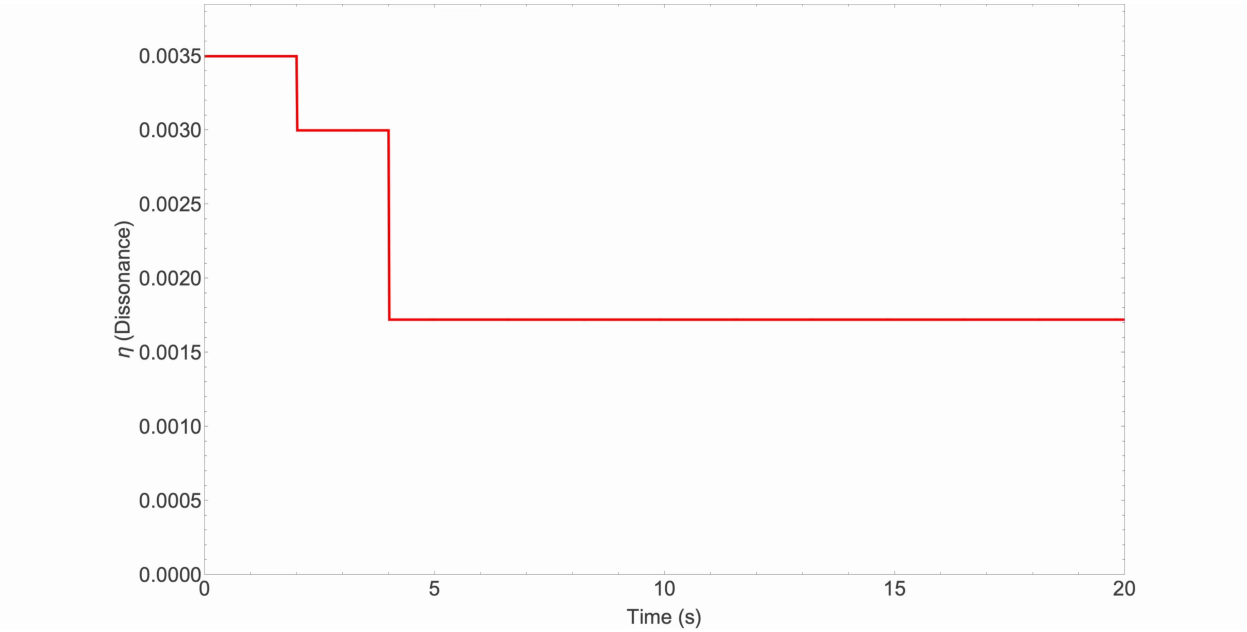


FIGURE 16. DISSONANCE η MAP OF "SUNRISE" FROM ALSO SPRACH ZARATHUSTRA, BY RICHARD STRAUSS.

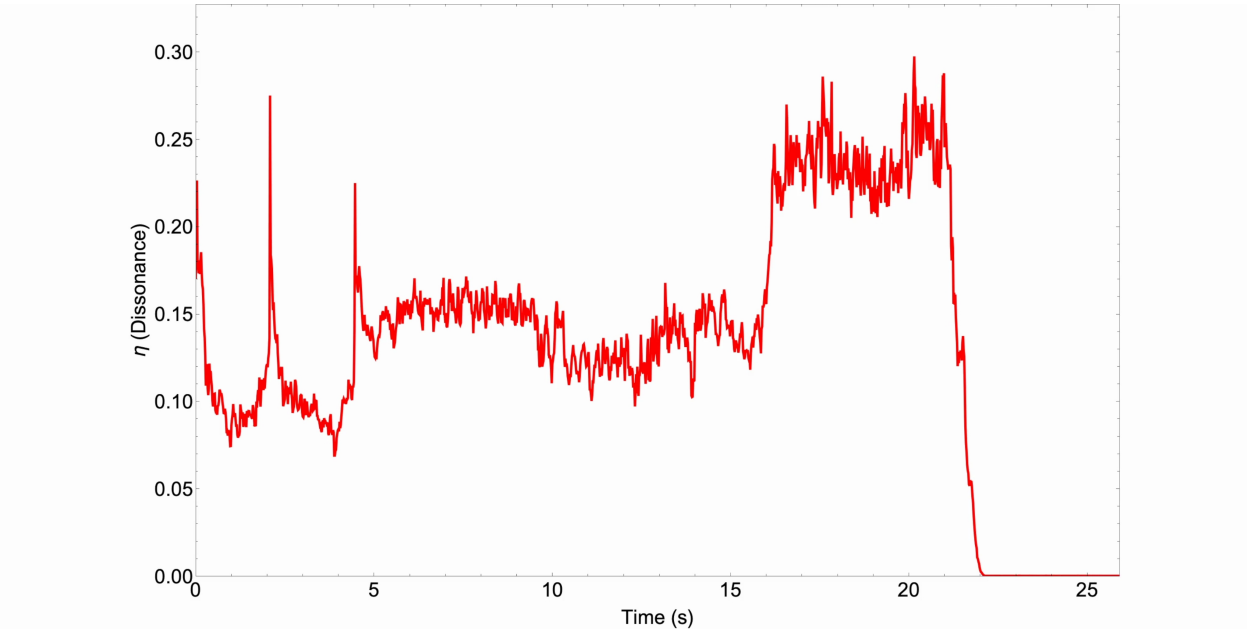


FIGURE 17. DISSONANCE η TERRAIN OF "SUNRISE" FROM ALSO SPRACH ZARATHUSTRA, BY RICHARD STRAUSS.

Smoothness Ω

Here the contrast changes character. On the map, Ω is silent except for two needle-thin spikes at the two chord changes (Figure 18): motion is punctual, and between the changes

nothing moves at all. In the terrain, Ω never stops—a continuous chatter of acoustic migration runs the full length of the passage, its largest spikes aligned to the actual attacks but embedded in constant, low-level motion (Figure 19). The trace runs a floor of about 0.03–0.07, and its two largest excursions — about 0.22 at 2.1 s and about 0.22 at 4.5 s — fall on the map's two notated chord changes. At 22 s the motion drops to a quieter band near 0.03 and stays there. The score's two harmonic events are real, and everything between them is the performance's own continuous micro-motion.

Together the four traces establish the method the next section carries into a variety of repertoire. The notation predicts a brief, tidy, resolving cadence. The performance delivers a continuously moving, noise-rich, and far more dissonant acoustic event. Its arrival registers not as a fall in η, which never comes, but as Ω settling into its quiet band after 22 s.

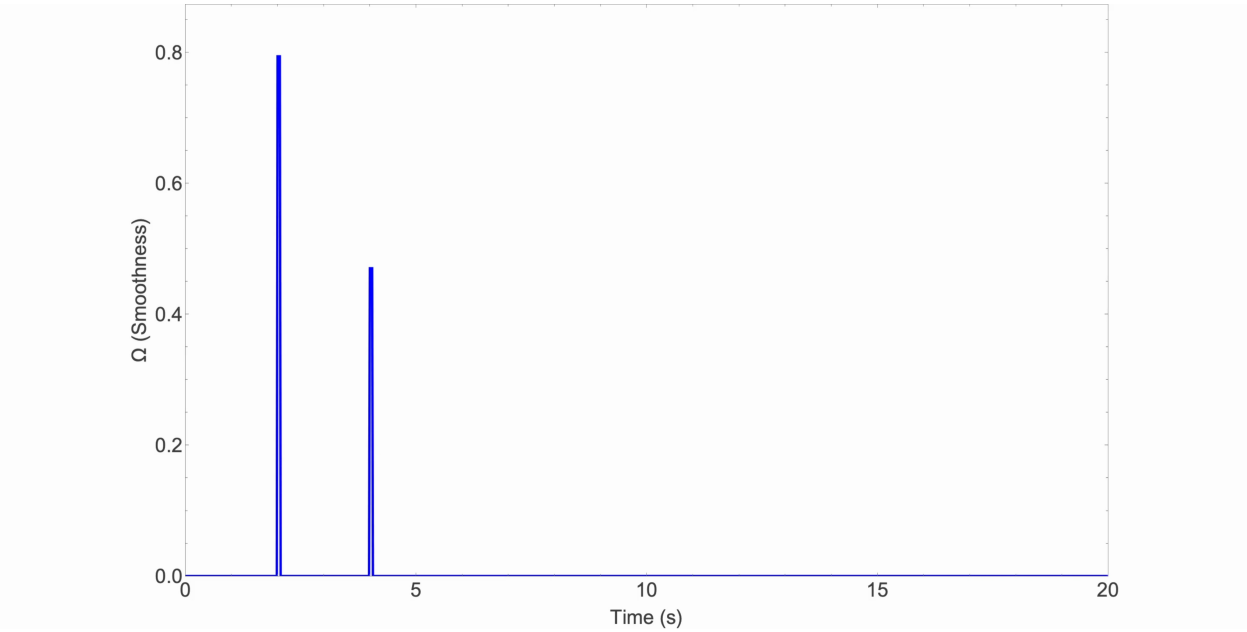


FIGURE 18. SMOOTHNESS Ω MAP OF "SUNRISE" FROM ALSO SPRACH ZARATHUSTRA, BY RICHARD STRAUSS.

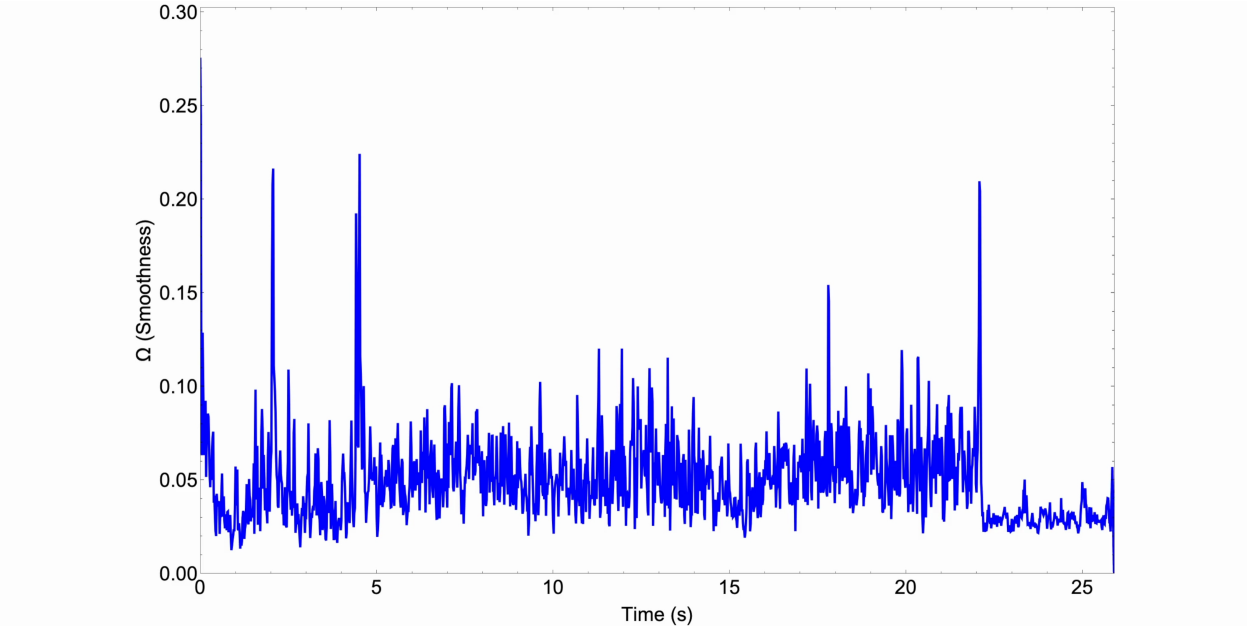


FIGURE 19. SMOOTHNESS Ω TERRAIN OF "SUNRISE" FROM ALSO SPRACH ZARATHUSTRA, BY RICHARD STRAUSS.

Part IV. Dynamic Analysis Across the Repertoire

Having demonstrated τ on isolated sonorities and η and Ω on Strauss's "Sunrise," we now apply the same procedure to the repertoire proper. The method is unchanged: for each excerpt, an ideal map trace (a MIDI rendering of the notated harmony) is set beside a real terrain trace (the recorded performance), and the distance between them isolates what the notation omits.

W. A. Mozart, Symphony No. 40: the baseline of functional tonality

Mozart is the control case. His harmonic language is precisely the one three centuries of theory were built to describe, so this is where the framework would merely echo the textbook if it were going to do so anywhere. Two excerpts are analyzed for both metrics (notation in Appendix B, Figures B1–B4).

The first excerpt is the movement's celebrated opening, measures 1–22.[12] A single bar of divided violas sets a murmuring, agitated accompaniment before the violins enter with the main theme—an urgent, restless melody spun from a repeated descending half-step sigh, cast as a tight antecedent–consequent period. The harmony stays close to home, circling the tonic and dominant of G minor and inflected by the theme's chromatic neighbors, until a sudden forte tutti near the end of the passage thickens the texture and begins pressing outward toward the relative major. Its tension is carried as much by rhythm, register, and orchestral weight as by any single chord.

The second excerpt is the exposition's lyrical second subject, measures 44–72, which arrives in the relative major, B-flat. In sharp contrast to the driving opening, the theme is graceful and songful—a chromatically descending line handed in dialogue between strings and woodwinds, sliding downward by half step over a stable B-flat major harmony. Here the chromaticism decorates a clear tonal center rather than unsettling it, and the passage eases the movement's tension before the closing material presses the exposition toward its cadence. Where the first theme is agitation, this is repose.

Dissonance η

In both excerpts the ideal map trace behaves exactly as the discrete model predicts: a staircase of flat plateaus, one per notated chord, stepping up and down as the harmony changes and toggling rapidly where the harmonic rhythm quickens (Figure 20). The real terrain trace[13]

---

[12] Wolfgang Amadeus Mozart, *Symphony No. 40 in G Minor, K. 550*, ed. H. C. Robbins Landon (Kassel: Bärenreiter, 1958).

[13] accentusmusic, "Lucerne Festival: Mozart – Symphony No. 40 (Berliner Philharmoniker, Simon Rattle)," YouTube video, 30:22, March 20, 2026, https://www.youtube.com/watch?v=Lk01TQnxsMk.

shares none of that flatness—it is volatile throughout, oscillating continuously within every plateau the map holds constant.

Two features of the contrast repay attention, and the first is visible only against the Strauss passage of the previous section. There, the notated triads were so nearly pure that the ideal map η all but vanished while the real terrain ran some eighty times higher. In Mozart the two traces occupy the same order of magnitude: the map runs from 0.00 to 0.11, the terrain from 0.02 to 0.26. The reason is instructive: Mozart's functional harmony already encodes genuine dissonance—dominant sevenths, suspensions, appoggiaturas—so the map itself predicts substantial roughness, and the terrain adds volatility rather than a wholesale jump in level. The size of the map–terrain gap is therefore not a fixed property of the method; it scales with how much real dissonance the notation already carries. What persists across both cases, regardless of magnitude, is the difference in shape: stepped versus volatile.

Second, the terrain does not preserve the map's ordering. Where the ideal map η cuts clean notches of near-perfect consonance at the simplest notated sonorities—dropping almost to zero in the second excerpt—the real terrain η dips at those points but does not hold the lows for long: its troughs are half-second dips rather than sustained plateaus. Register, doubling, ensemble vibrato, and attack noise inject roughness that the interval content cannot register, decoupling what the ear receives from what the notation ranks. This is the ground beneath the stronger claim that a low, densely doubled, heavily played major triad can read as rougher than a more openly voiced seventh chord elsewhere in the same passage: the plots make the decoupling plain.

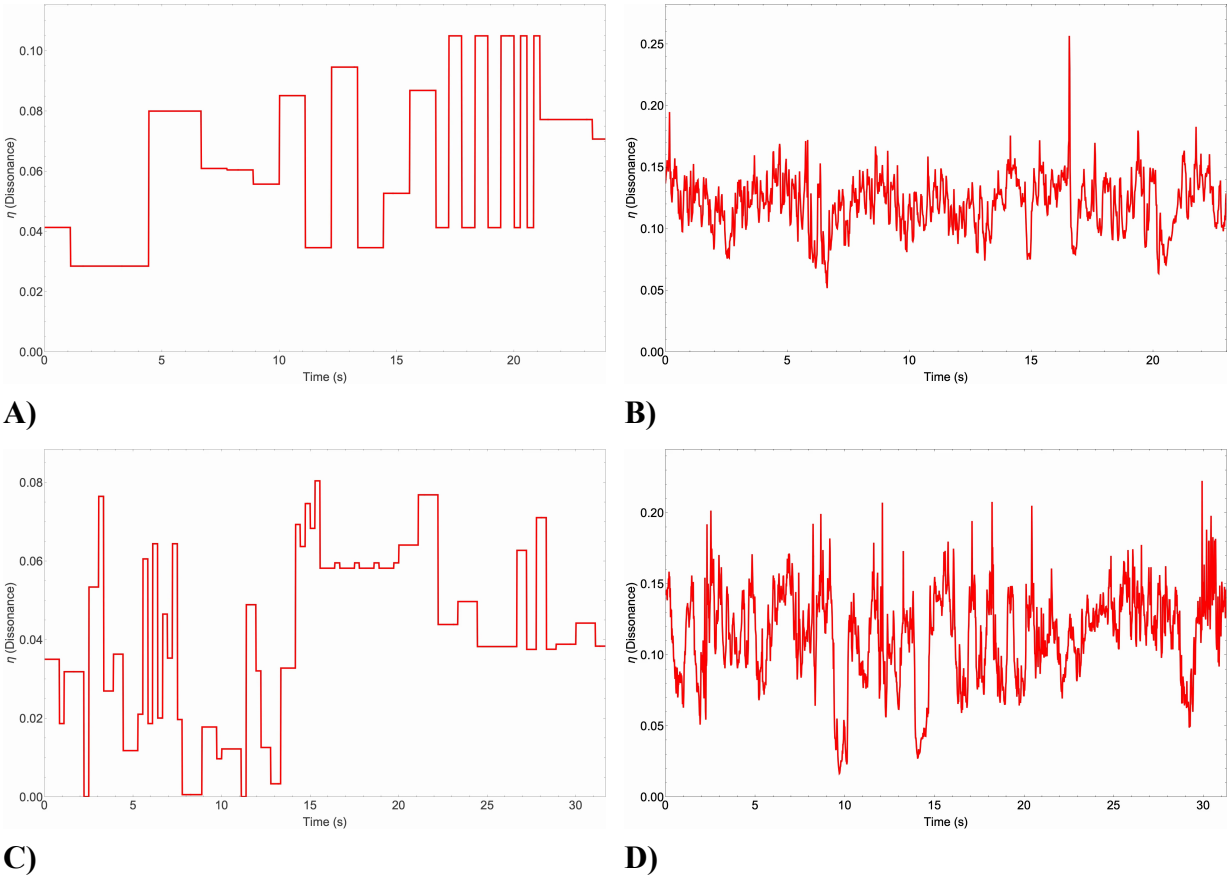


FIGURE 20. DISSONANCE η FOR TWO EXCERPTS FROM SYMPHONY NO. 40, BY W. A. MOZART: A) IDEAL MAP FOR EXCERPT 1 (SEE APPENDIX B, FIGURE B2), B) REAL TERRAIN FOR EXCERPT 1, C) IDEAL MAP FOR EXCERPT 2 (SEE APPENDIX B, FIGURE B4), D) REAL TERRAIN FOR EXCERPT 2.

Smoothness Ω

The ideal map Ω is, in both excerpts, a row of isolated spikes—one at each chord change, silence between — running together into a near-continuous band only where the notated harmonic rhythm quickens (Figure 21). The real terrain Ω never rests: a continuous chatter of acoustic migration runs the length of each excerpt, its largest spikes aligned to the notated changes but embedded in constant, low-level motion. Much of that motion answers to dynamics and orchestration rather than to harmony—swells, entrances, and articulations transport acoustic mass whether or not the chord changes underneath them. In the first excerpt a single instant at

16.6 s spikes simultaneously in both η and Ω, marking a strongly articulated structural event that the map registers, if at all, as merely one more chord change among many. The trace reaches about 0.35 there, against the 0.04–0.09 floor it holds elsewhere — the landscape moved that far within fifty milliseconds. Only an attack can do that.

The lesson of the baseline is that even in the repertoire the discrete model was built for, dissonance and smoothness are governed as much by register, orchestration, and dynamics as by the notated harmony. Everything the framework finds in the harder repertoire to follow is an amplification of this gap.

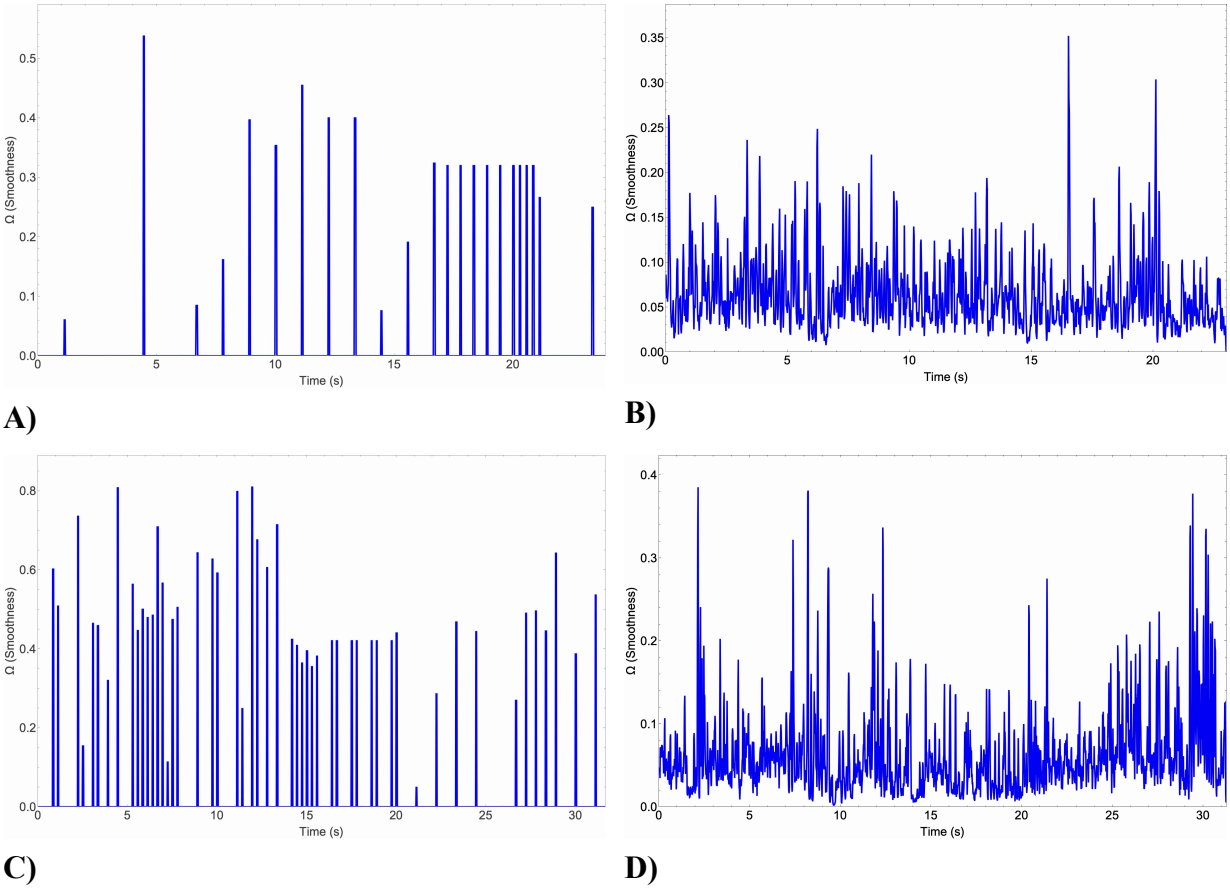


FIGURE 21. SMOOTHNESS Ω FOR TWO EXCERPTS FROM SYMPHONY NO. 40, BY W. A. MOZART: A) IDEAL MAP FOR EXCERPT 1 (SEE APPENDIX B, FIGURE B2), B) REAL TERRAIN FOR EXCERPT 1, C) IDEAL MAP FOR EXCERPT 2 (SEE APPENDIX B, FIGURE B4), D) REAL TERRAIN FOR EXCERPT 2.

Richard Wagner, *Tristan und Isolde*: the smearing of syntax

If Mozart is the control, Wagner is the first true test: a harmonic language built on withheld resolution, whose famous opening chord launches some four hours of music that refuses to settle. Four excerpts are analyzed[14]—three from the Prelude, one from the work's final cadence (notation in Appendix B, Figures B5–B8).

The first excerpt is the opening phrase of the Prelude, mm. 1–3, the most analyzed measures in the repertoire. The cellos rise through a yearning minor sixth and fall back

[14] Richard Wagner, *Tristan und Isolde*, in Full Score (New York: Dover Publications, 1973).

chromatically while the winds sound the famous "Tristan chord" (F–B–D♯–G♯); rather than resolving to any tonic, it slides chromatically onto the dominant seventh of A minor and simply stops, suspended over a fermata and silence. Nothing is affirmed—no key, no cadence—only an appetite for resolution that the whole opera will withhold. It is the founding gesture of musical longing.

The second excerpt, mm. 4–7, repeats that opening gesture higher, the first of the Prelude's rising sequences. The same shape—chromatic ascent in the winds over the cello's falling line, closing again on an unresolved dominant seventh—returns transposed upward, so that each repetition lifts the music another degree without ever arriving. This is Wagner's engine of "endless melody": tension renewed rather than released, the harmony always leaning forward. The ear is drawn higher and left, once more, hanging.

The third excerpt, mm. 8–11, climbs higher still and begins to intensify. The rising cells follow one another more closely and the chromatic voice-leading thickens as the music gathers toward the Prelude's first great swell; the harmony grows steadily less stable, each chord a chromatic step from the last with no cadence to anchor it. Here accumulation becomes palpable —the longing no longer merely stated but audibly building. The tonal ground has all but dissolved into continuous chromatic motion.

The final excerpt is the close of the concert *Liebestod*—the transfiguration of "Mild und leise," in which the work at last grants the resolution it has withheld across the entire opera. After a long chromatic ascent to its climax, the music subsides through a chain of suspensions onto a sustained, radiant B major, the goal toward which the drama has strained since the Prelude's first

unresolved chord. It is the most settled sonority in the whole excerpt set—a genuine tonal arrival after hours of evasion, the suspension of the opening finally released.

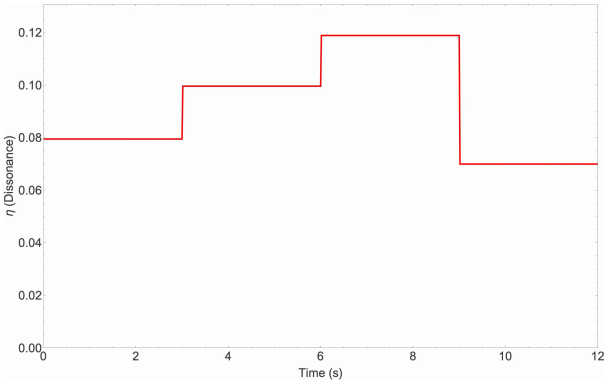
Dissonance η

The Prelude excerpts reveal what the Mozart baseline could not. In Mozart the ideal map η dropped to clean notches of near-consonance wherever the notated harmony simplified; in Wagner it never does. Across all three Prelude excerpts the ideal map trace holds an elevated plateau—stepping among values of roughly 0.045 to 0.119, but never descending toward the zero that marks true consonance (Figure 22). The notated harmony has no resting point: tension is not periodically released but sustained. The real terrain trace[15] confirms the reading—volatile, as ever, but volatile around a high mean, rarely and never durably falling to the low values that punctuated the Mozart performance.

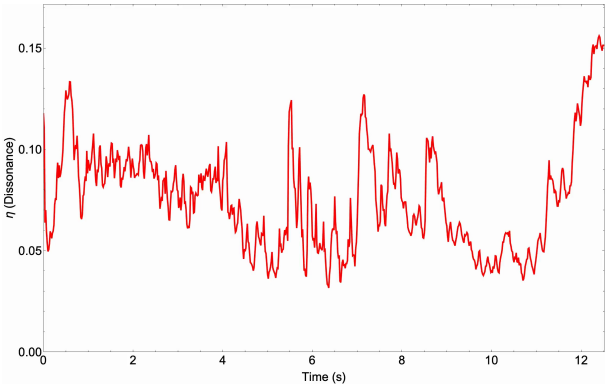
This yields a cross-repertoire observation worth stating plainly, because it is exactly the kind of thing the framework was built to make measurable. Reading the ideal map η floor across the pieces analyzed so far, a clear order emerges. Strauss's pure triads, a deliberately consonant reference point, sent it nearly to zero. Between Mozart and Wagner the floor rises: Mozart's functional harmony lifted it but still touched near-zero at every cadential consonance. Wagner's chromatic syntax holds it permanently above roughly 0.045 in the Prelude excerpts, with no consonant floor at all. The "emancipation of dissonance" that historians describe in prose appears here as a rising numerical baseline—the notated dissonance floor climbing from one style to the next. And as that floor rises, the magnitude of the map–terrain gap keeps shrinking: in Wagner

[15] Richard Wagner, *Tristan und Isolde: Prelude and Liebestod* (1996 Remastered Version), Berliner Philharmoniker, Conducted by Herbert von Karajan, EMI Classics, 1996, streaming audio.

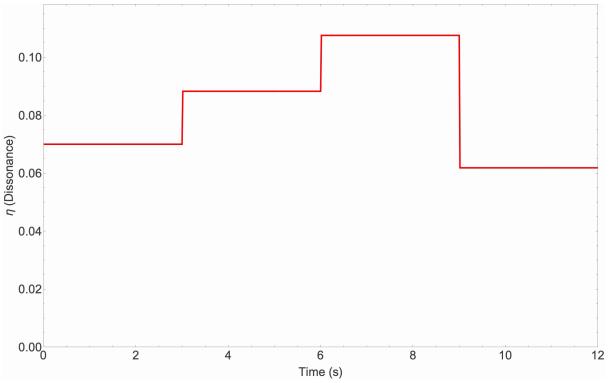
the ideal map and real terrain η occupy nearly the same band, so that what distinguishes them is almost entirely shape rather than level.



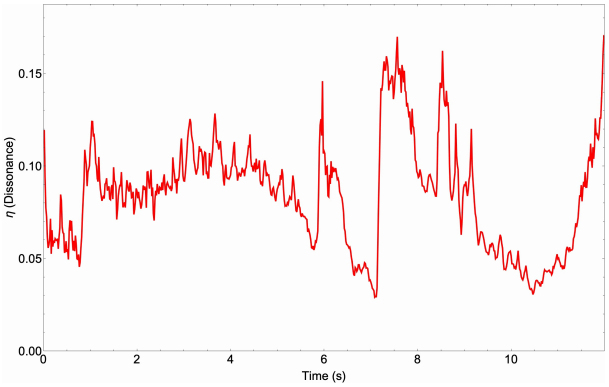
A)



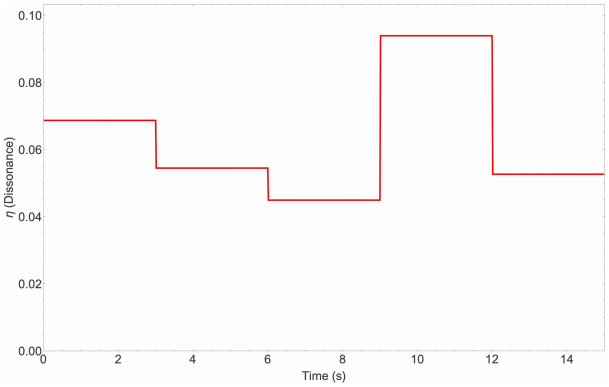
B)



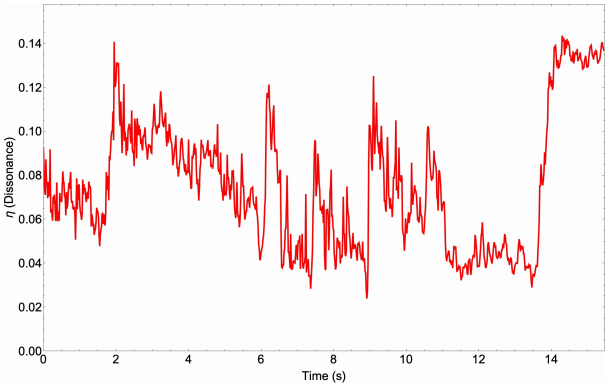
C)



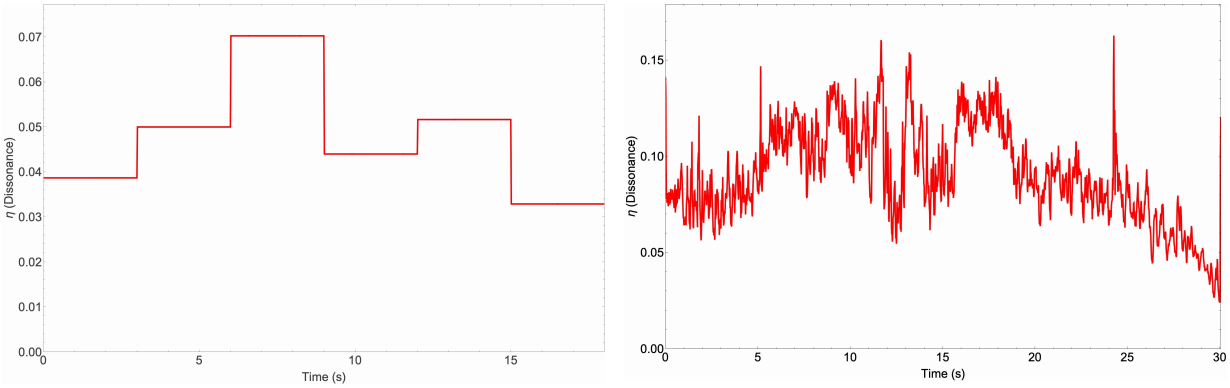
D)



E)



F)

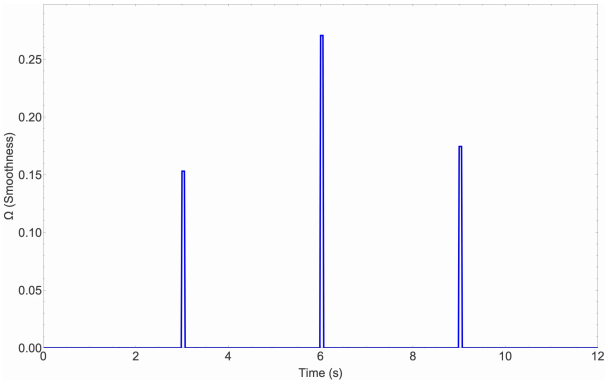


G) **H)**

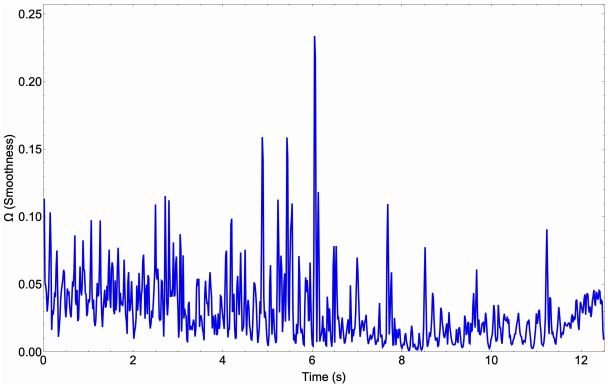
FIGURE 22. DISSONANCE η FOR FOUR EXCERPTS FROM TRISTAN UND ISOLDE, BY RICHARD WAGNER: A) IDEAL MAP FOR EXCERPT 1 (SEE APPENDIX B, FIGURE B5), B) REAL TERRAIN FOR EXCERPT 1, C) IDEAL MAP FOR EXCERPT 2 (SEE APPENDIX B, FIGURE B6), D) REAL TERRAIN FOR EXCERPT 2, E) IDEAL MAP FOR EXCERPT 3 (SEE APPENDIX B, FIGURE B7), F) REAL TERRAIN FOR EXCERPT 3, G) IDEAL MAP FOR EXCERPT 4 (SEE APPENDIX B, FIGURE B8), H) REAL TERRAIN FOR EXCERPT 4.

Smoothness Ω

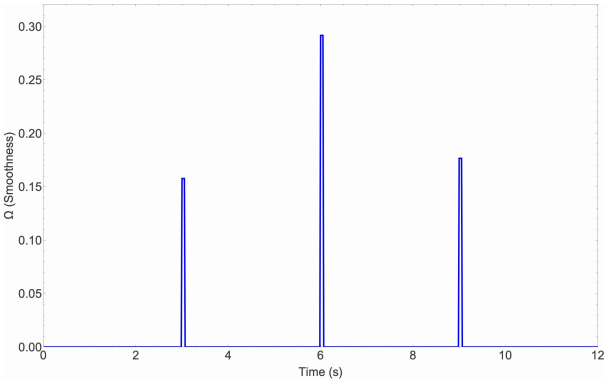
The ideal map Ω is again a row of clean isolated spikes at the notated chord changes (Figure 23). But the real terrain Ω does not merely add chatter between them—it smears the spikes into hills. Where a chord change in Mozart produced a single sharp real terrain spike embedded in low motion, in Wagner the elevated motion spreads across broad spans, as overlapping orchestration, dynamic swells, and constant chromatic voice-leading keep acoustic mass in motion between the notated changes. Across Figures 23B/D/F/H the terrain holds a low, evenly occupied floor of about 0.02–0.06 whose excursions — roughly 0.11 to 0.27 — are modest and spread through the span rather than gathered at the notated changes. The boundaries the score draws between one harmony and the next are, acoustically, blurred nearly to continuity—an apt image of a style whose whole rhetorical project is the evasion of clear arrival.


Ω (Smoothness)
Time (s)

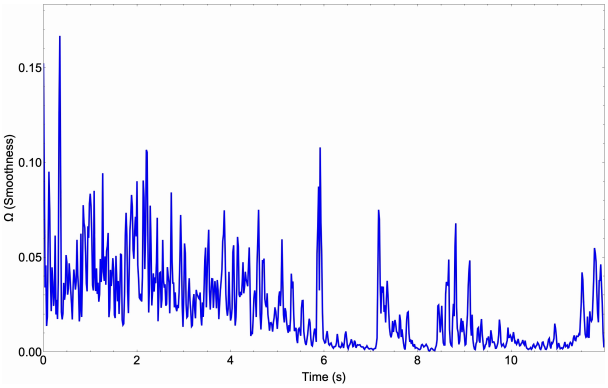
A)


Ω (Smoothness)
Time (s)

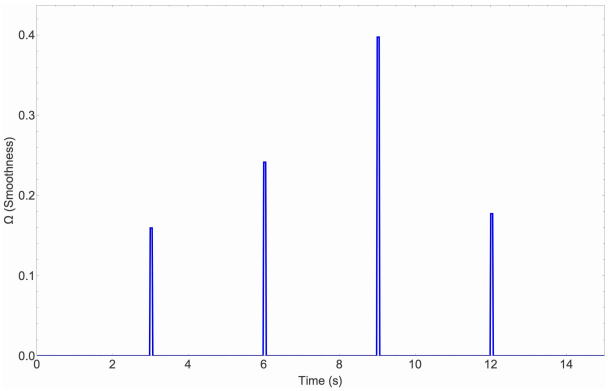
B)


Ω (Smoothness)
Time (s)

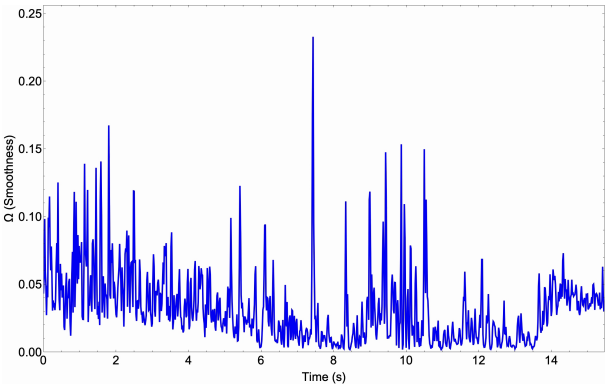
C)


Ω (Smoothness)
Time (s)

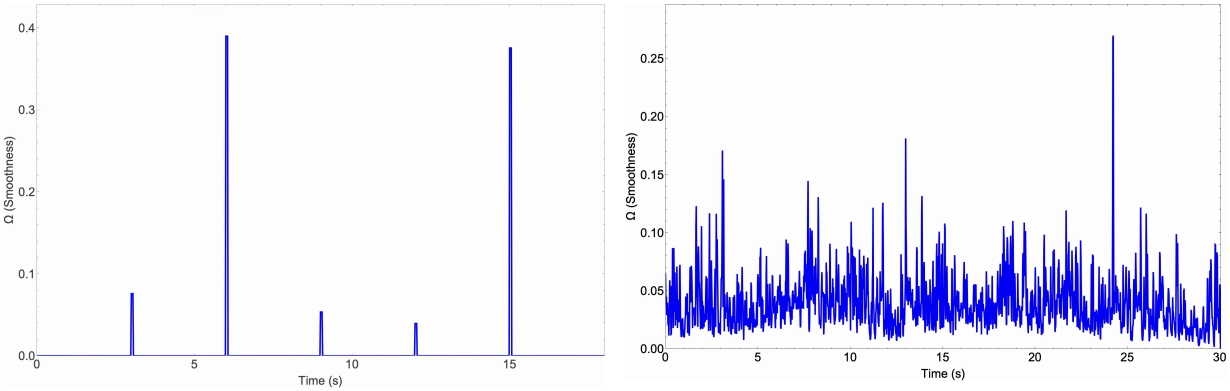
D)


Ω (Smoothness)
Time (s)

E)


Ω (Smoothness)
Time (s)

F)



G) **H)**

FIGURE 23. SMOOTHNESS Ω FOR FOUR EXCERPTS FROM TRISTAN UND ISOLDE, BY RICHARD WAGNER: A) IDEAL MAP FOR EXCERPT 1 (SEE APPENDIX B, FIGURE B5), B) REAL TERRAIN FOR EXCERPT 1, C) IDEAL MAP FOR EXCERPT 2 (SEE APPENDIX B, FIGURE B6), D) REAL TERRAIN FOR EXCERPT 2, E) IDEAL MAP FOR EXCERPT 3 (SEE APPENDIX B, FIGURE B7), F) REAL TERRAIN FOR EXCERPT 3, G) IDEAL MAP FOR EXCERPT 4 (SEE APPENDIX B, FIGURE B8), H) REAL TERRAIN FOR EXCERPT 4.

The structural payoff: the final cadence (excerpt 4)

The last excerpt is the resolution the Prelude spends the entire opera withholding. Here the ideal map η steps among lower values than the Prelude's (roughly 0.033 to 0.070), already reflecting a harmony beginning to consent to consonance; but it is the real terrain η that tells the story, rising to a turbulent 0.14 at the climax before declining across the final ten seconds to a settled 0.03 as the sonority resolves to B major and rings out (Figure 22H). The multi-hour harmonic suspension that began with the Tristan chord ends, in the terrain, as a measurable and sustained plunge in physiological tension—the framework registering, as a single downward drift, what the drama stakes its entire length upon.

Franz Liszt, *Nuages gris*: texture over syntax

Composed in 1881, Nuages gris ("Grey Clouds") is one of Liszt's late, experimental miniatures—a slow Andante in G minor that all but abandons functional cadence. The excerpt is

the piece's closing (mm. 33–44):[16] a rocking left-hand ostinato of low, oscillating sonorities beneath a spare, chromatic right-hand line, its harmonies sliding among minor, augmented, and whole-tone-tinged shapes rather than progressing by root motion. No clear tonic arrival ever comes; the music hovers, and the piece famously ends unresolved—a harmonic suspension later heard as a premonition of twentieth-century practice. It is a study in color and texture more than in chord-to-chord function.

Liszt's late miniature is where functional syntax begins to dissolve, and it is also the first solo-piano excerpt—which changes the shape of the terrain itself. One excerpt is analyzed. The notation is given in Appendix B in two forms: the actual arpeggiated, tied, and pedaled texture (the terrain), beside an idealized block-chord reduction (the map), Figures B9–B10.

Dissonance η and smoothness Ω

The ideal map traces behave as before—η a stepped plateau of moderate values (roughly 0.047 to 0.118, with two prominent steps up), Ω a row of discrete spikes at the notated changes. The real terrain traces look unlike anything in the orchestral excerpts.[17] Where Mozart's and Wagner's performed curves were continuous chatter—the signature of bowed and blown instruments sustaining between changes—Liszt's are a low floor punctuated by tall, narrow spikes. Both η and Ω spike sharply and briefly, and the spikes in the two metrics broadly coincide (Figure 24). Those coincident spikes are the piano's hammer strikes. Each spike rises from a floor near zero to between about 0.7 and 0.9 and falls back within a single frame: the

---

[16] Franz Liszt, *Nuages gris* (S. 199), in *Neue Ausgabe sämtlicher Werke*, Series I, vol. 12, ed. Imre Mező and Imre Sulyok (Budapest: Editio Musica Budapest, 1978).

[17] Franz Liszt, "Nuages gris, S. 199," on *Liszt*, André Laplante, piano, Analekta AN 2 9273, 1994, compact disc.

whole excursion completes inside fifty milliseconds, the signature of a struck attack rather than of a harmony changing. Maximum Ω across the terrains: Strauss 0.22, Wagner 0.27, Mozart 0.35, Collier 0.75, Varèse 0.79, Liszt 0.89. On a measure bounded at 1, a single piano hammer moves the virtual-pitch landscape nearly as far as it can go — further than any orchestral chord change in the set. The shape of the terrain has become a fingerprint of the instrument's physical sound production—struck-and-decaying rather than sustained—and the framework registers articulation and timbre as the dominant structural events, above the harmony itself.

Between the strikes, the real terrain η does not fall to the silence the ideal map reaches at the piece's end; it holds an elevated floor. This is the damper pedal: acoustic mass from each chord bleeds forward and overlaps the next, so the sound never fully clears.

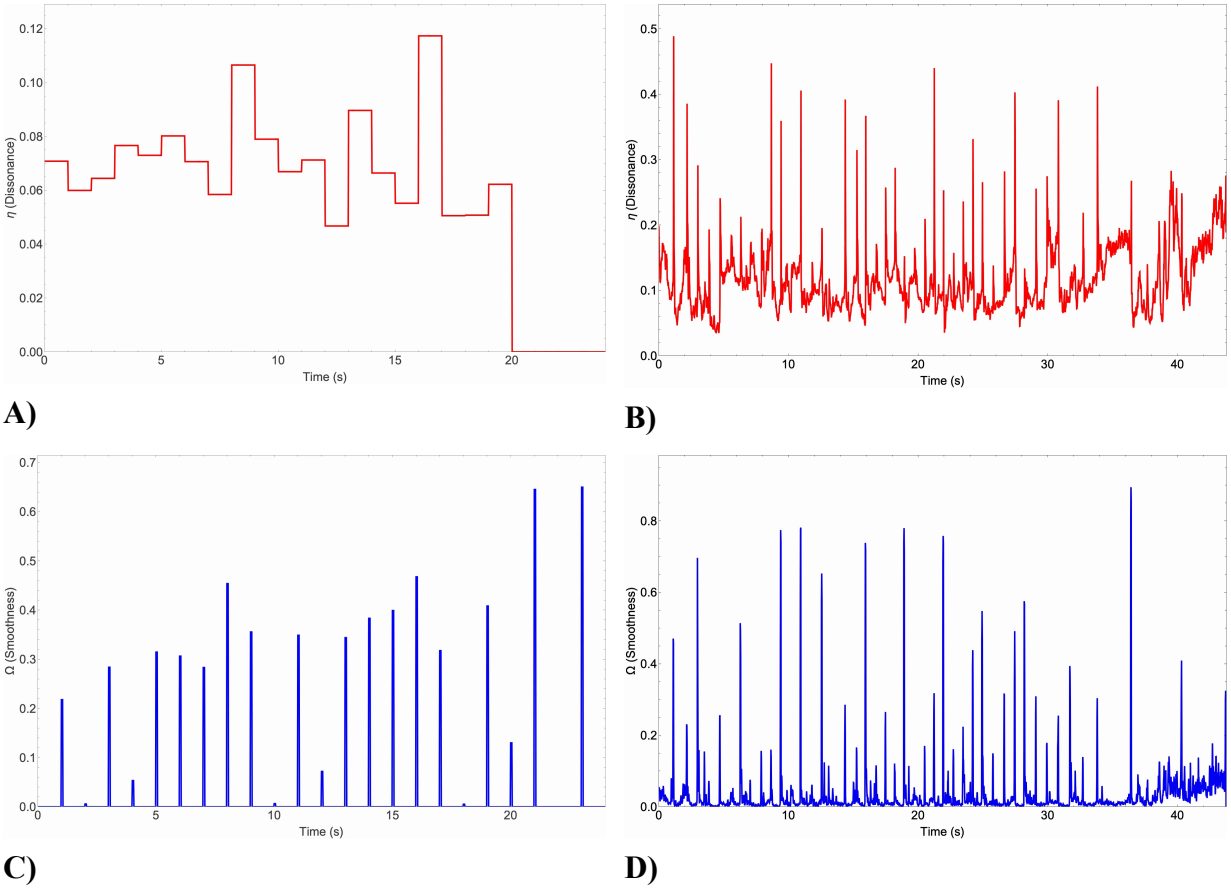


FIGURE 24. EXCERPT FROM NUAGES GRIS, BY FRANZ LISZT: A) IDEAL DISSONANCE η MAP (SEE APPENDIX B, FIGURES B9–B10), B) REAL DISSONANCE η TERRAIN, C) IDEAL SMOOTHNESS Ω MAP, D) REAL SMOOTHNESS Ω TERRAIN.

Jacob Collier, *In the Bleak Midwinter*: the structural role of tuning

Jacob Collier's a cappella arrangement is the framework's microtonal test, and here the ideal map trace disappears: the performance is freely and precisely microtonal, with no notation exact enough to render a meaningful MIDI map. The analysis is therefore of the terrain[18] alone—four excerpts, η and Ω—which is itself the point. Where no faithful map exists, the framework still measures.

Excerpt 1 (0:10–0:57) is the arrangement's introduction, before the first verse. It is mostly atmospheric: hushed, sustained vocal harmonies drift from one into the next without clear

[18] Jacob Collier, "In the Bleak Midwinter," self-released single, 2016, YouTube video, 6:12, https://www.youtube.com/watch?v=mPZn4x3uOac.

meter or cadence, and extramusical sounds—breaths, hums, and studio-processed textures—are woven through the choir. There is no tune yet, only color and a slowly shifting harmonic haze that establishes the arrangement's intimate, close-harmony sound-world.

Excerpt 2 (0:59–1:50) is the first verse. Here the Holst tune enters, sung in Collier's characteristically rich, jazz-inflected close harmony; the reharmonization is lush but still anchored to a clear key. The texture is fuller and more directed than the introduction, the melody now foregrounded against warm, chromatically colored chords.

Excerpt 3 (2:57–3:30) contains the arrangement's most celebrated moment: the modulation upward by a quarter tone, out of standard pitch altogether, followed by the start of the next verse in that raised tuning. The shift is executed as a smooth collective glide rather than a functional pivot—the entire choir bending microtonally upward together—so the music arrives in a tonal space unavailable to equal temperament. It is disorienting and luminous at once, the harmony intact but the ground silently raised beneath it.

Excerpt 4 (5:06–6:03) is the ending, reharmonized with the greatest complexity in the arrangement. The tune returns clothed in elaborate jazz voicings—extended and altered chords, dense chromatic voice-leading, and thick stacks of close-spaced vocal parts—as the piece gathers toward its final cadence. It is the most harmonically saturated writing of the whole track, luxuriant and intricate.

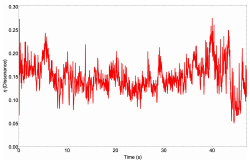
Dissonance η

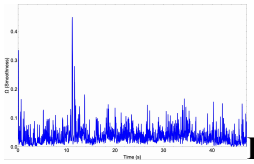
The striking feature is how low the dissonance runs for music whose notation, where it can be approximated at all, looks forbiddingly complex. Across the excerpts the real terrain η

holds a baseline near 0.14—comparable to Mozart's—and repeatedly drops into deep valleys around 0.03 to 0.06 (Figure 25). Those valleys are the harmonic core of the performance: by tuning chords toward just intervals, Collier aligns partials that equal temperament would leave beating, and the measured roughness falls accordingly. The claim that such valleys are unavailable in twelve-tone equal temperament is consistent with the tuning, though the present data include no 12-TET rendering to set beside it. What the data show directly is that extreme notational complexity coincides with low, and sometimes very low, acoustic dissonance.

Smoothness Ω

As in the Liszt, the real terrain **Ω** of excerpts 2, 3, and 4 is dominated by tall isolated spikes — reaching 0.75, 0.67, and 0.48 — over a moderate floor near 0.02–0.08 (Figure 25). Excerpt 1 is the exception, and it is the informative one: its terrain is almost spikeless, a continuous low chatter with a single excursion to 0.45 across forty-seven seconds. Excerpt 1 is also the only one of the four with no sung text — an atmospheric introduction of hummed and sustained vowels, before the tune enters. The same voices, the same studio, the same microtonal harmony, and the spikes appear exactly when words do. They are the phonetic attacks of the sung consonants, not harmonic change. The parallel with Liszt is exact in form: in both, the largest acoustic-mass excursions are events of articulation — there a hammer, here a consonant — rather than of harmony.





B)

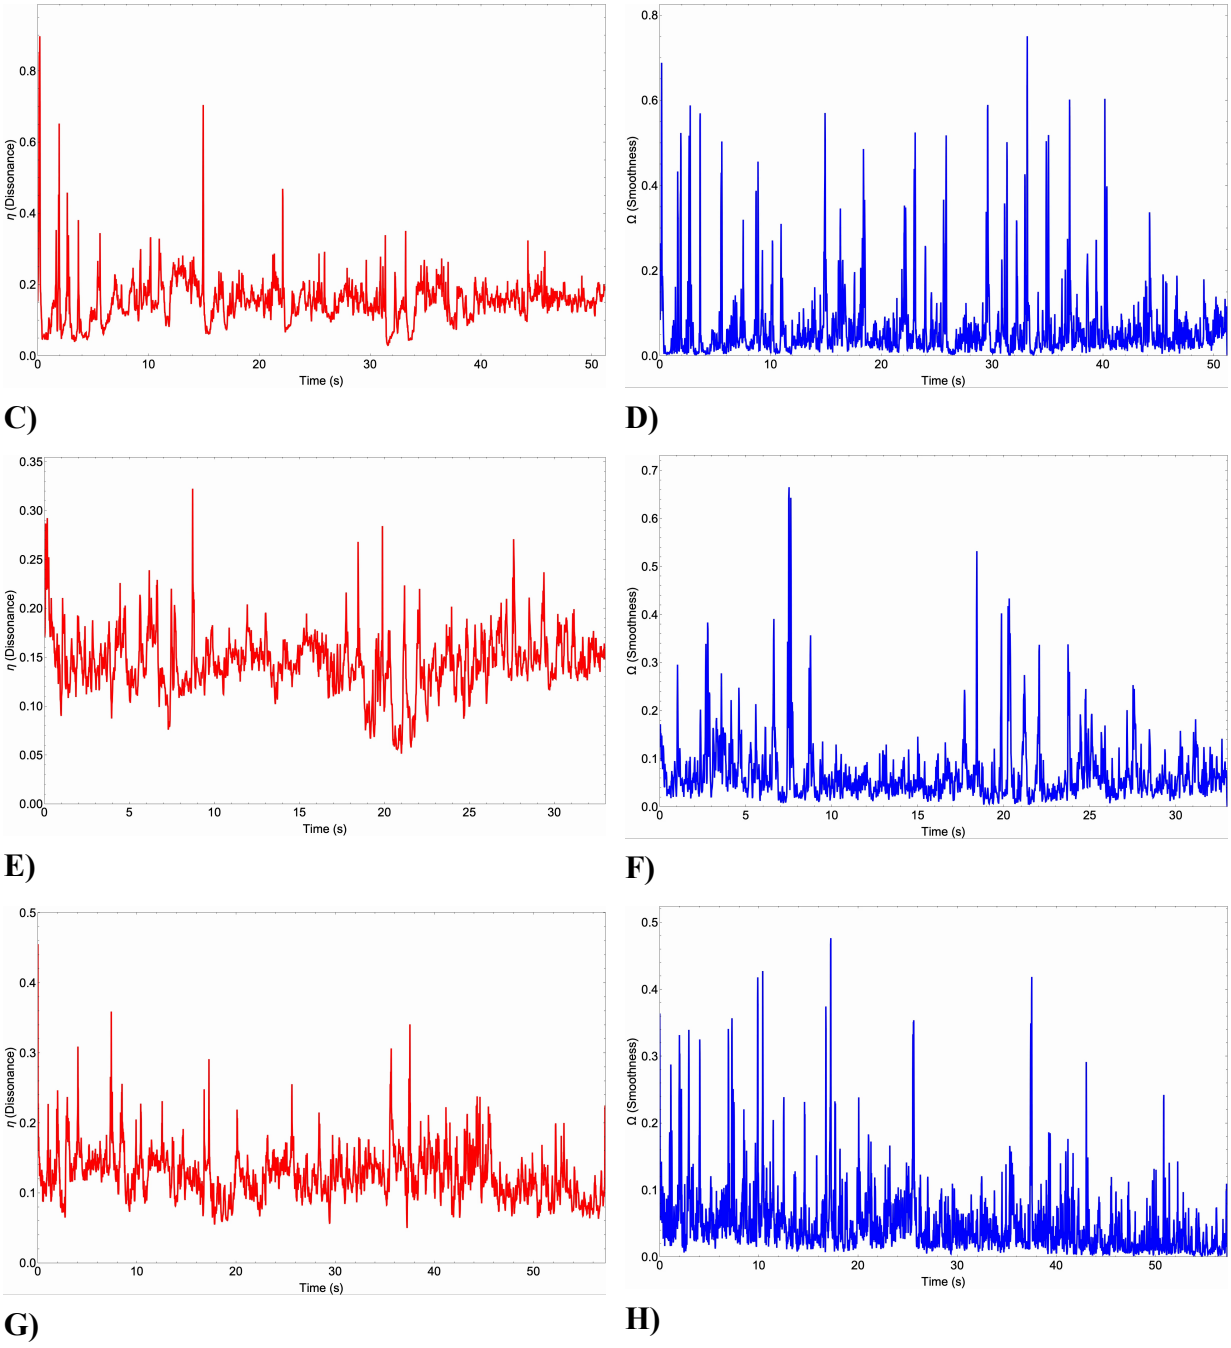


FIGURE 25. FOUR EXCERPTS FROM IN THE BLEAK MIDWINTER, BY JACOB COLLIER: A) REAL DISSONANCE η TERRAIN FOR EXCERPT 1 (0:10–0:57), B) REAL SMOOTHNESS Ω TERRAIN FOR EXCERPT 1, C) REAL DISSONANCE η TERRAIN FOR EXCERPT 2 (0:59–1:50), D) REAL SMOOTHNESS Ω TERRAIN FOR EXCERPT 2, E) REAL DISSONANCE η TERRAIN FOR EXCERPT 3 (2:57–3:30), F) REAL SMOOTHNESS Ω TERRAIN FOR EXCERPT 3, G) REAL DISSONANCE η TERRAIN FOR EXCERPT 4 (5:06–6:03), H) REAL SMOOTHNESS Ω TERRAIN FOR EXCERPT 4.

Synthesis

The lesson repeats the one the eight chords first taught in Part III: complexity on the page is not dissonance in the ear. Acoustically, Collier's microtonal writing sits in the same band as Mozart's — a comparable baseline with comparable dips — while looking, on the surface, like nothing Mozart could have written. The framework measures the sound, and the sound is smoother than the symbol.

Edgard Varèse, *Poème électronique*: the fluid dynamics of acoustic mass

Varèse is the limiting case, and the vindication of the whole approach. *Poème électronique* has no score in any conventional sense: it is organized sound—inharmonic masses and swept noise—for which notation is not merely approximate but inapplicable.[19] There is accordingly no ideal map trace, because there is nothing to render; yet η and Ω compute exactly as they did for a Mozart cadence (Figure 26). The same engine that measured functional tonality measures pure electronic sound, with no adjustment and no map.

The analyzed excerpt (2:00–2:48) is one of the most active stretches of the work. It is a dense montage of percussive attacks—timpani and snare drum among them—cutting across sustained electronic pitches, siren-like glissandi, and bursts of processed noise, with no pitch center or pulse to organize them. Material arrives in abrupt blocks, juxtaposed rather than developed, so that the passage reads as a churn of colliding sonic objects—among the busiest gestures in the piece.

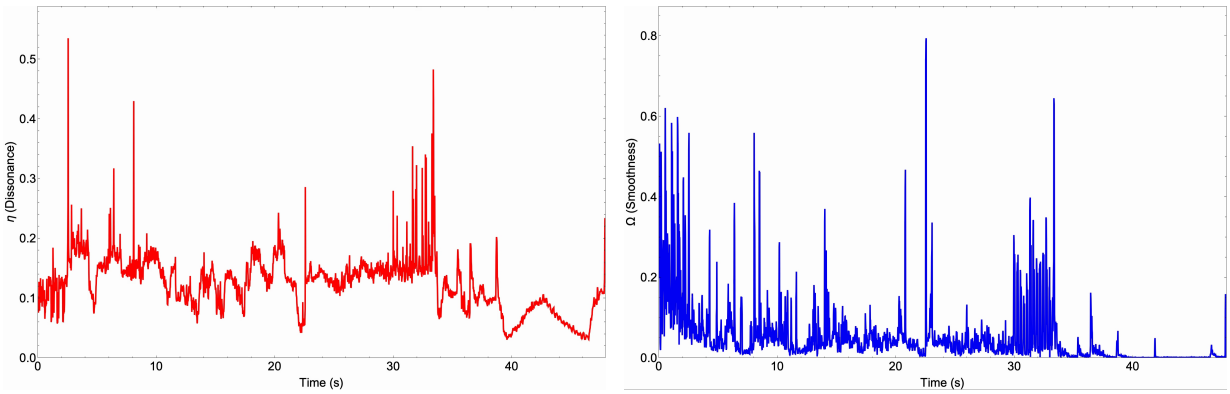
---

[19] Edgard Varèse, "Poème électronique," On *The Complete Works of Edgard Varèse,* Royal Conservatory Institute of Sonology transfer, Decca, 1998, compact disc.

Dissonance η and smoothness Ω

Both traces are volatile, as one expects of noise, but they are not random. η runs a baseline near 0.13 with sharp excursions past 0.5, and—tellingly—resolves in the final stretch into a low, smooth region around 0.03 to 0.10, a sustained near-consonant electronic texture. Ω tells the clearest story: dense clusters of tall spikes—turbulent passages of rapid acoustic-mass transport—alternate with calm, near-flat spans where little mass moves. η's smooth low stretch (39–47 s) falls inside Ω's calm span (34–48 s). The two independent measures agree on where the repose is. This alternation of activity and repose is arsis and thesis, tension and release, surviving in a medium with no chords, no meter, and no notes. Ω is tracking the physical migration of acoustic mass across the spectrum, and finding shape in it.

That is the framework's boldest claim made good: voice-leading, reconceived as the transport of acoustic mass, is not a property of notation but of sound, and applies as readily to Varèse's sound-masses as to a Mozart cadence. Where the map has vanished entirely, the territory remains measurable—and remains, recognizably, music.



A) **B)**

FIGURE 26. EXCERPT (2:00–2:48) FROM *POÈME ÉLECTRONIQUE*, BY EDGARD VARÈSE: A) REAL DISSONANCE η TERRAIN, B) REAL SMOOTHNESS Ω TERRAIN.

Part V. Conclusion: Grounding Synoptic Theory

Three tools, applied without modification across six works and four centuries, have held. It is worth drawing together what they found.

Three general results

First, a chord has no single root. It projects a continuous landscape of candidate fundamentals, and τ renders that landscape from the sound itself— distinguishing the fused C-major triad from the ten-note fifth-stack that shares none of its fusion, and both from the twelve-note cluster that nearly matches the fifth-stack in note count yet beats twenty-seven times as much. Second, physiological dissonance is set by register, timbre, spacing, and tuning at least as much as by interval class; η measures it as one normalized quantity, comparable across a triad, a Collier microtone, and a Varèse sound-mass alike. Third, harmonic motion is not a rulebook of prohibitions but the measurable transport of acoustic mass; Ω quantifies it for a Mozart cadence and for scoreless electronic noise on the same scale.

One arc

Read together, the repertoire tells a single story that no one piece could. In Strauss's pure triads the notated dissonance was almost nil while the performed sound carried real roughness—an eightyfold gap. In Mozart's harmonies the gap narrowed, because functional harmony already encodes genuine dissonance: the ideal map η lifted off the floor but still touched near-zero at harmonic consonances. In Wagner it narrowed further and the consonant floor vanished altogether—the notated harmony holds a permanent elevated plateau, tension sustained rather

than released, and the map–terrain difference reduces almost entirely to shape. The size of the gap, in other words, shrinks as the notation itself carries more dissonance. And from Mozart to Wagner the framework renders, as a rising number, the historical process theorists call the emancipation of dissonance.

Then the map itself begins to dissolve. In Liszt the notation's block-chord reduction cannot hold the pedaled, arpeggiated, floating texture the ear receives, and the terrain's structure passes to articulation—the hammer strike—rather than to harmony. In Collier the notation is only an approximation of a freely microtonal performance whose just-leaning tuning coincides with deep valleys in the measured dissonance. Measured, Collier's music proves acoustically closer to Mozart than to its own apparent complexity. And in Varèse the map is gone entirely—no score, no notes—yet η and Ω still find tension and release, arsis and thesis, in pure sound. The through-line that began with two chords of identical interval content sounding utterly different ends with music that has no notation at all still yielding to measurement.

What the framework is, and is not

None of this displaces traditional analysis. Tonalness theory addresses immediate apprehension—the bottom-up, moment-to-moment response of the ear to the acoustic signal—not the synoptic comprehension of a complete work's form, culture, and history. It does not explain why the *Tristan* Prelude means what it means, or how *Nuages gris* sits within Liszt's late style. What it offers is the layer beneath interpretation: the physiological reality on which any higher reading must ultimately rest. By folding critical band, frequency response, and virtual-pitch perception into the mathematics of the signal, a single framework measures the dissonance

of a Mozart cadence, a microtonal jazz voicing, and an electronic sound-mass on one scale—and so returns analysis, at last, to the sound itself, and not merely to the symbols that stand in for it.

Appendix A: Tonalness Equations

This appendix states the three measures the analysis uses — the tonalness function τ, the dissonance η, and the tonalness progression function (TPF) Ω, which measures smoothness — together with the auxiliary quantities each is built from, in their general continuous form. Thirteen equations are given in all. They are stated, not derived; the derivations, and the psychoacoustic models they rest on, are given in the companion paper.[20]

The dependency runs bottom-up. τ is built from a psychoacoustic Fourier function Y_h, which folds the critical band, the ear's frequency response, and virtual-pitch perception into ordinary Fourier analysis. Two readings are taken from τ over a short window: steady-state tonalness ψ_1 and tonalness variance ψ_2. The dissonance η is a ratio of the two at a single instant; the TPF Ω compares ψ_1 at two instants. Everything else below is scaffolding for those three.

Four constants recur. κ is a dimensionless scaling on the critical bandwidth. δ is a small positive regularizer that keeps a denominator or square root finite where the signal vanishes. It carries no interpretive weight. ζ is likewise small and positive, but it also sets the threshold below which a lower candidate counts as having no order-1 energy of its own (Equation 7). $T_c \approx 0.05$ s is the integration window, chosen to match the ear's temporal resolution for pitch.

The tonalness function, τ

The measures take a sound-pressure signal $P(t)$ and nothing else — no chord, no list of pitches, no score. Its ordinary Fourier spectrum is:

[20] Strohman, "Tonalness."

$$X(f) = 2\int_{-\infty}^{\infty} e^{-2\pi i t f} P(t)\, dt \tag{1}$$

Equation 1 is used nowhere in what follows. It is stated only as the object of contrast: the plain spectrum of Figures 3 and 5, against which the tonalness view of Figures 4 and 6 is set. In place of Equation 1, the framework uses a windowed transform evaluated at the *h*-th harmonic multiple of the candidate frequency *f*:

$$Y_h(f,t) = 2R(hf)\int_{-\infty}^{\infty} \frac{e^{-\frac{(u-t)^2}{2\hat{\sigma}_h(f)^2}} e^{-2\pi i (u-t) hf} P(u)}{\sqrt{2\pi}\hat{\sigma}_h(f)}\, du \tag{2}$$

At $h = 1$ this is an ordinary Gaussian-windowed transform of the signal at *f*, scaled by the ear's sensitivity there. At $h > 1$ it asks a different question: how much energy sits at the *h*-th harmonic of a fundamental that might be at *f*. It is this second reading that makes virtual pitch available to the mathematics — a fundamental the signal omits still leaves its harmonics behind, and Equation 2 finds them.

The transform's time window is the reciprocal of its bandwidth,

$$\hat{\sigma}_h(f) = \frac{1}{2\pi\sigma_h(f)} \tag{3}$$

so that a filter narrow in frequency is necessarily broad in time. The bandwidth itself is the Glasberg–Moore[21] equivalent rectangular bandwidth, scaled by κ and divided by the harmonic index:

[21] Brian R. Glasberg and Brian C. J. Moore, "Derivation of Auditory Filter Shapes from Notched-Noise Data," *Hearing Research* 47, no. 1–2 (1990): 103–38.

$$\sigma_h(f) = \kappa \left(\frac{24.7}{h} + 0.107939 f \right) \tag{4}$$

This is where the critical band enters. $\sigma_1(f)$, the $h = 1$ case, is the ordinary critical bandwidth at f, and it recurs below as the measure against which the band integrals of η, μ and Ω are taken. Integrating $df/\sigma_1(f)$ rather than df weights every region of the spectrum by how finely the ear resolves it, so a wide high-frequency band and a narrow low-frequency one contribute comparably.

The ear is not equally sensitive across the spectrum. R applies the standard A-weighting curve:[22]

$$R(f) = \frac{12194^2 f^4}{\left(f^2 + 20.6^2\right)\left(f^2 + 12194^2\right)\sqrt{\left(f^2 + 107.7^2\right)\left(f^2 + 737.9^2\right)}} \tag{5}$$

τ sums the harmonic evidence coherently:

$$\tau(f,t) = \left| \sum_{h=1}^{\infty} W_h(f,t) \left| Y_h(f,t) \right| \exp\left(\frac{i\Phi_h(f,t)}{h} \right) \right| \tag{6}$$

τ describes the strength with which the auditory system projects a virtual pitch fundamental at each frequency f and time t. The modulus is taken after the sum, not before, so harmonics whose phases agree reinforce one another and harmonics whose phases disagree cancel: a set of partials that genuinely belongs to one fundamental produces a tall, stable peak, while a coincidental

[22] International Electrotechnical Commission, *Electroacoustics—Sound Level Meters—Part 1: Specifications*, IEC 61672-1:2013, ed. 2.0 (Geneva: IEC, 2013); Harvey Fletcher and W. A. Munson, "Loudness, Its Definition, Measurement and Calculation," *Journal of the Acoustical Society of America* 5 (1933): 82–108.

alignment does not. Across the pitch range, τ is the landscape of candidate fundamentals that Part III reads. It depends on two auxiliary functions.

Inside Equation 6 we notice a coincidence gate function $W_h(f,t)$:

$$W_h(f,t) = \begin{cases} \sqrt{1 - \sum_{k=2}^{\infty} W_k(f/k,t)^2 \frac{|Y_1(f/k,t)|}{|Y_1(f/k,t)| + \zeta}} & \text{if } h = 1 \\ \frac{1}{h}\left(1 - \frac{|Y_h(f,t)|}{\sum_{g=1}^{\infty} |Y_g(f,t)| + \delta}\right) & \text{otherwise} \end{cases} \tag{7}$$

The gate admits a virtual pitch only when more than one harmonic reinforces it, vanishing where a candidate arises from a single harmonic alone. The two branches do different work. For $h > 1$, the $1/h$ coefficient makes each successive harmonic count for less, so a virtual pitch inferred from high harmonics is always subordinate in strength to a partial actually present in the signal. For $h = 1$, the spectral pitch at f surrenders weight to lower candidates at f/k — but only in proportion to whether those candidates have genuine order-1 energy of their own, which is what the ζ term tests. A real note lower down can claim credit from a partial above it. A coincidental harmonic ghost, with nothing sounding at its own fundamental, cannot. At first glance W_h appears self-referential, but the recursion terminates. W_1 depends only on W_k for $k > 1$, and those depend on Y_h alone.

Introducing the index h into the transform shifts the phase of each harmonic; Φ_h unwraps and conserves it so that the sum in Equation 6 remains coherent:

$$\Phi_h(f,t) = \arg\left(Y_h\left(f,t_0\right)\right) + \int_{t_0}^{t} \Im\left(\frac{\frac{\partial Y_h(f,u)}{\partial u}}{Y_h(f,u)}\right) du \tag{8}$$

Dissonance, η

In order to define the dissonance quantity η, we first define ψ_1:

$$\psi_1(f,t) = \max_{|u-t|<\frac{T_c}{2}} \tau(f,u) \tag{9}$$

ψ_1 is the peak of τ over the window T_c and serves as a proxy for perceived harmonicity at a given frequency and time. Taking the maximum rather than the mean is what separates pitch from rhythm: a landscape that holds its height across the window registers as a sustained harmonic presence, where one that merely flickers through it does not.

We also define ψ_2:

$$\psi_2(f,t) = \sqrt{\frac{1}{f} \max_{|u-t|<\frac{T_c}{2}} \left| \frac{\partial}{\partial u} \tau(f,u)^2 \right|} \tag{10}$$

ψ_2 is the rate of change of τ over the same window, frequency-scaled, and serves as a proxy for perceived roughness. Where two components fall close enough to collide within a critical band, τ beats rather than holding steady, and ψ_2 registers the flutter. The frequency scaling matters: the same absolute rate of change is heard as roughness low in the spectrum and as something closer to vibrato high in it.

ψ_1 and ψ_2 are independent readings of the same landscape. A sonority can be tightly fused yet rough, or dispersed yet smooth, and it is the pair that makes the eight chords of Part III separable. From these quantities we can define dissonance η:

$$\eta(t) = \frac{\int_0^\infty \frac{\psi_2(f,t)^2}{\sigma_1(f)}\, df}{\int_0^\infty \frac{\psi_1(f,t)^2}{\sigma_1(f)}\, df + \delta} \tag{11}$$

η measures the total roughness energy of a sonority normalized by its total harmonicity energy. Because it is a ratio of two band integrals taken over the same measure, it is independent of overall loudness and of the number of notes sounding: a single tension value for the whole sonority, comparable across a triad, a twelve-note cluster, and an electronic sound-mass alike.

Tonalness progression function, Ω

Where η reads one landscape, Ω compares two. It reuses ψ_1 from Equation 9 and requires one further auxiliary function μ.

$$\mu(t) = \sqrt{\int_0^\infty \frac{\psi_1(f,t)^2}{\sigma_1(f)}\, df} + \delta \tag{12}$$

μ is the norm of the landscape—the square root of its total harmonicity energy—at time t, measured in the same band integral η uses. It exists to be divided out. From this we can define Ω.

$$\Omega\left(t_1, t_2\right) = \frac{1}{2}\int_0^\infty \frac{1}{\sigma_1(f)}\left(\frac{\psi_1\left(f, t_1\right)}{\mu\left(t_1\right)} - \frac{\psi_1\left(f, t_2\right)}{\mu\left(t_2\right)}\right)^2 df \tag{13}$$
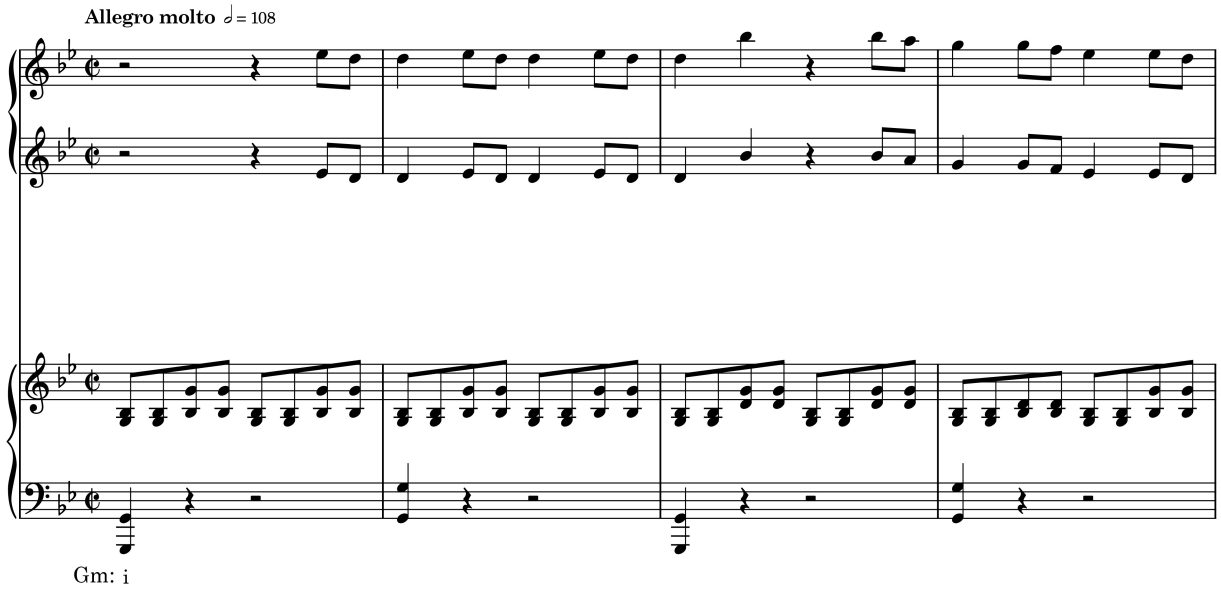
Ω measures the distance between two harmonicity landscapes: a single number for how far the sound masses are perceived to have moved from one moment to the next. Throughout Parts III and IV it is evaluated at $t_1 = t$ and $t_2 = t + T_c$ so that the plotted trace reports how far the landscape travels within one integration window.
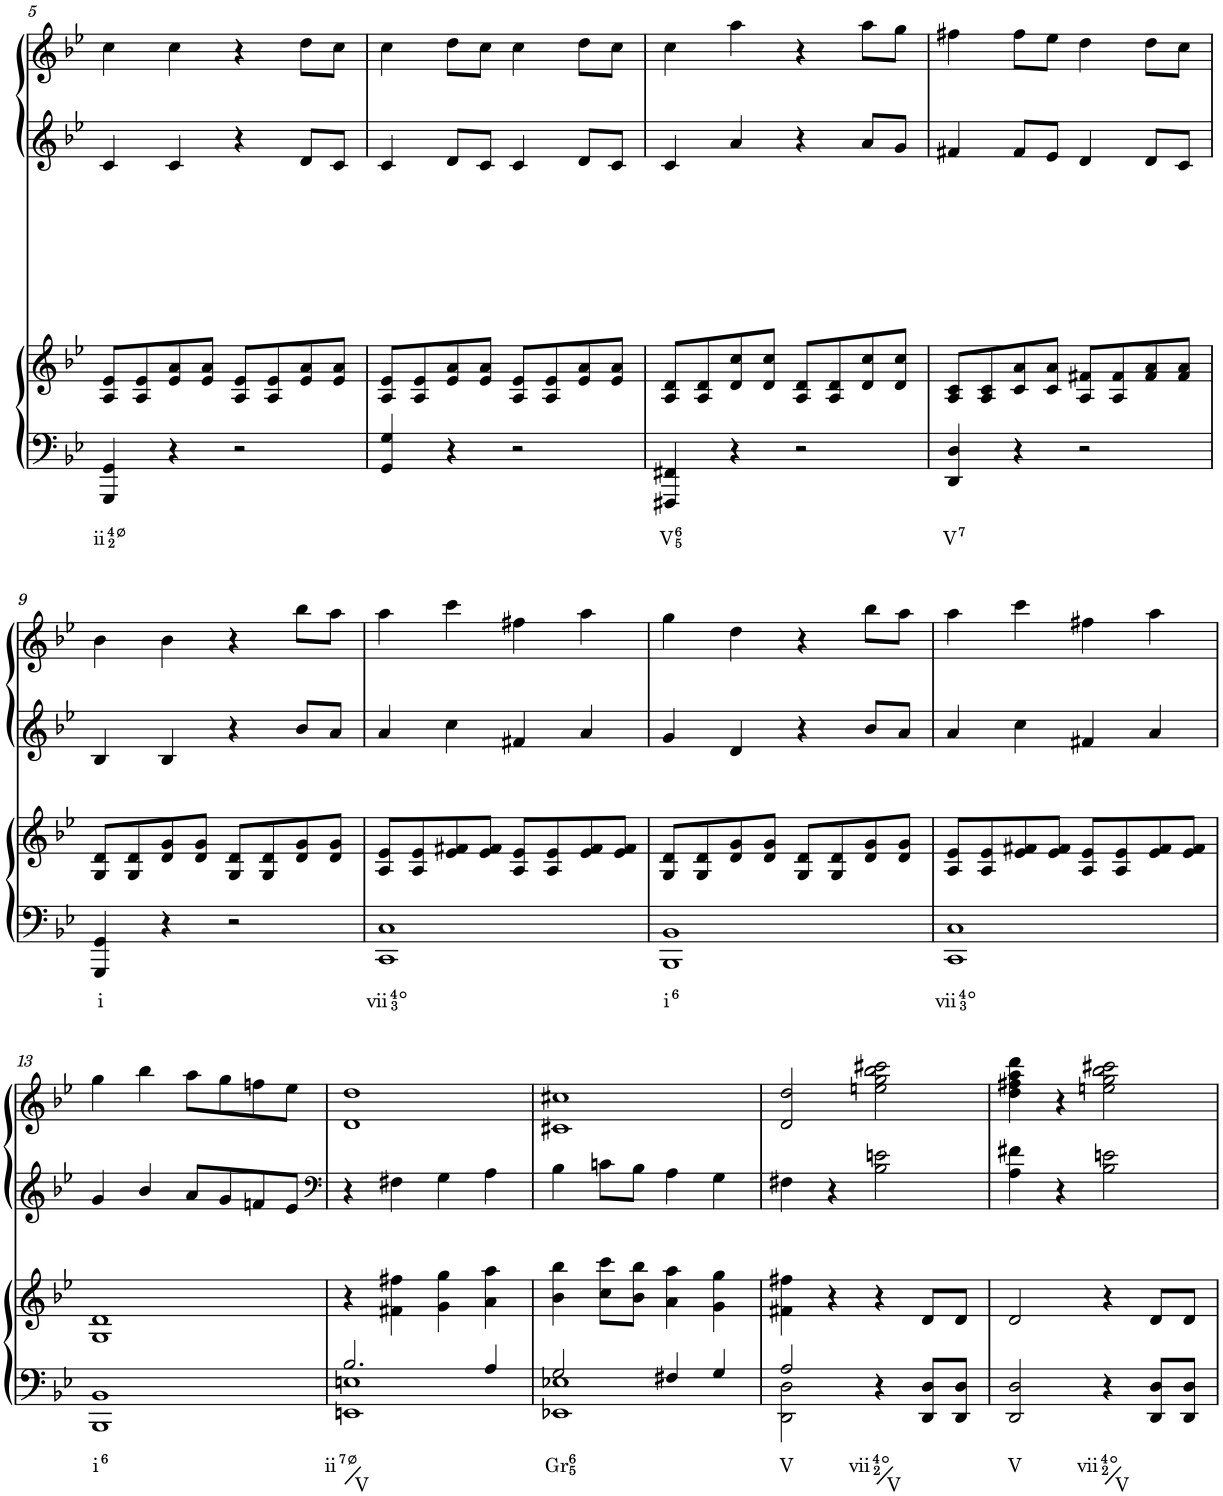
Normalizing each landscape by its own μ before differencing is what makes Ω a measure of harmonic motion rather than of loudness. Because each normalized landscape has unit energy by construction, Equation 13 is equivalent to one minus their overlap, and Ω is bounded between 0 and 1. It returns 0 when the two landscapes have the same shape, however much louder or

quieter one may be than the other, and 1 only when they share no harmonicity at all. A crescendo on a held chord therefore returns zero, and what the trace reports is the redistribution of acoustic mass across the spectrum.

Appendix B: Musical Examples

The notated sources for the analyses in Parts III and IV are collected here (except for Strauss, which already appeared in the main text). Several run to more than a single system, and are placed in the appendix to keep the main text readable. Examples for the Collier and Varèse analyses are omitted: the former exists only as an approximate transcription of a freely microtonal performance, and the latter has no score in any conventional sense.




5
ii 4 2 ø
V 6 5
V 7
9
i
vii 4 3 °
i 6
vii 4 3 °
13
i 6
ii 7 ø / V
Gr 6 5
V
vii 4 2 ° / V
V
vii 4 2 ° / V
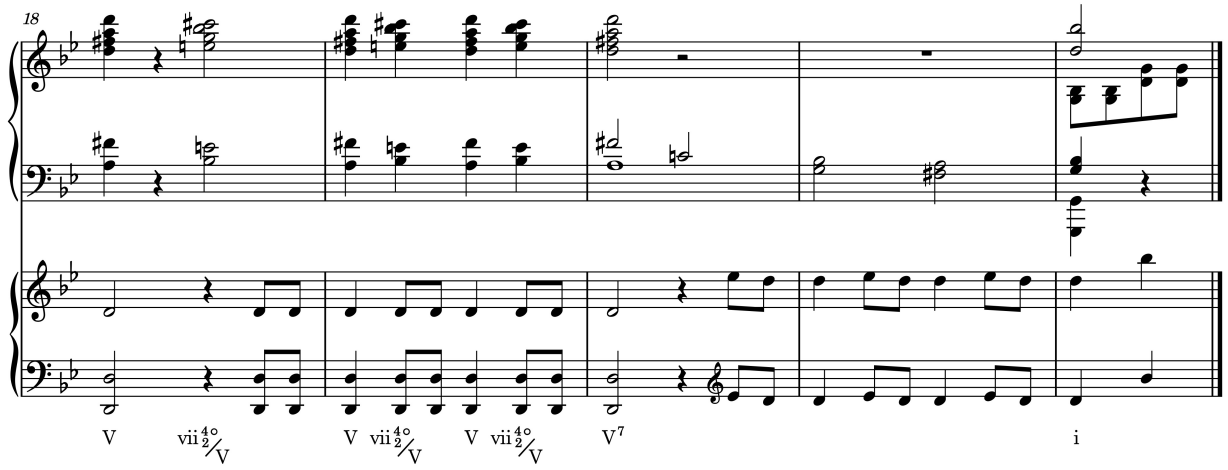


FIGURE B1. W. A. MOZART, SYMPHONY NO. 40 IN G MINOR, K. 550—EXCERPT 1.

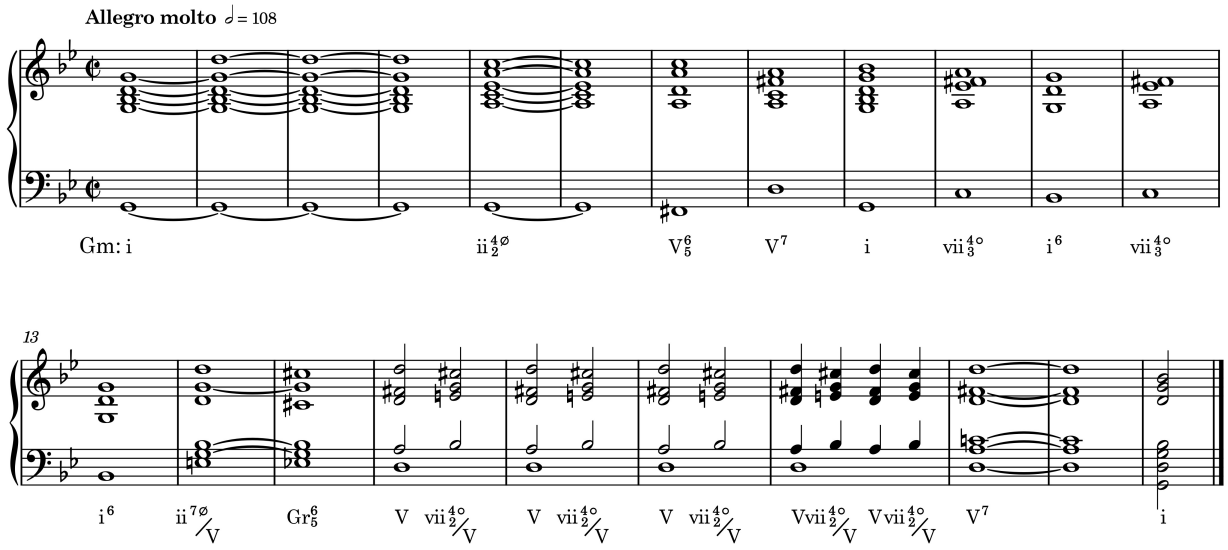
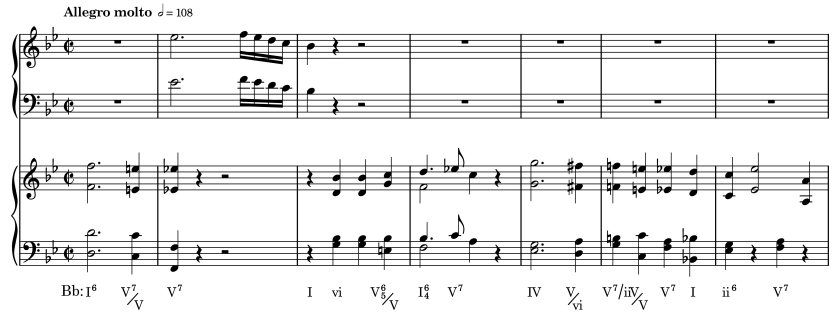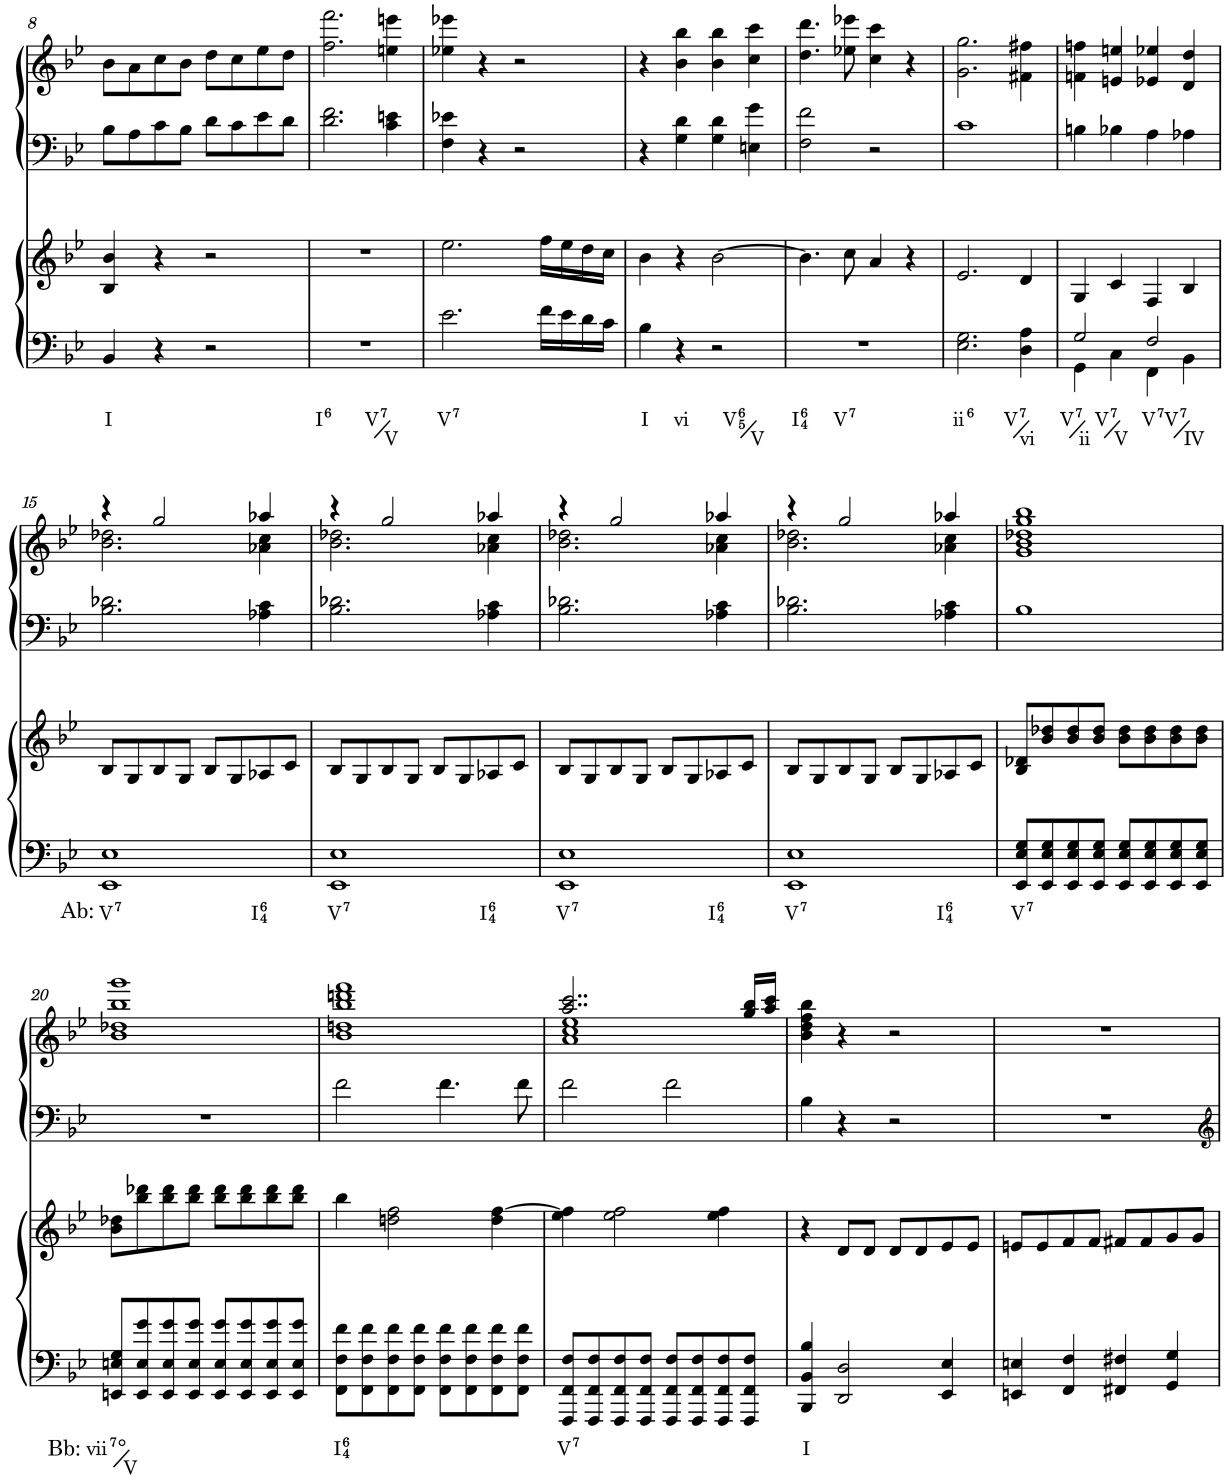


FIGURE B2. W. A. MOZART, SYMPHONY NO. 40 IN G MINOR, K. 550—EXCERPT 1, IDEALIZED BLOCK-CHORD REDUCTION.




8
I
I6
V7/V
V7
I
vi
V6 5/V
I6 4
V7
ii6
V7/vi
V7/ii
V7/V
V7V7/IV
15
Ab: V7
I6 4
V7
I6 4
V7
I6 4
V7
I6 4
V7
20
Bb: vii°7/V
I6 4
V7
I

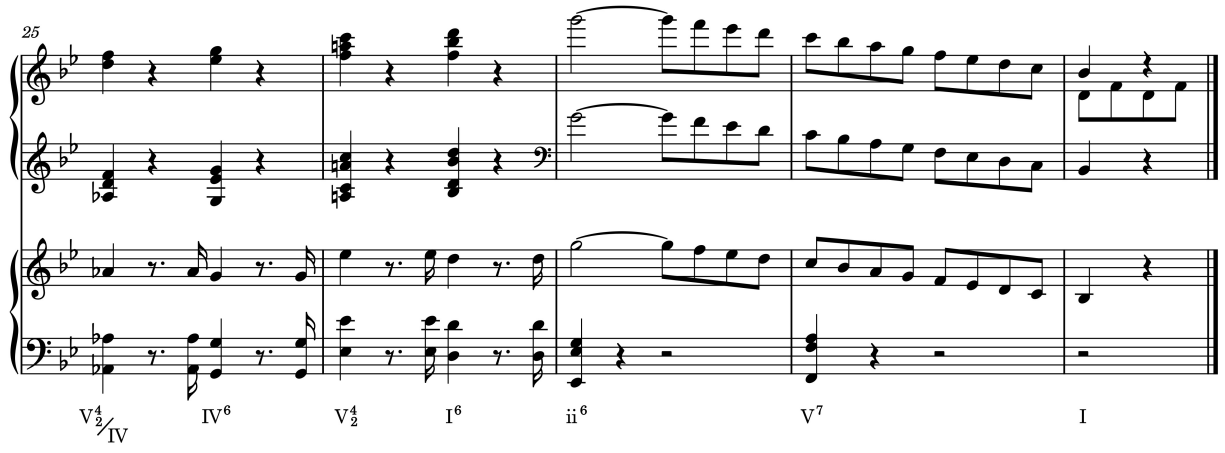


FIGURE B3. W. A. MOZART, SYMPHONY NO. 40 IN G MINOR, K. 550—EXCERPT 2.

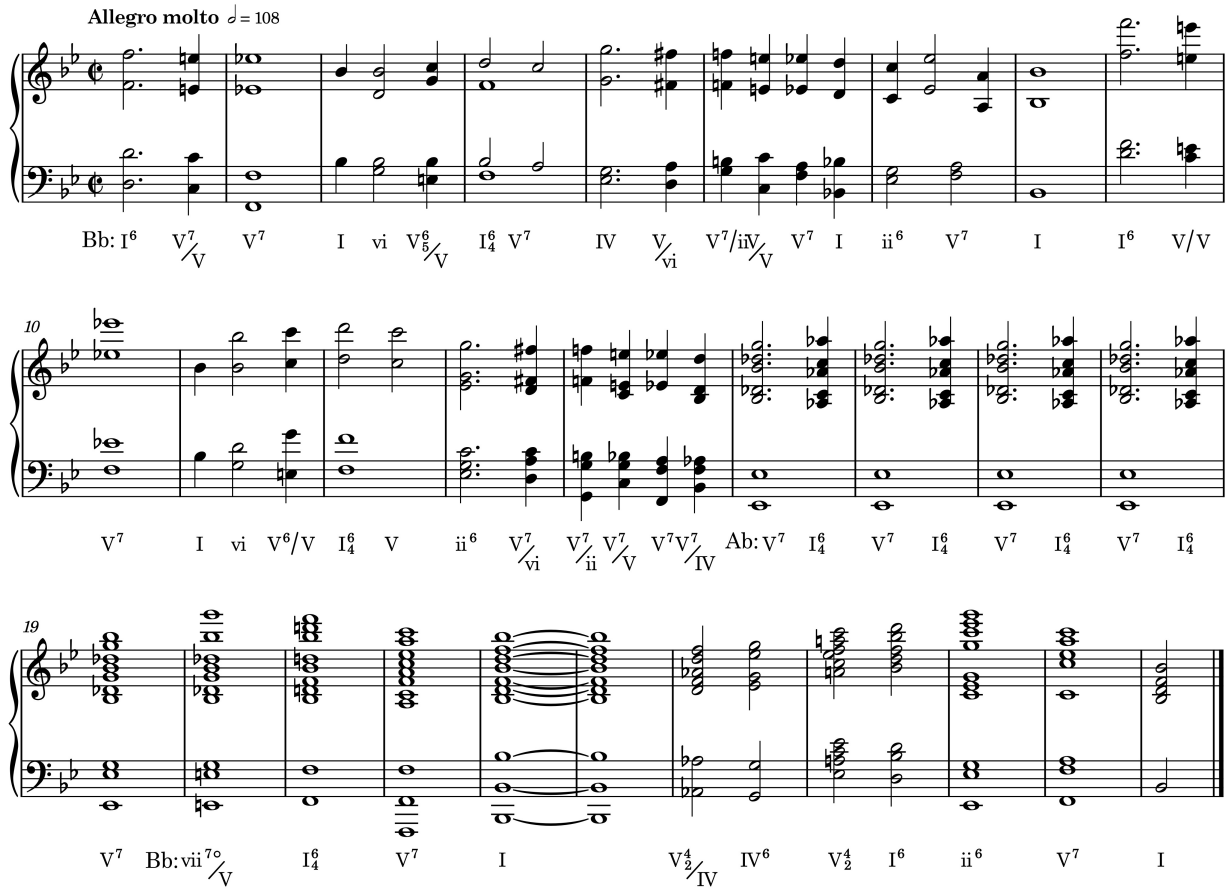


FIGURE B4. W. A. MOZART, SYMPHONY NO. 40 IN G MINOR, K. 550—EXCERPT 2, IDEALIZED BLOCK-CHORD REDUCTION.
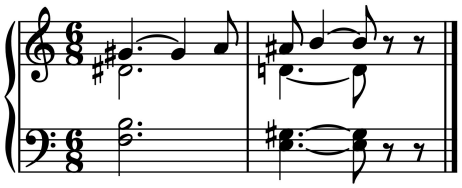


FIGURE B5. RICHARD WAGNER, *TRISTAN UND ISOLDE*—PRELUDE, EXCERPT 1.
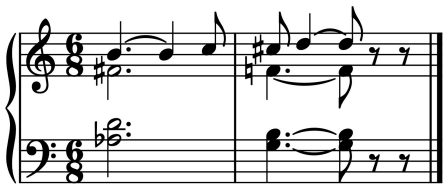


FIGURE B6. RICHARD WAGNER, *TRISTAN UND ISOLDE*—PRELUDE, EXCERPT 2.
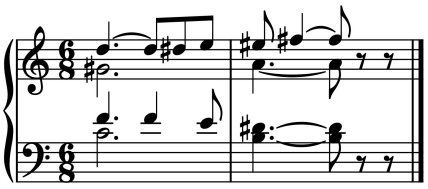


FIGURE B7. RICHARD WAGNER, *TRISTAN UND ISOLDE*—PRELUDE, EXCERPT 3.

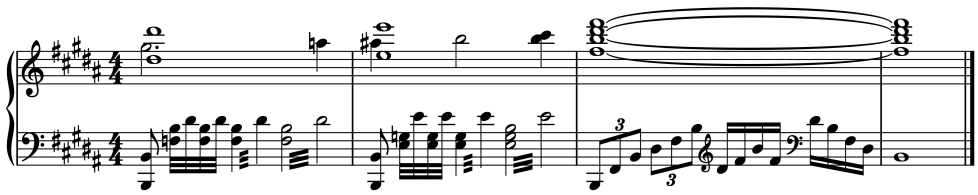


FIGURE B8. RICHARD WAGNER, *TRISTAN UND ISOLDE*—FINAL CADENCE, EXCERPT 4.
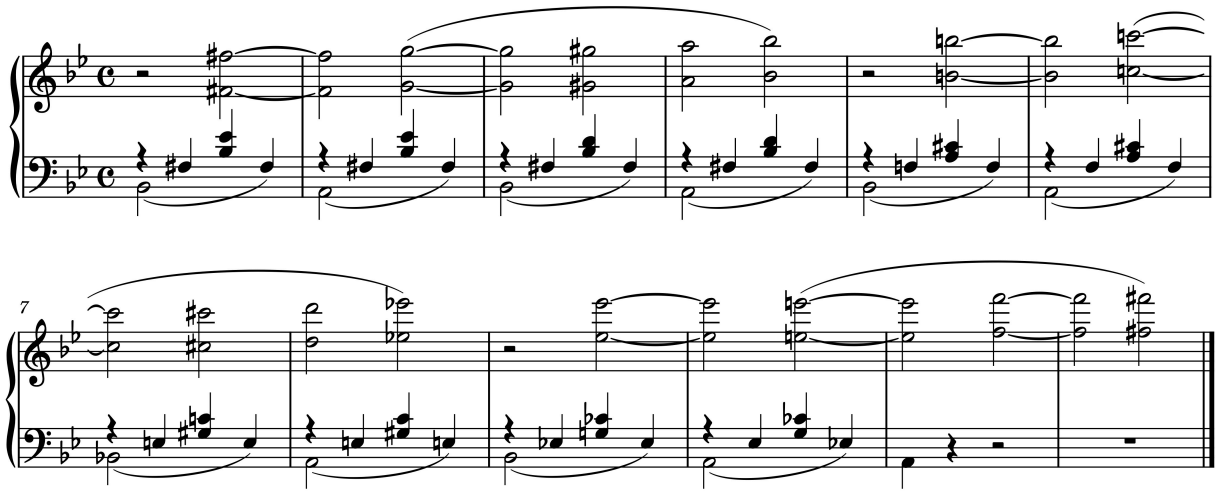


FIGURE B9. FRANZ LISZT, *NUAGES GRIS*—EXCERPT, ACTUAL NOTATION.
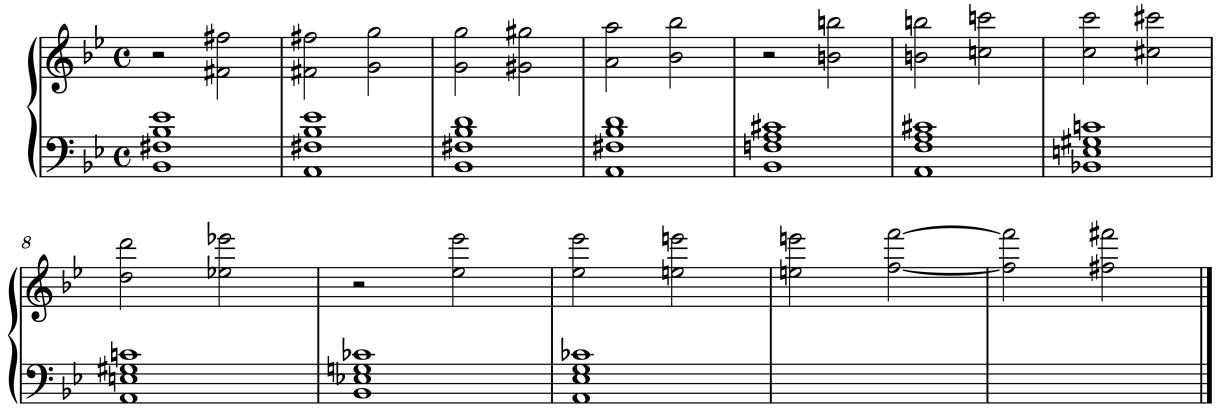


FIGURE B10. FRANZ LISZT, *NUAGES GRIS*—EXCERPT, IDEALIZED BLOCK-CHORD REDUCTION.

Bibliography

accentusmusic. “Lucerne Festival: Mozart – Symphony No. 40 (Berliner Philharmoniker, Simon Rattle).” YouTube video, 30:22. March 20, 2026. https://www.youtube.com/watch?v=Lk01TQnxsMk.

Berliner Philharmoniker. “Strauss: Also sprach Zarathustra / Dudamel · Berliner Philharmoniker.” YouTube video, 1:55. October 26, 2012. https://www.youtube.com/watch?v=Szdziw4tI9o.

Cohn, Richard. “Neo-Riemannian Operations, Parsimonious Trichords, and Their Tonnetz Representations.” *Journal of Music Theory* 41, no. 1 (1997): 1–66.

Collier, Jacob. “In the Bleak Midwinter.” Self-released single, 2016. YouTube video, 6:12. https://www.youtube.com/watch?v=mPZn4x3uOac.

Fletcher, Harvey, and W. A. Munson. “Loudness, Its Definition, Measurement and Calculation.” *Journal of the Acoustical Society of America* 5 (1933): 82–108.

Glasberg, Brian R., and Brian C. J. Moore. “Derivation of Auditory Filter Shapes from Notched-Noise Data.” *Hearing Research* 47, no. 1–2 (1990): 103–38.

Harrison, Peter M. C., and Marcus T. Pearce. “Simultaneous Consonance in Music Perception and Composition.” *Psychological Review* 127, no. 2 (2020): 216–44.

Helmholtz, Hermann von. *On the Sensations of Tone as a Physiological Basis for the Theory of Music*. Translated by Alexander J. Ellis. New York: Dover Publications, 1954.

International Electrotechnical Commission. *Electroacoustics—Sound Level Meters—Part 1: Specifications*. IEC 61672-1:2013, ed. 2.0. Geneva: IEC, 2013.

Lerdahl, Fred. *Tonal Pitch Space*. New York: Oxford University Press, 2001.

Liszt, Franz. Nuages gris (S. 199). In *Franz Liszt: Neue Ausgabe sämtlicher Werke* [New Edition of the Complete Works], Series I, vol. 12, edited by Imre Mező and Imre Sulyok. Budapest: Editio Musica Budapest, 1978.

———. “Nuages gris, S. 199.” On *Liszt*. André Laplante, piano. Analekta AN 2 9273, 1994. Compact disc.

Mozart, Wolfgang Amadeus. Symphony No. 40 in G Minor, K. 550. Edited by H. C. Robbins Landon. Kassel: Bärenreiter, 1958.

Parncutt, Richard. *Harmony: A Psychoacoustical Approach*. Berlin: Springer-Verlag, 1989.

———. “Revision of Terhardt’s Psychoacoustical Model of the Root(s) of a Musical Chord.” *Music Perception* 6, no. 1 (1988): 65–93.

Plomp, Reinier, and Willem J. M. Levelt. “Tonal Consonance and Critical Bandwidth.” *Journal of the Acoustical Society of America* 38, no. 4 (1965): 548–60.

Rameau, Jean-Philippe. *Traité de l’harmonie réduite à ses principes naturels*. Paris: Jean-Baptiste-Christophe Ballard, 1722.

Sethares, William A. “Local Consonance and the Relationship between Timbre and Scale.” *Journal of the Acoustical Society of America* 94, no. 3 (1993): 1218–28.

———. *Tuning, Timbre, Spectrum, Scale*. 2nd ed. London: Springer, 2005.

Strohman, Greg. “Tonalness: A Signal-Domain Transform for Virtual Pitch, with Implications for Musical Harmony.” Unpublished manuscript, last modified October 2026. PDF file.

Terhardt, Ernst. “Calculating Virtual Pitch.” *Hearing Research* 1, no. 2 (1979): 155–82.

Tymoczko, Dmitri. *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*. New York: Oxford University Press, 2011.

Varèse, Edgard. “Poème électronique.” On *The Complete Works of Edgard Varèse*. Decca, 1998. Compact disc.

Wagner, Richard. *Tristan und Isolde*. In Full Score. New York: Dover Publications, 1973.

———. *Tristan und Isolde: Prelude and Liebestod*. Berliner Philharmoniker, conducted by Herbert von Karajan. 1996 remastered version. EMI Classics, 1996. Streaming audio.