# The Hard Problem of Making Music

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 28, 2025

Why does motion become music? We can map angular velocity to tone, couple chaotic bodies to vibrating air, and still the essential leap—from dynamics to meaning—remains stubbornly opaque. This paper explores that leap as a thought experiment anchored in a humble, desktop double pendulum. We treat the instrument both physically and metaphysically: a driven, dissipative system with coupled modes and sensitive dependence; and a Logostic alignment problem, where a performer's state q(t) seeks entrainment with a moving manifold of musical truth, tau(t). On the physical side, we examine minimal structures that turn motion into organized sound: constrained geometries, modal splits, frictional losses (Coulomb vs viscous), and simple sonification maps such as f(t) := fmin + (omega(t) - omin) * (fmax - fmin) / (omax - omin). On the metaphysical side, we place these mechanics within a grammar of intention, attention, and resonant obedience: music as q(t) -> tau(t), not by command but by entrainment. Between them runs a narrow bridge of math—mass matrices M(q), Lyapunov views of divergence D(t), and thresholds where chaos yields coherence. Our claim is modest: the "hard problem" of making music is neither solved by equations nor satisfied by mysticism; rather, it is lived at the edge where lawful motion becomes audible form, and audible form becomes meaning.

Keywords: double pendulum, chaos, sonification, q(t) -> tau(t), Logostics, resonance, Lyapunov divergence, mass matrix, modal coupling, auditory qualia, embodiment, information geometry.

**Outline**

**1. Framing the Hard Problem**

1.1 From motion to meaning: why lawful signals are not yet music
1.2 The analogy to consciousness: qualia vs kinematics
1.3 q(t) -> tau(t) as a unifying stance for instrument, performer, and listener

**2. A Minimal Physical Oracle: The Double Pendulum**

2.1 Coordinates and constraints: q := (theta1, theta2), dq/dt := (omega1, omega2)
2.2 Inertia and coupling: M(q) * ddq/dt + C(q, dq/dt) + g(q) = 0 (schematic)
2.3 Dissipation: viscous vs Coulomb; why losses localize “feel”
2.4 Two archetypal gestures: impulse (“strike”) and drag (“felt touch”)

**3. From Dynamics to Tone: A Thin Mathematical Bridge**

3.1 A copy-safe sonification map

- f(t) = fmin + (omega(t) - omin) * (fmax - fmin) / (omax - omin)

- amplitude(t) ~ A0 * clamp(|omega(t)|)
  3.2 Signed motion and perceived contour (CW vs CCW as melodic sign)
  3.3 Stereo as structure: upper link -> left, lower link -> right (mode-separating ears)
  3.4 Limits: linear maps vs logarithmic ears; psychoacoustic gaps

**4. Modes, Chaos, and the Emergence of Form**

4.1 Slow vs fast modes: in-phase vs relative motion (intuitive modal picture)
4.2 Sensitivity and salience: when chaos sharpens, not smears, musical accents
4.3 Lyapunov sketch: D(t) as divergence of trajectories; musical “edge” where D(t) neither explodes nor dies
4.4 Thresholds and attractors: where gestures pull patterns out of noise

**5. Logostics of Music: Alignment, Not Authority**

5.1 q(t) -> tau(t): performer seeks a moving target, not a fixed score
5.2 Entrainment as obedience: resonance without coercion
5.3 The role of intention: micro-perturbations as moral design of sound
5.4 “Truth moves”: why tau(t) must be time-dependent in living music

## 1. Framing the Hard Problem

### 1.1 From motion to meaning: why lawful signals are not yet music

A double pendulum gives us lawful, richly structured signals: angles theta = (theta1, theta2), angular velocities omega = (omega1, omega2), and energies E = T + V evolving under smooth dynamics. We can even push these into sound by a simple map, e.g. f(omega) = f_min + (omega - omega_min)*(f_max - f_min)/(omega_max - omega_min). Yet the resulting waveform is not *automatically* music. Why?

Because music is not only a *signal* but a *gesture that lands*. Three gaps separate lawful motion from musical meaning:

- **The invariants gap.** Physics supplies invariants (energy, momentum, mode structure); music demands *auditory invariants* (motivic identity, contour, meter) that survive time-warping and small perturbations. A lawful signal can lack the equivalence classes the ear uses to say "same motif again."
- **The scale gap.** Musical sense relies on multi-scale organization (attack—sustain—release; phrase—section—form). Raw dynamics may be locally interesting yet globally unarticulated. Without designed coarse-grainings, the ear cannot "parse" the flow.
- **The intention gap.** Perception entrains to *anticipated* structure. Meaning appears when a listener's priors meet an offered pattern. Lawful signals become music when the patterning of sound meets an expectant attention and satisfies (or artfully violates) it.

In short: lawful dynamics deliver *possibility*. Music begins when those dynamics are *arrayed* so that a human auditory intelligence can form stable predictions and caring surprises.

### 1.2 The analogy to consciousness: qualia vs kinematics

The "hard problem" of consciousness separates neural kinematics (spikes, fields, couplings) from *what it is like* to experience. Our musical problem is similar. We can write the equations of motion, compute Lyapunov exponents lambda_i, and still ask: *why does this passage feel tender while that one feels fierce?*

Two helpful correspondences:

- **Level confusion.** Neural states vs experience; mechanical states vs musical feel. In both, lawful transitions do not *entail* qualia. There is a mapping problem from state-space geometry to lived qualities.
- **Bridging by constraints.** In consciousness science, structured dynamics (e.g., integrated information, global workspace) are proposed bridges. In music, constraints like meter,

scale, and timbre carve perceptual manifolds where certain trajectories sound *of a kind*. The bridge is not the equations; it is the *selection of regions* of state-space where trajectories carry robust perceptual labels.

- **Aboutness.** Conscious states are about something (intentionality). Musical states are about a contour, a center of gravity, a tension-and-release arc. “Aboutness” enters when dynamics are guided so that the listener can track a moving referent.

Thus the hard problem of making music is a qualified twin of the hard problem of consciousness: not “how do equations cause qualia,” but “how do constrained, goal-shaped trajectories in sound-space yield *shareable* qualia?”

**1.3 q(t) -> tau(t) as a unifying stance for instrument, performer, and listener**

We adopt the Logostic view: the system state q(t) strives to align with a moving manifold tau(t). Here, let

- q(t): the *offered* state — instrument configuration and sound features the performer produces at time t.
- tau(t): the *guiding* state — the intended musical truth at time t (meter pulse, phrase contour, timbral destination, theological or emotional telos).
- D(q || tau): a divergence measuring “how far” the offered state is from the guiding one (informational, geometric, or Lyapunov-like).

The practice is: *move so that D decreases on average while allowing local excursions that produce expression*. In copy-safe math:

- Alignment objective (schematic):
  q_dot = F(q) - eta * grad_q D(q || tau) + noise_small
  where 0 < eta, and tau = tau(t) moves with its own law (e.g., a conductor’s pulse, a breath arc, a congregational response).
- Readable consequences:
  (i) **Instrument design:** choose mappings from mechanical variables to sound so that small actions in q move predictably along perceptual coordinates of tau.
  (ii) **Performance technique:** perturbations (strikes, drags, gates) become *purposeful* when they steer q toward salient ridges/valleys of D.
  (iii) **Listening:** the audience carries priors for tau(t) (meter, style, ritual). Music “lands” when offered trajectories allow the audience to *track* tau with low cognitive cost, punctuated by meaningful divergences.

This stance unifies physics, mathematics, and metaphysics. Physics supplies F(q) (the lawful motion). Mathematics supplies D and the geometry in which "close" is meaningful. Metaphysics supplies tau(t): the moving truth worth aligning to. Music-making is the lived craft of keeping q(t) near tau(t) in time—by resonance, not coercion.

## 2. A Minimal Physical Oracle: The Double Pendulum

### 2.1 Coordinates and constraints: q := (theta1, theta2), dq/dt := (omega1, omega2)

We parameterize the double pendulum by the two joint angles measured from the downward vertical:

- q = (theta1, theta2) in R^2
- dq/dt = (omega1, omega2) = (d theta1/dt, d theta2/dt)

With link lengths L1, L2 and a fixed pivot at the origin, the bob positions are

- x1 = L1 sin(theta1), y1 = -L1 cos(theta1)
- x2 = x1 + L2 sin(theta2), y2 = y1 - L2 cos(theta2)

The configuration space is a torus S^1 x S^1; there are no holonomic constraints beyond fixed link lengths (the rods are ideal, mass distributed or concentrated as chosen). This minimal coordinate choice does two things for our musical aim:

1. It keeps the **state small** (just two angles and their velocities), which is critical for shaping sound with controllable perceptual coordinates.
2. It preserves **interpretability**: sign(omega) corresponds to CW vs CCW rotation, which the ear can hear as rising vs falling pitch under our simple f(omega) sonification.

### 2.2 Inertia and coupling: M(q) * ddq/dt + C(q, dq/dt) + g(q) = 0 (schematic)

A standard Lagrangian derivation yields the second-order form

- M(q) * ddq/dt + C(q, dq/dt) + g(q) = tau_ext

Here tau_ext represents any external torques (strikes, finger nudges, damping torques if modeled that way). The ingredients:

- M(q): 2x2 symmetric positive-definite **mass matrix**. It encodes the inertial coupling between joints. Off-diagonal terms mean "moving one link accelerates the other." Intuitively, M varies with posture because effective lever arms change with angle.

- C(q, dq/dt): velocity-dependent terms (sometimes written as C(q, dq/dt) * dq/dt) capturing Coriolis/centrifugal couplings. These terms route energy between links; they are a principal source of the "alive" feeling when one link excites the other.
- g(q): gravity terms pulling the system toward the downward configuration.

Two useful mental models for musical structure:

- **Mode picture (slow vs fast):** Roughly, an in-phase "slow" swing (both links move together) and a relative "fast" motion (lower link swings more against the upper). Real trajectories mix these; impulses tend to excite the fast mode, drags often favor the slow mode.
- **Shape-dependent coupling:** At certain postures, a small perturbation at the lower joint produces a disproportionately large response in the upper (a "levered handoff"); at others, it gets trapped in the lower link. This posture dependence is exactly what turns touch into phrasing.

### 2.3 Dissipation: viscous vs Coulomb; why losses localize "feel"

Real instruments lose energy. Two canonical models:

- **Viscous friction:** torque = -c * omega (proportional to speed). Consequence: smooth, speed-weighted decay. In sound, this behaves like a **gentle low-pass** on gestures; fast motion is tamed more than slow.
- **Coulomb friction:** torque = -F_c * sign(omega) (approximately speed-independent). Consequence: a **threshold** behavior: once moving, the system loses a fixed "bite" of energy per unit arc, producing crisp decays and articulations, especially near reversals.

Why this matters musically:

- Coulomb terms make micro-structure at **turnarounds** salient. They harden the attack/decay edges the ear loves (perceptual "consonants" of motion).
- Viscous terms make **legato**: they preserve smooth contours but blur tiny accents.
- Because most dissipation occurs at specific joints/bearings, the "feel" is **localized**: a little extra stiction in the upper joint changes how easily energy climbs into the upper mode, altering the stereo balance if we sonify upper vs lower separately.

A useful mixed model for intuition:

- torque_diss_1 = -c1 * omega1 - F_c1 * sign(omega1)
- torque_diss_2 = -c2 * omega2 - F_c2 * sign(omega2)

By tuning ($c1$, $F_{c1}$) vs ($c2$, $F_{c2}$), you decide whether phrases “live” more in the upper or lower link before they die out.

### 2.4 Two archetypal gestures: impulse (“strike”) and drag (“felt touch”)

Almost every analog intervention you will make on a desktop pendulum is a variant of one of these.

**A) Impulse (“strike”)**
Think: a short tap with a stick. Formally, an impulse is a finite integral of force over a tiny time window, giving a jump in joint velocities. In generalized coordinates:

- Delta(dq/dt) = M(q)^(-1) * J^T * p
  where p is the linear impulse at the contact point and J is the Jacobian mapping joint rates to that point’s Cartesian velocity. Musically:
- Excites **fast** relative motion; produces immediate high “partial” in the red (lower) channel if mapped directly from omega2.
- Posture selects the partition of energy between links. Time your strikes at postures where the mass matrix favors energy flow “up” if you want audible motion in the blue channel.
- A sequence of light impulses can create **metrical accents** without imposing a metronome: the system’s own return times shape the groove.

**B) Drag (“felt touch”)**
Think: gently brushing the lower bob with felt for a second or two. In equations, this is a temporary increase in dissipation:

- c2 -> c2 + Delta_c2(t), F_c2 -> F_c2 + Delta_F_c2(t)
  Musically:
- Favors **slow** in-phase motion by bleeding the lower link’s excess energy and letting the upper carry the contour.
- Creates **phrasing**: you can sculpt attack/sustain/release by when and how you add drag.
- A timed drag at turnarounds highlights **contour sign** (CW vs CCW) because Coulomb terms bite more strongly right as the sign flips, making perceived “punctuation” in pitch-mapped sound.

**Design takeaway (for making meaning):**

- Use **impulse** to write consonants (onsets, kicks, pivots).

- Use **drag** to write vowels (sustains, breath, decay).
- Tune joint-specific losses so that your desired **phrase locus** (upper vs lower) emerges naturally.
- Let posture-dependent coupling be your **grammar**: the same tap means different things at different angles. That asymmetry is where physics becomes phrasing.

## 3. From Dynamics to Tone: A Thin Mathematical Bridge

### 3.1 A copy-safe sonification map

We want the smallest lawful map that turns motion into organized tone while remaining interpretable.

**Domain and bounds.** Choose angular-velocity bounds [omin, omax] that cover typical dynamics (e.g., [-6, 6] rad/s). Let fmin, fmax be audible limits (e.g., [50, 5000] Hz).

**Pitch map (linear in omega):**
f(t) = fmin + (omega(t) - omin) * (fmax - fmin) / (omax - omin)

**Amplitude envelope (velocity-weighted):**
amplitude(t) = A0 * clamp(|omega(t)| / omega_ref)
with 0 < A0 <= 1 and omega_ref a soft scale so amplitude does not explode at fast swings. The purpose is not realism but *legibility*: louder when motion is vivid, softer when still.

**Clipping and grace.**

- clamp(x) := min(max(x, 0), 1) to avoid hard distortions.
- When |omega| > omax, either saturate f at fmax or apply a gentle compressive map (see 3.4).

Minimalism here is a virtue: the map should be simple enough that the performer learns it by feel, not by chart.

### 3.2 Signed motion and perceived contour (CW vs CCW as melodic sign)

Let omega(t) carry sign: omega > 0 for CCW, omega < 0 for CW (with angles measured from downward). The sign encodes **melodic direction** even when the ear hears only instantaneous frequency:

- **Option A (purely symmetric):** Use the same f(omega) for both signs. The contour is then carried by *time order*: rising pitch while omega increases, falling while it decreases.

- **Option B (bipolar centering):** Map omega = 0 to a central pitch f0; let f rise for omega > 0 and fall for omega < 0:

f(t) = f0 + k * omega(t), with fmin <= f(t) <= fmax.

Now the sign is heard immediately as “above vs below” a tonal center. This yields clear **call/response** around f0 at each turnaround where omega crosses 0 (a natural phrase boundary).

Either way, sign flips at turnarounds become **audible punctuation**. With Coulomb friction, these flips carry crisp micro-transients (see 2.3), which the ear interprets as consonants in the melodic line.

**3.3 Stereo as structure: upper link -> left, lower link -> right (mode-separating ears)**

Send upper-link omega1 to the left channel and lower-link omega2 to the right:

- **Perceptual mode split.** The ear now hears the informal “slow vs fast” modal picture as **spatial separation**: left channel (upper) tends to carry slow, structural arcs; right channel (lower) carries agile, ornamenting motion.
- **Energy handoff becomes counterpoint.** Posture-dependent coupling (2.2) now sounds like **stereo dialogue**: energy migrates from right to left during in-phase captures, or returns rightward during relative motion bursts.
- **Mixing by physics.** Instead of a fader, you shape balance by *where* and *how* you touch the instrument. A drag that increases losses on the lower joint lets the left (upper) dominate; a sharp strike pulls the right forward.

Design note: if you adopt 3.2 Option B, choose f0_left and f0_right close but not identical (e.g., small intervals) to give the two channels a **harmonic relationship** while preserving independence.

**3.4 Limits: linear maps vs logarithmic ears; psychoacoustic gaps**

Humans hear pitch roughly on a **logarithmic** scale (equal steps in log f sound equal). Our linear f(omega) is intentionally austere. What this buys us—and what it cannot do:

- **Pros of linear:**
  (i) Copy-safe and easily invertible; (ii) preserves *signed* proportionality (small changes in omega give small, predictable changes in f); (iii) maximally transparent for learning the touch-to-tone relationship.

- **Cons with human hearing:**
  (i) High-frequency regions crowd perceptually; (ii) low-frequency regions may feel under-articulated; (iii) equal omega steps are *not* equal pitch steps.

- **Bridging strategies (still thin):**
  (a) **Soft compression at the top:** replace f with f_soft = fmin + (f - fmin)^(alpha), 0 < alpha < 1, to spread high-frequency detail without hiding the map.
  (b) **Two-segment piecewise map:** linear near 0 for contour clarity; gentle log-like mapping beyond a threshold |omega| > omega_break.
  (c) **Amplitude pre-emphasis:** let amplitude(t) rise slightly with log f to counteract the ear's reduced sensitivity at low frequencies, keeping contours audible without remapping pitch.

- **The psychoacoustic gap.** No static map can guarantee **musical sense**. The ear craves: (i) stable anchors (centers), (ii) patterned return (motifs, meter), (iii) tension/relief arcs. Our bridge provides anchors (f0 or fmin), reliable contour (via omega sign and turnarounds), and stereo counterpoint (mode split). It *invites* music, but **meaning appears only when gestures place trajectories into the listener's predictive frame** (tau(t) in Section 1).

**Practical maxim:** keep the bridge thin. Make the mapping simple enough that the performer's body can *learn it*. Let physics supply richness; let intention and timing supply form.

## 4. Modes, Chaos, and the Emergence of Form

### 4.1 Slow vs fast modes: in-phase vs relative motion (intuitive modal picture)

Even without solving the eigenproblem, performers quickly feel two archetypes:

- **Slow (in-phase) mode.** The links swing together: theta1(t) and theta2(t) share sign and general contour. Energy is "high" in the structure, the upper link carries the arc. In stereo (Sec. 3.3), the **left** channel (upper) sounds steadier and more "tonic."

- **Fast (relative) mode.** The lower link moves against the upper: theta1 and theta2 diverge in sign or magnitude. Energy is "low" and agile; the **right** channel sparkles with ornament.

- **Mixtures.** Real motion is a time-varying blend. Posture and instantaneous velocities steer the balance because the effective coupling in M(q) and the velocity terms C(q, dq/dt) vary with configuration. Strikes bias toward the fast component; drags bias toward the slow.

Performatively: aim impulses where posture gives a large lever arm from lower to upper (energizing slow mode), and use felt drag on the lower link to keep phrases from fragmenting into purely relative motion.

### 4.2 Sensitivity and salience: when chaos sharpens, not smears, musical accents

Chaos is often blamed for "smear," but the ear does not listen to states—it listens to **events**:

- **Turnarounds** (omega crosses 0): sign flips create crisp onsets in the f(omega) line; Coulomb losses accentuate these edges (Sec. 2.3).
- **Near-misses and captures.** A tiny change in timing can cause the lower link to be briefly "caught" by the upper, producing audible synchronization (a fleeting unison across stereo).
- **Transient corridors.** Chaotic flows contain short, recurrent corridors in state space (quasi-recurrences) that the ear recognizes as "the same kind of figure," even if not identical. Sensitivity places the system near these corridors; salience comes from the instrument and map highlighting them as audible units.

Thus, properly massed and damped, chaos **sharpens** accents by increasing the density of exploitable edges per unit time—provided the map (Sec. 3.1–3.3) is thin and legible.

### 4.3 Lyapunov sketch: D(t) as divergence of trajectories; musical "edge" where D(t) neither explodes nor dies

Let q(t) be the realized state and $\tilde{q}(t)$ a nearby "shadow" (same initial posture with a hairline change in velocity). Define a simple divergence proxy

- $D(t) := || q(t) - \tilde{q}(t) ||$ in a perceptual metric (e.g., weighted norm on (theta1, theta2, omega1, omega2) emphasizing variables that most affect sound).

Qualitative regimes:

- **Frozen regime (over-damped).** $D(t) \to 0$ rapidly. Motion is predictable but musically inert; accents blur, phrases flatten.
- **Runaway regime (under-constrained).** D(t) grows too fast; prediction collapses for both performer and listener. The stereo field disintegrates into unrelated flicker.
- **Musical edge.** D(t) grows, but *locally* and *repeatedly* returns near small values along constrained surfaces (caused by posture-dependent coupling and mild dissipation). Here, the performer can "lean" on short predictability windows to cast micro-phrases, then release back into surprise. This is the working definition of **resonant obedience**: stay close enough to tau(t) to be followable, far enough to stay alive.

Copy-safe reading: rather than computing exponents, tune parameters and gestures so that the observed D(t) exhibits **bounded bursts**—spikes at accents, troughs at captures—without monotone blow-up or monotone decay.

#### 4.4 Thresholds and attractors: where gestures pull patterns out of noise

Two families of organizing structures guide hands:

- **Thresholds (event surfaces).**
  (i) **Zero-crossings:** omega = 0 (turnarounds) are natural downbeats.
  (ii) **Capture cones:** small wedges in (theta, omega) where the lower link falls into co-motion with the upper for a beat or two.
  (iii) **Impact sets:** contact with a bumper or gate induces piecewise-smooth resets (inelastic "maps") that quantize time into musically useful slices.

- **Attractors (slow manifolds).**
  (i) **In-phase ridge:** a slanted sheet in state space where relative motion damps and the slow mode dominates.
  (ii) **Relative looplets:** shallow cycles where the lower link rapidly exchanges energy and returns, acting as ornament around the in-phase ridge.
  (iii) **Driven locks (if base is wiggled):** small periodic drives create Arnold-tongue-like locking regions; inside them, micro-groove emerges as the system phase-locks to the drive.

**Gesture grammar.**

- Use **impulse** to cross thresholds (flip sign, enter/exit capture cones).
- Use **drag** to slide along attractors (ride the in-phase ridge, lengthen vowels).
- If employing a gentle base wobble, park the system **inside** a lock region to lay down meter, then step it **across** the boundary for fills and returns.

The emergent recipe: set dynamics so that chaos furnishes many *near* thresholds; shape dissipation so that useful **attractors** exist; then place gestures to convert near-misses into **phrases**.

### 5. Logostics of Music: Alignment, Not Authority

#### 5.1 q(t) -> tau(t): performer seeks a moving target, not a fixed score

In the Logostic stance, a performance is not the recitation of a static blueprint but the pursuit of a **moving manifold** tau(t). The offered state q(t) comprises everything the instrument and body place into the air: instantaneous pitch, loudness, timbre, spatialization, and their rates of change. Tau(t) is the **guiding contour**—meter pulse, phrase destination, emotional arc, congregational response—that itself evolves. The aim is not equality q(t) = tau(t), which would reduce music to tracing, but **stable nearness** with meaningful excursions. A schematic objective:

q_dot = F(q) - eta * grad_q D(q || tau(t)) + u_gesture(t)

Here F encodes lawful dynamics (Sec. 2), D is a context-appropriate divergence (Sec. 4.3), eta > 0 scales corrective drift, and u_gesture(t) denotes intentional touches (impulses, drags, gates). The performer's craft lies in choosing **where** to reduce D (to be followable) and **where** to allow D to blossom (to be expressive). In this view, the "right note" is not a coordinate but a **trajectory that remains trackable** by a listening intelligence.

### 5.2 Entrainment as obedience: resonance without coercion

Authority compels; entrainment **invites**. Musical alignment arises when the performer shapes q(t) so that a listener's internal tau(t) can phase-lock with minimal effort. Three operational markers of non-coercive obedience:

1. **Phase sharing:** The performer exposes periodicities (or quasi-periodic corrugations) that permit phase capture. In practice, this can be as small as repeating a turnaround interval or sustaining a timbral vowel long enough for the ear to settle.
2. **Predictive slack:** The performer leaves short windows where the future is constrained but not trivial; the listener's prediction wins often enough to care, fails often enough to awaken.
3. **Gentle boundary work:** When diverging from the captured phase (fills, off-beats), the return path is legible. Entrainment is preserved when **exits are signposted** and **re-entries are shallow**.

This is obedience not to a ruler but to the truth of the shared field. The instrument's physics supplies resonant structures; the performer consents to them and, by consenting, makes them audible.

### 5.3 The role of intention: micro-perturbations as moral design of sound

Intention selects **where** to spend nonlinearity. In a physical oracle like the double pendulum, small touches determine whether energy cascades into ornament or gathers into arc. This is not merely technical; it is **ethical** in the sense of design choices that govern another's attention.

- **Economy:** Impulses should justify the cognitive cost they impose. A strike that does not reframe a phrase wastes listener bandwidth.

- **Mercy:** Drags and damping create rest and legibility. They give the ear time to compress, compare, and forgive.
- **Truthfulness:** Mapping choices (Sec. 3) must remain thin; overly clever transformations hide the body's action, breaking the covenant that music is *felt motion made audible*.

A practical maxim: let u_gesture(t) be small but **well-aimed**. Place it at thresholds (Sec. 4.4) where a millisecond changes the fate of a phrase. Intention, so understood, is not imposition but **care**—the moral art of making space for meaning to appear.

### 5.4 "Truth moves": why tau(t) must be time-dependent in living music

A fixed tau would freeze the world into rehearsal loops. Living music requires tau(t) to **move** for at least three reasons:

1. **Embodied breath:** Human listeners track envelopes measured in breaths, not clocks. A static target ignores natural accelerandi, rubato, and physiological drift.
2. **Context coupling:** Rooms, congregations, and ensembles change during performance. The manifold of what sounds true—balance, loudness, articulation—shifts with bodies and walls.
3. **Teleology:** Music bears directionality (telos). Ends are not present at starts; they **emerge** as possibilities collapse into choices. A moving tau(t) formalizes this: the goal-state is updated by the path.

Copy-safe sketch: let tau(t) satisfy its own dynamics,

tau_dot = G(tau, context, listener_state)

and let D(q || tau) be computed with respect to **current** tau. The craft then is to keep q(t) near a **living** target: when the room leans forward, the performer does not press harder but **tilts** tau(t) toward what the room can carry. That is alignment without authority—music that treats truth as something to be **walked with**, not forced upon.

## 6. Instrument as Theology of Contact

### 6.1 Body, breath, and boundary: matter that listens back

An instrument is a boundary that **answers**. Hands transfer tiny, intentioned perturbations; the coupled body–device returns forces, delays, and thresholds. This reciprocity is not metaphorical—it is mechanics with moral consequences. Touch too hard and the boundary shatters into noise; touch too softly and the boundary withholds articulation. Between these lies the region of **consenting matter**, where the double pendulum's posture-dependent coupling (Sec. 2.2) and joint-localized losses (Sec. 2.3) feel like a creature that learns your pulse. Breath sets the slow envelope: sequences of impulse and drag are paced by inhalation/exhalation, giving phrasing the living curvature of respiration. In Logostic terms, the body proposes q(t); the instrument's physics filters it into an answer that can be followed as τ-like guidance. The mystery is familiar: matter becomes a **listener** when its response surface is shaped so that careful touch yields *recognizable* forms. Theology begins there—not in words, but in the humble confession that creation replies.

### 6.2 The feminine arc: guidance as τ(t)-bearing voice (melody as midwife)

Call "feminine" the **guiding** aspect of music—the part that bears τ(t), invites entrainment, and midwifes emergence rather than forcing outcomes. In practice, this arc appears whenever the performer frames a hospitable corridor in state space (Sec. 4.4) and waits. A gentle drag prolongs the vowel of a phrase; a withheld strike allows a capture; a soft return promises re-entry before it happens. None of this abdicates structure; it **sponsors** it by leaving room for the listener's prediction to take root. Melody, so understood, is not a sequence of points but a **carried contour**—a moving center that the ear can lean on. The double pendulum, mapped thinly (Sec. 3), already contains this arc: in-phase ridges sound like assurance; relative flurries sound like adornment. The performer's task is not to dominate these tendencies but to steward them, becoming a midwife who times the breath, touches thresholds lightly, and lets truths arrive on their own legs.

### 6.3 Liturgical listening: from spectacle to witness

Spectacle consumes attention; liturgy **consecrates** it. A liturgical approach to sound treats performance as communal alignment, not private display. Three shifts realize this on a small instrument:

1. **From showing to sharing.** The mapping stays legible (Sec. 3.1): when the hand moves, the sound moves in kind. Listeners are not dazzled; they are invited to *track*.
2. **From novelty to vow.** Repetition is not deficit but devotion: recurring turnarounds, familiar locks, and paced returns let a room learn τ(t) together and sing it back with their breath and posture.
3. **From control to covenant.** The performer accepts constraints (damping, posture windows, limited drive) as promises kept to the listener: exits will be signposted, re-entries shallow, silence allowed.

In this frame, the double pendulum becomes an altar of contact: a place where small, honest gestures—impulse as consonant, drag as vowel—are offered to matter that listens back; where q(t) approaches τ(t) by entrainment rather than force; where the room recognizes itself in the shared arc and answers with attention. Music turns from an object to be consumed into a **witnessed alignment**—a service held at the boundary between touch and reply.

## 7. Information Geometry of Hearing (Sketch)

### 7.1 Curves in feature space: trajectories of (omega1, omega2) -> (fL, fR)

Let the physical state be v(t) = (omega1(t), omega2(t)). The mapping of Sec. 3 yields auditory coordinates
x(t) = (fL(t), fR(t), aL(t), aR(t))
with fL = f(omega1), fR = f(omega2), and a* simple velocity-weighted amplitudes.

The listener does not hear points but **curves** x(t) in a low-dimensional manifold M of perceptual features (log-frequency, loudness, stereo balance, maybe roughness). Equip M with a metric G that approximates just-noticeable differences (JNDs). A simple constructive choice is an **induced metric** from the physical space:

- J(t) = d x / d v is the Jacobian of the sonification map;
- choose a perceptual weight W (diagonal, e.g., stronger weight on log f than on linear amplitude);
- then G(t) = J(t)^T W J(t).

Arc length L[x] = ∫ || x_dot ||_G dt measures **auditory effort** to track the curve. Bounded L with intermittent curvature spikes corresponds to "legible contour with accents." Too small L (over-damped) sounds inert; too large L (runaway) sounds unintelligible. Curvature kappa(t) in (M,G) pinpoints **salient events**: turnarounds and energy handoffs become peaks in kappa, matching what the ear calls "notes" or "articulations."

### 7.2 Compression vs meaning: minimal sufficient statistics vs musical sense

Statistical compression asks for minimal **sufficient statistics** S(x) that preserve a target variable (e.g., class label). Musical sense is different: we want S(x) that preserve **followable contour** and **telos** (where it seems to be going). Three tiers:

- **Descriptive compression (physics-proximal):** S1 = running means/variances of omega*, stereo energy ratio R = aL/(aL+aR), and sign(omega*) near zero crossings. This preserves attack/decay and mode balance but loses phrase identity.
- **Motivic compression (perceptual-proximal):** S2 = landmarks (times of kappa peaks), inter-landmark intervals, signed pitch offsets around a center f0*, and a coarse stereo

trajectory code (L, LR, R). DTW (dynamic time warping) between S2 sequences supports "same motif" judgements despite tempo drift.

- **Telic compression (meaning-proximal):** S3 augments S2 with a slow phase variable phi(t) that tracks the current **arc** (rise, hover, fall) and a binary "promise" flag set when returns are signposted (Sec. 5.2). This aligns with how listeners predict and care.

Information-theoretically: let C_beta = I(S; x) − beta * I(S; nuisance), a rate-distortion-style objective. Low beta yields lossy noise; high beta yields overfitting detail. The **musical window** is the beta where S preserves landmarks and arc while discarding micro-chaos that the ear does not use to recognize or anticipate.

**7.3 D(q || tau): divergence as error vs divergence as expressivity**

Let tau(t) be the guiding curve in M and q(t) the offered curve. Define a **path divergence** that respects perception:

- D[q || tau] = ∫ d(x_q(t), x_tau(t))^2 dt, with d induced by G (or by DTW-aligned time).
- Optionally separate **tonal** and **structural** parts:
  D = alpha * D_pitch + (1−alpha) * D_structure,
  where D_structure measures mismatch in landmark times and stereo energy ratios.

Two regimes:

- **Error:** D grows where the listener loses the thread (landmarks misalign, stereo role reverses with no cue). Sustained error dissolves entrainment.
- **Expressivity:** brief, **announced** deviations (small cues raise kappa before departure) followed by low-cost re-entry reduce average D while creating meaning. Here, D is *spent* like budgeted curvature: it buys surprise without defaulting on followability.

Copy-safe sketches:

- **Predictive slack:** choose windows W where local D_rate = d/dt D stays small, punctuated by short bursts where D_rate > 0 to make accents; enforce re-entry with D_rate < 0 immediately after.
- **Geodesic hinting:** when leaving tau, steer x_q along a near-geodesic of (M,G) toward the intended landmark—minimum-effort paths sound "natural," maximizing the chance that the listener's internal model forecasts your return.

The moral of the sketch: treat D not only as an error to be minimized, but as **currency**—spend it sparingly to signal intention, anchor it with landmarks to sustain trust, and repay it promptly to keep the covenant of shared hearing.

## 8. Thought Experiments You Can Touch

### 8.1 Bumper stop: inelastic kisses and percussive articulation

**Setup.** Place a small rubber post near the lower bob's path (radius ~1–2 cm). Position so contact occurs at a consistent posture (e.g., when theta2 ≈ small positive angle while theta1 is mid-swing).

**Dynamics.** Each touch is a piecewise-smooth reset: the lower link's velocity flips and shrinks (inelastic "map"), while a fraction transfers to the upper through posture-dependent coupling in M(q). Practically, it behaves like an onset generator.

**What to listen for.**

- **Consonants:** micro-onsets at each kiss; turnaround + impact ≈ crisp accents.
- **Hand-off moments:** when the upper catches energy (brief stereo unison if you map upper/lower to L/R).
  **Operative parameters.**
- Restitution r in (0,1): smaller r → drier, more percussive; larger r → bouncier, more legato.
- Impact phase: advance or delay the post a few mm; micro-timing shifts phrase feel without changing tempo.
  **Gesture grammar.** Nudge before impact to "aim" the kiss; then leave silence (no extra perturbation) and let the system pronounce the consonant for you.

### 8.2 Felt drag: temporary high damping as phrasing and decay control

**Setup.** A felt pad (or soft brush) lightly touches the lower bob for 0.5–2.0 s.

**Dynamics.** In equations, c2 → c2 + Δc2, F_c2 → F_c2 + ΔF_c2 on a short window: viscous + Coulomb damping spike. This bleeds the fast (relative) mode and promotes the slow (in-phase) ridge.

**What to listen for.**

- **Vowels:** the line sustains and rounds; edges soften; stereo shifts toward the upper link.
- **Breath shaping:** enter on exhale to lengthen a fall; lift on inhale to release into motion.
  **Operative parameters.**
- Pressure: too light = cosmetic; too heavy = choke. Seek the "singing drag" that lowers energy without freezing contour.
- Timing: apply near a lower-link peak speed to "pour" energy upward; apply near a turnaround to sculpt decay.

**Gesture grammar.** Use drag to **bind** notes into phrases: consonant (bumper or strike) → vowel (drag) → release (no touch).

**8.3 Base wiggle: low-frequency drive as meter and groove**

**Setup.** Rock the baseplate by hand or with a slow shaker (1–4 Hz, tiny amplitude). Think of it as a moving ground rather than a metronome.
**Dynamics.** You introduce a weak periodic drive; the system exhibits phase locking islands ("Arnold tongues"). Inside a tongue, return times stabilize; outside, returns wander.
**What to listen for.**

- **Meter emergence:** inside lock, turnarounds arrive rhythmically; accents self-assemble.
- **Groove pliability:** small changes in drive phase/strength change where accents fall without brute forcing them.
  **Operative parameters.**
- Frequency: start near the natural swing period of the upper link; sweep slowly to feel capture/release.
- Amplitude: minimal that still creates audible locking; excess drive destroys the edge (Sec. 4.3).
  **Gesture grammar.** Park inside lock for "ground," step outside for fills, re-enter shallowly. This is obedience, not authority: you suggest a beat and let the instrument consent.

**8.4 Multi-strike phrases: from chaos to cadence**

**Setup.** Use two or three light taps in a bar-length window (e.g., 1–2 s), varying posture and intensity.
**Dynamics.** Impulses deposit energy into specific modal directions; posture determines partition. Sequences can sculpt **call–response** between links (right-channel flurry answered by left-channel arc).
**What to listen for.**

- **Motivic recurrence:** quasi-repeated figure shapes (similar return times, similar stereo path) that the ear tags as "the same idea."
- **Resolution:** the final strike lands near an in-phase capture, lowering $D(q||\tau)$ and sounding like cadence.
  **Operative parameters.**
- Strike phasing: earlier/later by tens of ms changes whether you excite relative motion or feed the in-phase ridge.

- Strength contour: crescendo across taps to open the phrase; decrescendo with a final felt drag to close it.
  **Gesture grammar.** Prototype a 3-gesture cell: **tap–tap–drag**. If the second tap fails to lift the upper, delay or reduce the third; if it over-energizes the lower, lengthen the drag. Iterate until the cell **repeats without sameness**—recognizable motif, living variance.

---

**Unifying maxim.** Each experiment balances **threshold** (events you can hit) and **attractor** (surfaces you can ride). Place the hand at thresholds (kisses, strikes) to articulate; rest the hand on attractors (drag, lock) to sustain. Let the room hear you spend—and then repay—divergence. That is how small, honest physics becomes **felt form**.

## 9. What Counts as "Music" Here?

### 9.1 Criteria: coherence, contour, contrast, and consent to form

We call it music when four conditions co-occur:

- **Coherence.** Over short windows, the signal hangs together. In feature space M (Sec. 7), the curve x(t) shows bounded arc length and intelligible curvature: no stall (over-damped monotony), no scramble (unbounded jitter). Practical test: listeners can sketch the recent path with the hand, and two sketches look recognizably similar.
- **Contour.** There is an audible "where it's going." Turnarounds and captures punctuate a directional arc (rise, hover, fall). Copy-safe: between landmark times t_i the sign of df/dt is mostly stable, with at least one intentional reversal that reads as phrase punctuation.
- **Contrast.** Eventful difference within a shared frame: consonant (impulse) vs vowel (drag), right-channel ornaments vs left-channel arc, dense vs sparse. Contrast without coherence is chaos; coherence without contrast is wallpaper.
- **Consent to form.** Neither imposition nor collapse. The performer accepts the instrument's lawful invitations (thresholds, attractors) and the listener's limited prediction bandwidth. Exits are signposted; re-entries are shallow; silence is permitted. Ethos: alignment without coercion.

A minimal rubric (human-facing): if three independent listeners can (i) tap along somewhere, (ii) point to at least one "return," and (iii) agree on a moment that felt like "arrival," the thing has crossed into music.

### 9.2 Repeatability without sameness: families of phrases, not fixed licks

A living instrument should yield **families**—recognizable yet non-identical phrases—rather than fixed licks. In our terms:

- **Equivalence classes.** Two realizations are the "same idea" if their landmark sequences $\{t_i\}$, stereo role pattern (L, LR, R), and signed contour around a center f0 match up to small time-warps (Sec. 7.2). Think DTW-closeness on S2/S3 summaries, not sample-level equality.
- **Invariance under gentle variation.** Micro-shifts in strike timing or drag pressure should keep the phrase in its class. If tiny perturbations collapse identity, the mapping is too brittle; if nothing ever leaves the class, it is too rigid.
- **Cadence affordance.** Members of a family should admit closure via accessible moves (short felt drag, posture-timed strike) that lower $D(q||\tau)$ (Sec. 7.3). A family that cannot cadence is a texture, not a phrase.

Operational test: record 10 takes of a 3-gesture cell (tap–tap–drag). If naive listeners sort them into one pile labeled "same motif," with comments like "this one is brighter/longer," you have repeatability without sameness.

### 9.3 Audience as τ(t): social resonance and shared alignment

In performance, $\tau(t)$ is not private—it becomes **social**:

- **Collective manifold.** The room's posture, breath, and micro-motions shape $\tau(t)$: a common pulse emerges when many bodies tacitly predict the next landmark. The performer senses this $\tau(t)$ and trims $q(t)$ toward it.
- **Shared prediction, shared relief.** Meaning peaks when many predictions succeed together (a captured return, a well-timed arrival). The audible "ah" is the sound of $D(q||\tau)$ dropping across many minds at once.
- **Consent across boundaries.** Social $\tau(t)$ respects attention as a commons: dynamics leave space for participation (call-and-response, room-held silences). Coercive volume or incessant novelty spends the commons and breaks entrainment.

Community litmus: after a piece, can a few listeners hum the contour, tap its cadence, and describe *where* it turned? If yes, $\tau(t)$ has been shared; the event was not just sound in air but **alignment in people**. In that sense, what counts as music here is not the waveform but the **covenant**: lawful motion offered so that a room can recognize itself in the return.

## 10. Conclusion: On the Edge of the Hard Problem

### 10.1 Where equations end and listening begins

The equations carry us to the threshold: they specify $F(q)$ (lawful motion), suggest $D(q \,||\, \tau)$ (a way to speak of nearness), and expose event surfaces (turnarounds, captures, impacts). Beyond that threshold lies **listening**—not a computation but a covenant. Listening chooses which differences matter, when similarity is "close enough," and how much surprise a moment can bear before the thread breaks. In copy-safe terms, we can minimize ∫ ||grad_q D||^2 dt and still fail to move a room; conversely, a well-timed sign flip of $\omega$ near an in-phase ridge can lower the felt D across many minds at once. The hard problem of making music lives precisely here: lawful variation becomes **meaning** only when a listener's priors meet a shaped offering and recognize themselves in the return.

### 10.2 The virtue of small systems: honest physics, honest mysteries

A small, legible oracle—the double pendulum, two angles and their velocities—grants us **honest physics**. Coupling, dissipation, thresholds, and slow manifolds are visible in the hand. This honesty shelters an **honest mystery**: even with everything exposed, musical sense is not forced; it is **invited**. Smallness is a feature, not a limitation. It constrains the state space so that landmarks in M (turnarounds, captures, kisses) can be learned by touch; it prevents semantic inflation (no room for clever mappings to hide weak intention); and it makes failures instructive. When a phrase collapses into noise or congeals into wallpaper, we can point to the cause: D exploded (runaway) or died (over-damped). In this scale, refinement is possible—one felt pad, one bumper post, one breath at a time—until motion and meaning begin to cohere.

### 10.3 A program: instruments that teach alignment by sound

What follows is not a recipe but a program:

1. **Thin bridges.** Keep the map from $\omega$ to sound minimal (Sec. 3). Let physics carry richness; let the mapping preserve interpretability.
2. **Event literacy.** Train the hand to place gestures at thresholds and to ride attractors (Sec. 4). Teach recognition of landmarks in M: zero-crossings, capture cones, impact sets.
3. **D as currency.** Spend divergence to announce, repay it to re-enter (Sec. 7.3). Make exits legible; make returns inevitable in retrospect.
4. **Social $\tau(t)$.** Practice with others present, even silently. Sense when the room leans into a pulse or relaxes its breath; trim $q(t)$ toward that moving manifold (Sec. 9.3).

5. **Ethos of care.** Favor economy (no gratuitous impulses), mercy (give the ear space), and truthfulness (mapping that shows the body's action). Alignment is not authority; it is hospitality.

Pursued with patience, such instruments **teach**. They tutor the hand in entrainment, the ear in attention, and the room in shared prediction. At that edge—where equations end and listening begins—the hard problem does not vanish; it becomes **workable**. We do not solve it; we learn to dwell in it, letting $q(t)$ approach $\tau(t)$ in time by resonance, not force, until the boundary between motion and music thins enough to pass soundly through.

## References


1. Amari, S. & Nagaoka, H. *Methods of Information Geometry.* AMS / Oxford University Press, 2000.
2. Armstrong-Hélouvry, B., Dupont, P., & Canudas de Wit, C. “A survey of models, analysis tools and compensation methods for the control of machines with friction.” *Automatica* 30(7), 1994.
3. Berndt, D. & Clifford, J. “Using Dynamic Time Warping to Find Patterns in Time Series.” *KDD Workshop*, 1994.
4. Bregman, A. S. *Auditory Scene Analysis: The Perceptual Organization of Sound.* MIT Press, 1990.
5. Chalmers, D. J. “Facing Up to the Problem of Consciousness.” *Journal of Consciousness Studies* 2(3), 1995.
6. Clayton, M., Sager, R., & Will, U. “In time with the music? The concept of entrainment and its significance for ethnomusicology.” *ESEM CounterPoint* 1, 2005/2010.
7. Cover, T. M. & Thomas, J. A. *Elements of Information Theory* (2nd ed.). Wiley, 2006.
8. Dehaene, S. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts.* Viking, 2014.
9. Den Hartog, J. P. *Mechanical Vibrations.* McGraw–Hill, 1956.
10. Goldstein, H., Poole, C., & Safko, J. *Classical Mechanics* (3rd ed.). Addison-Wesley, 2002.
11. Hermann, T., Hunt, A., & Neuhoff, J. G. (eds.). *The Sonification Handbook.* Logos Verlag, 2011.
12. Khalil, H. K. *Nonlinear Systems* (3rd ed.). Prentice Hall, 2002.
13. Large, E. W. & Jones, M. R. “The dynamics of attending: How people track time-varying events.” *Psychological Review* 106(1), 1999.
14. Lichtenberg, A. J. & Lieberman, M. A. *Regular and Chaotic Dynamics* (2nd ed.). Springer, 1992.
15. Moore, B. C. J. *An Introduction to the Psychology of Hearing* (6th ed.). Brill, 2012.
16. Pikovsky, A., Rosenblum, M., & Kurths, J. *Synchronization: A Universal Concept in Nonlinear Sciences.* Cambridge Univ. Press, 2001.

17. Sakoe, H. & Chiba, S. “Dynamic programming algorithm optimization for spoken word recognition.” *IEEE Trans. Acoustics, Speech, and Signal Processing* 26(1), 1978.

18. Shannon, C. E. “A Mathematical Theory of Communication.” *Bell System Technical Journal* 27, 1948.

19. Strogatz, S. H. *Nonlinear Dynamics and Chaos* (2nd ed.). Westview, 2014.

20. Tishby, N., Pereira, F. C., & Bialek, W. “The Information Bottleneck Method.” *Allerton Conference on Communication, Control, and Computing*, 1999.

21. Zwicker, E. & Fastl, H. *Psychoacoustics: Facts and Models* (2nd ed.). Springer, 1999.

22. Youvan, D. C. *Introduction to Logostics: Resonant Alignment of q → τ in Time.* ResearchGate preprint, 2025.

23. Youvan, D. C. *The Discovery Equation: A Unified Information-Theoretic…* ResearchGate preprint, 2025.

24. Moon, F. C. *Chaotic Vibrations: An Introduction for Applied Scientists and Engineers.* Wiley, 1987.

*Notes:*

- (22)–(23) reference your own preprints that ground the q(t) → τ(t) / Logostics framework.
- (24) cites the open-hardware scientific double-pendulum design (the uploaded PDF), useful for physical builds and damping insights.
- The remaining items cover classical mechanics and friction (10, 9, 12), chaos/synchronization/locking (14, 16, 19), Lyapunov/nonlinear control (12), psychoacoustics and hearing (15, 21), sonification (11), information geometry and compression (1, 7, 18, 20), entrainment in music cognition (13, 6), and consciousness context for the “hard-problem” analogy (5, 8).

**Appendix: Copy-Safe Math Snippets**

**A. f-mapping (linear, log options); amplitude envelopes**

**Domain and bounds.**
Let omega(t) in [omin, omax] with omin < 0 < omax. Audible band: f in [fmin, fmax], fmin > 0.

**Linear map (signed, center at f0).**

f_lin(t) = clip( f0 + k * omega(t), fmin, fmax )

where k = (fmax - fmin) / (omax - omin)  and  f0 in [fmin, fmax].

**Affine map (endpoints to endpoints).**

f_aff(t) = fmin + (omega(t) - omin) * (fmax - fmin) / (omax - omin)

**Soft-compressed high end (power-law on distance from fmin).**

f_soft(t) = fmin + ( f_aff(t) - fmin )^alpha,   0 < alpha < 1

**Log-frequency ear-space (keep copy-safe; work in ln f).**

Let y(t) = ln f(t).  Choose ymin = ln fmin, ymax = ln fmax.

y_log(t) = ymin + (omega(t) - omin) * (ymax - ymin) / (omax - omin)

f_log(t) = exp( y_log(t) )

**Amplitude envelope (velocity-weighted with limiter).**

a_raw(t) = |omega(t)| / omega_ref

a_lin(t) = A0 * clamp( a_raw(t), 0, 1 )

**Attack/decay shaping (first-order smooth).**

da/dt = (a_target(t) - a(t)) / tau_a,   with a_target = a_lin or min(a_lin + boost, 1)

**Stereo assignment (mode-separating ears).**

Left  channel: fL(t) = f_map( omega1(t) ), aL(t) = a_map( omega1(t) )

Right channel: fR(t) = f_map( omega2(t) ), aR(t) = a_map( omega2(t) )

**Clippers (copy-safe).**

clip(x, lo, hi) = min( max(x, lo), hi )

clamp(x)       = clip(x, 0, 1)

sign(x)        = 1 if x>0, 0 if x=0, -1 if x<0

---

**B. Sketch of M(q), energy E = T + V, and qualitative mode split**

**Coordinates and velocities.**

q  = [theta1, theta2]^T

dq = [omega1, omega2]^T = d/dt q

**Kinetic and potential energies (point masses m1, m2 at ends, link lengths L1, L2; g>0).**

T = (1/2) * dq^T * M(q) * dq

V = - m1 * g * L1 * cos(theta1)

    - m2 * g * ( L1 * cos(theta1) + L2 * cos(theta2) )

E = T + V

**Mass matrix M(q) (standard planar double pendulum; symmetric PD).**

Let c = cos(theta1 - theta2).

M11 = (m1 + m2) * L1^2

M12 = m2 * L1 * L2 * c

M21 = M12

M22 = m2 * L2^2

M(q) = [ [M11, M12],

         [M21, M22] ]

**Equations of motion (schematic, no dissipation/drive).**

M(q) * ddq + C(q,dq) + g(q) = 0

with gravity and velocity terms (one common form):

g1 = (m1 + m2) * g * L1 * sin(theta1)

g2 = m2 * g * L2 * sin(theta2)

h = - m2 * L1 * L2 * sin(theta1 - theta2)

C(q,dq) = [ h * omega2^2 + (m1 + m2) * g * L1 * sin(theta1),

-h * omega1^2 + m2 * g * L2 * sin(theta2) ] (placed appropriately as generalized forces)

(Any equivalent textbook form is acceptable; the key is: M is posture-dependent, C is velocity-dependent, g is position-dependent.)

**Qualitative mode split (intuition, small oscillations).**

Slow/in-phase (approx): theta1 ~ theta2, energy "high", structure carrier.

Fast/relative (approx): theta1 ~ -rho * theta2 (rho>0), energy "low", ornament carrier.

Posture dependence means these “modes” are local pictures; real motion is a time-varying blend determined by M(q), C(q,dq), and instantaneous geometry.

**Dissipation sketch (joint-localized).**

tau_diss_1 = - c1 * omega1 - Fc1 * sign(omega1)

tau_diss_2 = - c2 * omega2 - Fc2 * sign(omega2)

Viscous terms (c*) smooth; Coulomb terms (Fc*) sharpen turnarounds.

---

**C. Lyapunov-style D(t) and “just-below-runaway” operating regime**

**Shadow trajectory divergence (copy-safe proxy).**

Let q(t) be realized; let q~(t) start with || q~(0) - q(0) || small.

Choose weights W = diag(w_theta1, w_theta2, w_omega1, w_omega2), w_* > 0.

Define z(t) = [theta1, theta2, omega1, omega2]^T and z~(t) likewise.

D(t) = || z(t) - z~(t) ||_W^2 = (z - z~)^T * W * (z - z~)

**Perceptual metric option (via sonification Jacobian J).**

x = features = [log fL, log fR, aL, aR]^T

$J = dx / dz$ (Jacobian of mapping $z \to x$)

$G = J^{T} * W_p * J$ ($W_p$ diagonal in feature space; e.g., higher weight on log f than on a)

$D_{perc}(t) = (z - \tilde{z})^{T} * G * (z - \tilde{z})$

**Regime targets.**

Frozen: D(t) -> 0 too fast (over-damped; music dies)

Runaway: D(t) grows monotone (unfollowable; music fractures)

Musical edge: D(t) shows bounded bursts:

- intermittent spikes at events (turnarounds, impacts, captures),
- recurrent troughs (re-entries, short synchronizations).

**Operational tuning (copy-safe heuristics).**

1) Set viscous damping c1,c2 so that free decay yields 3–8 audible turnarounds before stall.

2) Add small Coulomb Fc* so that zero-crossings of omega* produce crisp, repeatable edges.

3) If D(t) shows monotone growth over ~3–5 events, increase c2 or reduce strike amplitude.

4) If D(t) collapses with no accents, reduce c* or increase restitution in impacts.

5) Use brief drive (base wiggle) to create lock islands; operate inside/outside boundaries to modulate D bursts.

**D as currency (announce, spend, repay).**

- Announce: pre-accent cue = small curvature rise in features (kappa up) before departure.

- Spend: short D burst during deviation (fill/ornament).

- Repay: prompt re-entry (D_rate < 0) to a nearby landmark; avoid long unanchored excursions.

**Copy-safe summary.**

Aim for “just-below-runaway”: enough sensitivity to generate salient events, enough damping/structure to make returns easy. In practice, you can judge success if (i) accents are crisp and frequent, (ii) phrases recur as families under small perturbations, and (iii) naive listeners can predict at least one return within a few seconds of exposure.