What Can a Voice Keep?
Chord profiles, voice leading, and exact tuning decisions

Oleksiy Babanskyy

Abstract

Which features of a chord tell a player whether a particular note can last? Two progressions with identical ordered pair-sum profiles give different answers for a three-chord pedal. When parts have names and registers, an explicit four-profile family connects previously isolated assignments under fixed holds. With the same four profiles, different endpoints and no protected parts, expanding the baseline halves the minimum total motion. A constructed five-part phrase shows why a tempered pitch may have a just realization at either join but none through both. In the first eleven measures of Bruckner's *Locus iste*, a bass-relative tuning problem requires at least six changed written attacks and admits four minimizing assignments. Requiring all selected pairs to approach their targets gives a sharp worst error of a quarter syntonic comma in the continuous model, attained by a fixed meantone scale. A final 24-pulse example identifies a missing attack when the preceding rhythm is known and exactly one permitted swap occurred. Each conclusion specifies what must stay fixed and what a player is allowed to change.

1 What can one voice keep?

Imagine singing one note while three successive chords change around you. A *pedal* here means one pitch available in all three chords, so a single part can sustain it. Start with just the note names. Number the keys of the chromatic scale from C as 0 to B as 11. Adding 1 raises a note by one semitone and adding 12 returns to the same *pitch class*. An actual C4 and C5 have the same class and different registers.

Take the four-note chords $X = \{0, 1, 3, 7\}$ and $Y = \{0, 3, 4, 10\}$. For each one, play the original chord, raise its notes by one semitone, then play the original chord raised by four semitones. Reduce modulo 12 at each event. Figure 1 contains the resulting six chords. Follow one horizontal row of dots across a panel.

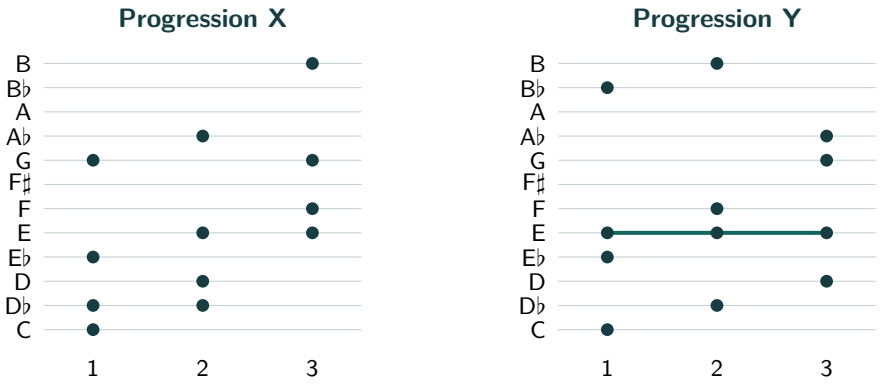

Figure 1: Each column is one chord. Read upward for pitch and from left to right for the three events. The line through E in the right panel marks a possible pedal. For X the first two chords share D-flat, the last two share E, and the first and last share G. No row spans all three.

For Y, class 4 belongs to all three chords. Choose E4 at each event and assign it to the same player. The three other parts can take the remaining notes in each chord, since this example imposes no other register or movement restrictions. For X, any sustained note would have to belong to the intersection of all three class sets, which is empty. This already rules out a pedal in every octave. Each pair of events in each progression still shares exactly one class.

Why compare these particular progressions? At each event the two chords give the same answer to a stronger counting test. Add the pitch classes of every ordered pair of notes in the chord, including a note paired with itself, reduce each sum modulo 12 and count each label. Section 2 derives the common counts and explains what else they preserve. The pedal question shows exactly where more information is needed. The transpositions are fixed at $0, 1, 4$ for both progressions. Letting each chord choose its own transposition would change the question.

Reading paths. For a score decision, start with the passing E in Section 4 or the soprano entrance over a held bass in Section 5. For the algebra behind the opening picture, continue to the chord profiles and the two proofs about named parts in Section 3. Appendix A locates each experiment and gives a first action in the existing Studio. The examples and their hypotheses can also be read directly in these pages.

2 What does a chord description preserve?

Write $C_n = \mathbb{Z}/n\mathbb{Z}$ for the pitch classes of an octave divided into n equal steps. A chord profile $f\colon C_n \to \mathbb{Z}_{\geq 0}$ counts its sounding occurrences by class. A repeated class contributes its multiplicity even if the occurrences occupy different octaves or parts. The support records which classes occur. The mass $m = \sum_x f(x)$ counts all occurrences.

Consider the labelled sum and difference profiles

$$Q_f(s) = \sum_{x \in C_n} f(x) f(s - x), \qquad \mathrm{AC}_f(d) = \sum_{x \in C_n} f(x) f(x - d).$$

Both use ordered pairs of occurrences, including the diagonal. If two different occurrences share a class, they contribute separately. The sum coordinate s is the parameter of the inversion $x \mapsto s - x$. For a set with no repeated class, $Q_f(s)$ therefore counts the tones common to the set and its image under this inversion. For multiplicities, the formula counts pairs by a product of weights. A multiset intersection would instead use the smaller of each pair of weights. A fixed inversion point of class x contributes $f(x)^2$ ordered pairs, of which $f(x)$ are diagonal. Two different classes exchanged by the inversion contribute in both orders.

Translations and inversions of the chord supply a different comparison. Under $x \mapsto x + b$, the sum label shifts by $2b$ and the difference label stays fixed. Under $x \mapsto b - x$, the sum label s goes to $2b - s$ and the difference label d goes to $-d$. These statements follow by applying the action to both occurrences in each ordered pair. Since $\mathrm{AC}_f(d) = \mathrm{AC}_f(-d)$, the difference profile is invariant under both actions. The set of all translated and inverted chords will be called the T/I orbit. Equality of Q here means equality at every label with the origin fixed.

Example 2.1 (Equal sums across corresponding chords)**.** For the two sets in Figure 1 (Section 1), add the four diagonal pairs and double each of the six unordered pairs. Both calculations give the same vector

$$(1, 2, 2, 2, 2, 0, 1, 2, 2, 0, 2, 0)$$

in labels $0, \ldots, 11$. The respective circular gap multisets are $\{1, 2, 4, 5\}$ and $\{1, 2, 3, 6\}$. A translation preserves the circular order up to rotation, and an inversion reverses it, so these different gap multisets exclude T/I equivalence. Applying the same translations $0, 1, 4$ to each set shifts the corresponding labels of Q by $0, 2, 8$. Equality therefore holds at each corresponding event. The triple intersections in Figure 1 still give different answers to the question of a common pedal.

For the comparison with autocorrelation, let $\widehat{f}(k) = \sum_x f(x)\exp(2\pi ikx/n)$. Multiplying and reindexing finite sums gives $\widehat{Q}_f(k) = \widehat{f}(k)^2$ and $\widehat{\mathrm{AC}}_f(k) = |\widehat{f}(k)|^2$. Hence equal labelled Q implies equal AC, since $\widehat{\mathrm{AC}}_f(k) = |\widehat{Q}_f(k)|$ and the transform is injective. The reverse implication can fail even for single occurrences. The one-note chords at classes 0 and 1 in C_{12} both have one difference pair at 0. Their only sum pairs occur at 0 and 2, respectively.

The arithmetic paper [1] studies nonnegative convolution-square fibres on $C_4 \times C_{p^a}$ for an odd prime p and $a \geq 1$, with labelled coordinates and a fixed origin. The application below uses one example in $C_{12} \cong C_4 \times C_3$. It needs the precise list of profiles in that example and a model for the motion of named parts. The first part is algebraic. The next section introduces the second part.

Proposition 2.2 (The four profiles in the example)**.** *Identify $H = \{0, 3, 6, 9\} \subset C_{12}$ with C_4 by $j \mapsto 3j$. For each integer $t \geq 1$, the following are all the nonnegative real profiles on C_{12} with the same labelled sum profile as A_t. They vanish outside H. Their entries in the order $0, 3, 6, 9$ are*

$$A_t = t(2, 2, 2, 0), \quad B_t = t(2, 0, 2, 2), \quad C_t = t(1, 3, 1, 1), \quad D_t = t(1, 1, 1, 3).$$

Their common Q on H is $4t^2(2, 2, 3, 2)$.

Proof. The formula for Q follows by convolving A_t with itself on H. If $g \geq 0$ has the same Q, then $g(x) > 0$ implies $Q_g(2x) \geq g(x)^2 > 0$. Thus $2x \in H$, and divisibility by 3 gives $x \in H$. On the four coordinates, the Fourier coefficients of A_t are $6t, 2ti, 2t, -2ti$ in frequencies $0, 1, 2, 3$. Equality of the squares fixes the first coefficient to $6t$ by nonnegativity. Reality pairs the signs of frequencies 1 and 3, and frequency 2 can change sign independently. The inverse transforms of $(6t, 2ti\epsilon, 2t\delta, -2ti\epsilon)$, for $\epsilon, \delta \in \{1, -1\}$, are

$$\frac{t}{2}(3 + \delta,\ 3 - \delta + 2\epsilon,\ 3 + \delta,\ 3 - \delta - 2\epsilon).$$

The four choices give exactly the four displayed nonnegative profiles, all with integral entries for integral t. □

For $t = 1$, the four profiles each contain six occurrences, yet two have three occupied classes and two have four. The move from three classes to four is possible inside the same labelled fibre precisely because repeated occurrences carry weight. This is the concrete arithmetic input to the next comparison.

3 Which voice can move?

Each of the $6t$ named parts now occupies one of four *registered* levels $0, 1, 2, 3$, separated by three semitones. A state $x = (x_v)_v$ specifies the level of every part. Fixing the level of a protected part throughout a path is called pinning it. The histogram forgets these identities and gives the counts of Proposition 2.2. Adjacency requires $x \neq y$ and $|x_v - y_v| \leq 1$ for each part. Thus no part moves more than three semitones per edge. The distance between levels is linear, with no octave wrap. Unisons, crossings and simultaneous changes are allowed. When rearrangements are allowed, an edge may preserve the histogram. When they are forbidden, an edge must change it. An edge costs $3L(x, y)$ semitones in total, where $L(x, y) = \sum_v |x_v - y_v|$.

We compare the baseline with histograms A_t, B_t against the graph also allowing C_t, D_t. The endpoints, part identities, pins, step cap and costs are the same in both graphs. Parsimonious voice-leading graphs already illustrate how additional chord types can connect components [4]. Here a vertex retains the parts through unisons and crossings, where an unlabeled chord sequence can allow more than one matching [7]. We ask how many labelled states connect and how much the same parts have to move.

Theorem 3.1 (A bridge with protected parts)**.** *Pin* $2t$ *designated parts at level 1 and impose no other pins. There are* $\binom{4t}{2t}$ *baseline states, all isolated. In the graph with all four profiles, allowing rearrangements connects all these states. Each has exactly* $\binom{2t}{t}^2$ *neighbours of profile* C_t*. With rearrangements forbidden the baseline states stay in distinct components.*

Proof. Only A_t of the baseline profiles can accommodate the pins. Its $4t$ free parts occupy levels 0 and 2, $2t$ each, so their allocation determines the count. A direct change of any free part would jump two levels. The only new admissible profile is C_t. It has t free slots at each level. The $2t$ parts from A_t's level 0 can reach only levels 0 and 1, so they fill all those free slots. The remaining free parts fill levels 2 and 3. Choosing the t parts that move up in each of the two groups proves the neighbour count. In the reverse direction the parts occupying free slots at levels 0 and 1 must return to level 0, which gives a unique A_t neighbour for each C_t state.

If rearrangements are allowed, all C_t states are connected. Temporarily distinguish the slots at each level and join slots at consecutive levels. The resulting slot graph is connected. Swaps of occupants across its edges realize arbitrary permutations, since the transposition of two endpoints of a path is obtained by swapping along the path and then back along all but its last edge. Each swap moves two free parts by one level, so it is an allowed rearrangement. Reading the resulting slot assignments as levels gives paths of C_t states. Every A_t state attaches to this connected set. When rearrangements are forbidden there are no edges within C_t. Each C_t state then attaches only to its unique A_t neighbour, so it cannot connect two different baseline states. □

Example 3.2 (Six named parts)**.** For $t = 1$, protect parts 3 and 4 at level 1. One can exchange the levels of parts 1 and 5 along the displayed path.

$$\begin{array}{ll} A_1 & (0,0,1,1,2,2) \\ C_1 & (1,0,1,1,2,3) \\ C_1 & (2,0,1,1,1,3) \\ A_1 & (2,0,1,1,0,2). \end{array}$$

The middle edge keeps the same histogram but exchanges two part assignments. Each edge costs six semitones, for a total of 18. To see optimality, observe that the endpoint variation is four levels and that every path between these distinct baseline states must visit C_1, where some part reaches level 3. Both endpoints of that part's trajectory are at most 2. Its detour adds at least two levels to the endpoint variation. A two-level cap would instead allow the direct move.

Theorem 3.3 (A cheaper path without protected parts)**.** *Use no pins and allow rearrangements. Divide the parts into six named blocks of size* t*, and write a tuple of six levels for the state in which every part of a block takes its indicated level. For the endpoints*

$$x = (0,0,2,2,3,3), \qquad y = (0,2,0,3,2,3)$$

of profile B_t*, the baseline minimum cost is* $36t$ *semitones. With all four profiles the minimum is* $18t$ *semitones. Parts in the same block may act independently along a competing path.*

Proof. Any edge from A_t to B_t carries the parts at levels $0,1,2$ to $0,2,3$, respectively. The A_t parts at zero must fill the B_t zero slots, and the A_t parts at 1 must then fill the B_t slots at 2. Consequently there is a unique A_t neighbour for each B_t state, obtained by the relabelling of levels $0 \mapsto 0,\ 2 \mapsto 1,\ 3 \mapsto 2$. Denote this projection by π and extend it by the identity on A_t. An inter-profile edge costs $4t$ levels and collapses under π. A within-B_t edge can only move parts between levels 2 and 3, so the projected edge has the same cost in levels. The A_t edges retain their cost too.

The two specified endpoints have different sets of parts at zero, which no within-B_t edge can change. A baseline path must therefore cross between the profiles at least twice. Its within-profile edges cost at least the endpoint variation of the projected path, $L(\pi(x), \pi(y)) = 4t$. The total

is at least $8t + 4t = 12t$ levels, or $36t$ semitones. The path through $\pi(x)$ and $\pi(y)$ attains this bound because their four changed blocks differ by one level each.

Every path in the expanded graph costs at least $3L(x, y) = 18t$ semitones by the triangle inequality in each part. The two intermediate tuples $(0, 1, 2, 3, 3, 3)$ and $(0, 2, 1, 3, 3, 3)$ both have profile D_t. Insert them in that order between x and y. Each edge moves two blocks by one level, making $6t$ levels in total. Both lower bounds concern all labelled paths, including paths that split the blocks. □

For the autocorrelation control, keep the baseline A_t, B_t and every other path rule fixed. The expanded comparison graph admits all registered states with histogram f satisfying $\mathrm{AC}_f = \mathrm{AC}_{A_t}$. Section 2 shows that this retains every Q-admissible vertex and may add others. The two gains still occur. In Theorem 3.1 the expanded Q graph connects all baseline states, so the same paths connect them in the AC graph. In Theorem 3.3 the AC graph can attain $18t$, and no extra vertex can lower the coordinatewise bound. These conclusions concern the headline endpoint comparisons. They leave open any equality of the full path sets or neighbour counts.

A separate control requires the parts to stay in a fixed nondecreasing order at every state. Each histogram then admits one realization. The two ordered baseline profiles, when both compatible with the pins, are adjacent at the one-level cap and attain the endpoint variation directly. If only one is compatible, it provides no pair of distinct baseline endpoints. Even a new tied optimum via C_t or D_t is impossible in the ordered comparison. The first $2t$ parts of both baseline endpoints stay at zero, while only the first t are at zero in C_t or D_t. A shortest path at the endpoint variation would have to be monotone in each part. This control changes the permitted part order, whereas forbidding rearrangements changes the edge rule.

The comparison with T/I asks another question. In a separate finite substitution bank with matched voicing and context rules, the Q-constrained minimum did not improve on the T/I minimum in any of 2,339 comparisons. This bounds the observed advantage in that bank. It gives neither an inclusion between the T/I orbit and the labelled Q fibre nor a universal cost inequality.

3.1 Giving a line to a named part

The preceding graphs fix the endpoints but let the intermediate chords vary. An arrangement can instead fix the event at which a given part must receive a note. Consider a constructed six-event phrase with six movable parts and two accompaniment parts. The request is to give the high E–F cell to part 3 at events 2 and 3, with part 1 taking a held G4 there. The exterior assignments and both accompaniment parts stay fixed. Within the cell, scheduled attacks may use the pitches in their fixed menus. Each of the three original holds of part 3 and each of the three holds of part 4 has a separate release permission. Part 2 may lower its written pitch by one octave at either inner event only with that event's permission. The register and step limits below apply to every choice. At events 1 and 4 part 1 carries the line, while at events 2 and 3 part 3 carries it. Require the line and one of parts 2, 4 and 6 to sum to 132 in registered MIDI pitch at each of those four events, using C4 as MIDI 60.

At event 1 the line is D4 (62), requiring B-flat4 (70), and at event 4 it is E4 (64), requiring A-flat4 (68). At both exterior events only part 2 can supply the note. Requiring the same partner throughout the four events therefore forces part 2 to take the inner answers to E5 (76) and F5 (77). These are A-flat3 (56) and G3 (55). The original part 2 notes are A-flat4 and G4 at those events. Two separately named octave permissions are necessary to lower them. Part 3 originally has a G4 pedal, and all three of its protected continuations into, inside and out of the cell must be released. In the middle it changes from E5 to F5, so its line transfer specifies two distinct attacks.

The following registered witness attains exactly these five concessions. The last five rows retain their original pitches. The final two columns give each part's allowed register and one-step

cap in semitones.

	0	1	2	3	4	5	Register	Cap
V_1	60	62	67	67	64	60	$[55, 67]$	10
V_2	72	70	56	55	68	72	$[55, 72]$	14
V_3	67	67	76	77	67	67	$[64, 77]$	10
V_4	65	65	65	65	65	65	$[55, 68]$	10
V_5	48	50	52	53	52	48	$[48, 53]$	4
V_6	60	58	56	55	56	60	$[55, 60]$	4
Accomp. 1	43	43	45	45	43	43	$[43, 45]$	2
Accomp. 2	60	60	59	59	60	60	$[59, 60]$	1

Part 1 holds its new G4 between events 2 and 3. The four required sums are $62 + 70, 76 + 56, 77 + 55, 64 + 68$. The displayed part motions contribute 14, 34 and 20 semitones, respectively. The unchanged parts contribute 26 and the largest one-step move is 14. The registered domains permit each displayed pitch and transition. The unchanged F4 pedal in part 4 and both accompaniment holds are preserved. Together with the forced-concession argument, this gives an inclusion-minimal set in the fixed-accompaniment high-cell query. If the partner can change from event to event, part 6 already has the two inner notes, so the three releases alone suffice. Testing only pitch classes would accept the original part 2 notes at the inner events because they differ by octaves from the required answers. This is why the arrangement decision needs a named, registered partner even after selecting a chord profile.

4 Can a note last through the whole phrase?

A held note imposes an equality across time. A listener may hear the same nominal note in two harmonies, while an arranger has to decide whether a particular player can keep one frequency through both. In just intonation (JI) a target interval is an exact frequency ratio, such as a fifth of 3/2. In a tuning with n equal divisions of the octave (EDO), a step changes frequency by a factor $2^{1/n}$. For a ratio $r > 0$ the signed difference is $1200 \log_2 r$ cents. An octave is 1,200 cents. Write

$$\kappa = \frac{81}{80}, \qquad c = 1200 \log_2 \kappa \approx 21.5063 \text{ cents}.$$

We use the syntonic comma κ as a frequency ratio, and c for its size on the logarithmic cent scale. Continuous tuning with memory and exact score constraints have different objectives. An adaptive tuner that compromises between intervals can yield an audible realization where exact held-pitch constraints have no solution [14].

4.1 A held voice and a returning chord

Suppose each chord template v gives a positive ratio a_{vi} for each named slot i, with the octave already included. The frequency of that slot is $x_v a_{vi}$, where x_v is an unknown root frequency. The template root need not be sung. Keeping slot i of u sounding through slot j of v imposes $x_v / x_u = a_{ui} / a_{vj}$. If several slots are protected on the same change, their ratios must agree. An edge with neither a protected slot nor a prescribed root ratio is free and does not supply a gain.

Proposition 4.1 (Consistency of fixed gains)**.** *On a finite graph of locally compatible changes with fixed gains in a multiplicative subgroup of the positive reals, a solution exists, with one free root per component, exactly when the oriented product on each closed walk is* 1*. Additional prescribed roots must agree with one common scaling of the propagated roots per component.*

Proof. Ratios of actual roots telescope on a closed walk. Conversely, choose a spanning tree in each component, give a tree root value 1, and propagate the gains along its edges, inverting a

gain on a reverse traversal. Each edge outside the tree closes a walk whose product is 1, so that edge also agrees with the propagated values. All solutions differ from these values by one positive factor per connected component. A prescribed root fixes the factor there, so multiple prescribed roots can be checked by comparing their required factors. This is the balance criterion for gain graphs [15]. □

For a three-voice realization of the progression C–F–Dm–G–C, with a common C reference, the following ordered slot ratios give one concrete use of the criterion. The table names each hold by the voice in which it occurs.

Template	Slots (V_1, V_2, V_3)	Hold into next template	Root gain
C	$(2, 5/2, 3)$	V_1 (C5)	4/3
F	$(3/2, 2, 5/2)$	V_2, V_3 (F5, A5)	5/6
Dm	$(2, 12/5, 3)$	V_1 (D5)	4/3
G	$(3/2, 5/2, 2)$	V_3 (G5)	2/3

For example, on F–Dm the two hold equations give $2/(12/5) = (5/2)/3 = 5/6$. Every change is locally consistent, but the product on return is 80/81. Honouring all five prescribed continuations produces an ending C chord whose frequencies are all a factor 80/81 below their opening values. If the last G5 is restarted for the return to the opening frequencies, the first three changes can stay as calculated, and the last G5 goes from 80/27 to 3 times the reference. Alternatively, quarter-comma meantone sends this product defect to identity, with different frequencies throughout. The progression and its root drift are the familiar comma-pump example discussed by Hansen [5]. The registered slot choices specify which note and voice can be changed. Closing the entire progression at a fixed pitch asks more than playing each harmony with the right internal ratios.

4.2 Moving a held E to another player

Mi di passaggio is a constructed five-part phrase lasting 3.25 seconds. In the source, A begins the E line at 0.50 seconds, when B attacks the lower G. In the arrangement, B takes the E line from 0.50 to 2.50 seconds and A takes that lower G. The transfer takes place at the written attack. P is B's designated partner. B should sound at ratios $5/4, 7/6, 3/2$ above P, respectively. All pitches in Table 1 are ratios to C4 at 264 Hz. A dash marks a rest.

Window	Time (s)	A	B	P	Q	H
W_0	0–0.50	9/8	–	3/4	7/8	1/2
W_1	0.50–1.25	3/4	$e, e\kappa$	1	7/8	1/2
W_2	1.25–1.75	3/4	$e, e\kappa$	$p, p\kappa$	$p\kappa$	1/2
W_3	1.75–2.50	3/4	$e, e\kappa$	$27/32, 5/6$	5/6	2/3
W_4	2.50–3.25	4/3	–	27/32	5/6	2/3

Table 1: Pitch menus in the transferred phrase. Here $e = 5/4$, $p = 15/14$. P's written choices at W_2, W_3 are $p\kappa$ and 27/32. The lower pitch in W_2 is already in the menu. The alternate pitch in W_3 requires permission.

For a pitch $r = 2^a 3^b 5^d 7^h$, the integer map $k(r) = 24a + 38b + 56d + 67h$ prescribes the absolute 24-EDO step $k(r)$, which sounds at $264\, 2^{k(r)/24}$ Hz. This uses the usual generator coordinates for an integer temperament map [10]. It sends κ to zero and sends both e and $e\kappa$ to step 8. The target 7/6 maps to five steps, or 250 cents. Thus this example uses a target unavailable on the embedded 12-step grid. We compare two separate realizations on the same reference. Project every entry of Table 1 by k. All the choices within a cell have the same image. In part order A,

B, P, Q, H, the five rows of the registered 24-EDO path are

	A	B	P	Q	H
W_0	4	–	−10	−5	−24
W_1	−10	8	0	−5	−24
W_2	−10	8	3	3	−24
W_3	−10	8	−6	−6	−14
W_4	10	–	−6	−6	−14

In B's three sounding windows the target steps above P are $8, 5, 14$. The native 24-EDO realization can keep every written continuation. The JI lift must use pitches in the declared menus, project to the registered 24-EDO path and conserve the same frequency across each protected join, except where a concession explicitly permits a reattack.

Name the permissions b_2, b_3 for reattacking B at W_2, W_3, a_3 for taking P's 5/6 at W_3, and p_4 for reattacking P on return to 27/32 at W_4. A permission opens a choice, whether or not the chosen realization uses it. Without permissions P's written 27/32 continues from W_3 into W_4. The full list of protected joins is as follows. A and B each continue from W_1 through W_3. P continues from W_3 to W_4. Q continues across W_0, W_1 and across W_3, W_4. H continues from W_0 through W_2 and across W_3, W_4. Other entries after rests or unprotected boundaries are attacks. A necessary JI reattack is retained as a fresh articulation in its 24-EDO rendition even when the two projected step numbers coincide. A cap of 25 cents applies to any pitch change at an added reattack, and a separate allowance bounds the deviation of each JI pitch from its matched 24-EDO step.

Proposition 4.2. *With P fixed as the partner and* 35 *cents of deviation allowed, the inclusion-minimal permission sets are exactly* $\{b_2\}$, $\{b_3\}$ *and* $\{a_3, p_4\}$. *At* 20 *cents of deviation only* $\{b_2\}$ *remains minimal.*

Proof. The target in W_1 fixes B at e. If P remains at its written 27/32 in W_3, the third target fixes B at $(3/2)(27/32) = e\kappa$. The possible middle pairs (B, P) are (e, p) and $(e\kappa, p\kappa)$. Each join is possible in isolation, using a different middle pair. With B held through both joins, neither middle pair can satisfy both conditions.

All feasible trajectories must fall into the three cases below. The high middle pitch requires b_2, and its extension that keeps P's written values uses only b_2. The low middle pitch followed by a high B pitch requires b_3, and taking P at p in the middle and its written pitch in W_3 uses only b_3. If B stays low, the third target requires P at 5/6 and hence a_3. P's fixed 27/32 at W_4 then requires p_4. In each witness every added pitch jump is one comma, less than 25 cents, and every deviation is within 35 cents. None of the three necessary permission sets contains another.

The values p and $p\kappa$ both map to step 3. The first differs from that step by about -30.557 cents, and the second by about -9.051 cents. The 20-cent allowance excludes the low middle pair. The first witness satisfies the tightened allowance across all five parts, with maximum deviation about 18.826 cents. For the exact cap check, if r maps to step k, then $\max(r^{24}/2^k, 2^k/r^{24})^{50} \le 2^D$ is equivalent to deviation at most D cents. □

Choice	B in W_1, W_2, W_3 (Hz)	P at W_2, W_3, W_4	Added attack
Early	330, 334.125, 334.125	$p\kappa, 27/32, 27/32$	B at 1.25 s
Late	330, 330, 334.125	$p, 27/32, 27/32$	B at 1.75 s
Keep B held	330, 330, 330	$p, 5/6, 27/32$	P at 2.50 s

The first two choices use the fewest permissions. The third protects B's legato at the expense of changing P's line. With a jump cap of 20 cents none of the named-P repairs works, since each must move a protected pitch by one comma. For the wider 35-cent allowance, changing the partner by window to P, P, Q gives a different task with a held B, since Q already has 5/6

at W_3. The compatible-anchor control keeps P as partner and writes 5/6 for P in *both* W_3 and W_4. Then the low middle pair extends to a held B with zero permissions at 35 cents. If only the final anchor is changed to 5/6, a held B still requires a_3 to select P's alternative in W_3. These controls identify the separate roles of the partner, P's written choice and its final anchor.

A convenient complete check uses the admissible registered states D_t in each window, with the named identities, pitch choices and any persistent partner stored in the state. Write $E_t \subseteq D_t \times D_{t+1}$ for valid joins, and iterate $R_{t+1} = \{y : \exists x \in R_t,\ (x, y) \in E_t\}$ from the allowed initial states. By induction R_t contains exactly the states that end a legal prefix. Applying the final conditions tests for a whole phrase, as in relational consistency on a chain [2]. A nonempty E_t for every individual t fails to enforce the shared state at the middle window. The companion replay enumerates all 32 raw pitch paths in this construction before testing permissions and thresholds. The sets in Proposition 4.2 are inclusion minima, corresponding to minimal correction sets when disallowed actions are represented by soft prohibitions [9].

5 Where should a choir change pitch?

The opening eleven measures of Anton Bruckner's *Locus iste* (WAB 23) give two distinct questions. First, if we tune selected intervals exactly against the bass without changing any sustaining pitch, which written attacks need new pitches? Second, if we prescribe targets between all selected consonant pairs, how small can the worst interval error be? The second question changes the set of comparisons and the objective.

The score used here is Fig. 2 in Rosenzweig et al. [12], which credits Brian Marble's edition. We cross-checked it with Alexander Reuter's edition [3]. The aligned symbolic data also corroborate the written events. Our prescribed ratios are part of a tuning experiment. They are independent of the singers' recorded frequencies or any attribution of temperament to Bruckner.

Time q is measured in quarter notes from the beginning of m. 1. The section $[0, 44)$ has 109 written note occurrences, 25 rests, and 37 simultaneous-state windows, one of which is silent. One sounding occurrence keeps the same frequency from its attack to its release. This preserves 14 continuations across other voices' attacks. The three anchors are soprano C5 and bass C3 at $q = 0$, fixed at 528 and 132 Hz, and soprano C4 at $q = 43$, fixed at 264 Hz. This last anchor is inside the work. The excerpt stops before m. 12.

Let b_i be the base frequency ratio of written occurrence i to 264 Hz. Table 2 gives the fixed just-intonation base for each pitch class. Restore the written octave to obtain b_i. In the first question, assign frequency $264\, b_i \kappa^{z_i}$, with $z_i \in \{-1, 0, 1\}$. Constrain every pitch to lie within 40 cents of its 12-EDO frequency at the same reference. For a pair oriented from the lower written pitch to the higher, impose the ratios

semitones mod 12	0	3	4	5	7	8	9
simple ratio	1	6/5	5/4	4/3	3/2	8/5	5/3.

Multiply by the required power of two for a compound interval. Other interval classes carry no target. In the first question, if the bass is sounding we compare it with each sounding upper part whose interval class appears in the table.

	C	D	E	F	F♯	G	A	B
Base ratio (octave of C4)	1	9/8	5/4	4/3	45/32	3/2	5/3	15/8
Meantone offset $4x$	0	−2	0	1	−2	−1	1	−1

Table 2: The fixed just-intonation base and the later continuous witness. An offset x multiplies the base ratio by κ^x.

5.1 The fewest changed written attacks

Proposition 5.1. *Under the exact bass-relative targets, fixed anchors, original three-pitch palette and all written holds, the minimum number of written notes whose pitches differ from the fixed JI base is six. Exactly four pitch assignments attain it.*

Proof. For each of the 75 bass-relative comparisons between occurrences i, j, write r_{ij} for the prescribed registered ratio. Exactness imposes

$$\kappa^{z_j - z_i} = r_{ij}\, b_i / b_j.$$

In this excerpt each right-hand side is an integral power of κ, as checked by exact division. Take the graph with a vertex for each of the 109 written notes and an edge for each comparison, including repeats. The occurrence's vertex stays fixed through a hold. The edge differences determine the offsets in each connected component up to a common additive integer, as in Proposition 4.1. For each component, propagate the relative offsets, check all edges, and restrict the at most three root choices by the note palettes, the band and any anchors. Count nonzero exponents once per note. The objective is a sum over components, so minima add and the numbers of minimizing choices multiply.

There are 36 components. In 31 of them the all-zero choice is feasible and is the unique choice of zero cost. The other components have sizes $4, 4, 4, 4, 2$, and are localized by the times shown below. The middle entries list the edit counts for all feasible root choices.

Component location q	Counts at feasible roots	Minimum	Minimizing roots
12	$3, 1$	1	1
$13\frac{1}{2}$	$3, 1$	1	1
14	$2, 2$	2	2
16	$3, 1$	1	1
24–25	$1, 1$	1	2

At each of $12, 13\frac{1}{2}, 16$, the tenor offset must lie one comma above the other three, so raising the tenor alone uses one change. At 14 the soprano and tenor have the same offset, one comma above the alto and bass. Raising one pair or lowering the other each costs two. In the last component, the soprano at 25 must lie one comma below the offset of the bass that attacks at 24. Either of the two choices that change one of those notes is legal. This gives $1 + 1 + 2 + 1 + 1 = 6$ and $2 \cdot 2 = 4$. A second computation obtains the same count by extending consecutive windows and charging an edit at each new written attack. □

The score instructions are as follows. In each minimum, raise the tenor A3 at $q = 12, 13\frac{1}{2}, 16$ by one comma. At $q = 14$, either raise the soprano F4 and tenor A3, or lower alto D4 and bass D3. Independently, lower the soprano D5 at $q = 25$, or raise the bass F3 at its written attack at $q = 24$ and hold it through $q = 25$. Unlisted attacks use the base pitches.

A causal policy that chooses pitches at the current attacks and keeps sounding notes fixed attains the option of raising the soprano and tenor at 14 and lowering the soprano at 25. Against the base D3 of 148.5 Hz, the tenor fifth gives 222.75 Hz and the soprano compound minor third gives 356.4 Hz. At $q = 25$, the bass is already holding F3 at 176 Hz. The exact compound major sixth gives the soprano $176(10/3) = 1760/3$ Hz. The alternative in the preceding paragraph instead starts the bass at 178.2 Hz one quarter note earlier, and permits the soprano's base 594 Hz because $178.2(10/3) = 594$. Neither option adds a reattack. An optimum for a different objective, summed motion across the voices, alters 69 written pitches from the base while retaining zero added reattacks, exact ratios and exact return. That cost does not penalize such pitch changes. It therefore answers a different score question.

5.2 The least compromise between all targeted pairs

Now we prescribe the same table of targets for all sounding pairs to which it applies, including pairs above the bass. This produces 148 pair/window comparisons. The following instances conflict within the same sounding window.

q	Indirect route	Indirect ratio	Direct ratio
$28, 29\frac{1}{2}, 34$	C3–G4–D4–F5	$3(3/4)(12/5) = 27/5$	$16/3$
38	C3–G3–D4–F4	$(3/2)(3/2)(6/5) = 27/10$	$8/3$

In each row the two products differ by κ. The route through G4 and D4 traverses the required D4–G4 fourth in reverse. Figure 2 gives a way to check the second row by hand. The contradiction is simultaneous. A new articulation cannot make both ratios from C to F hold in the same window.

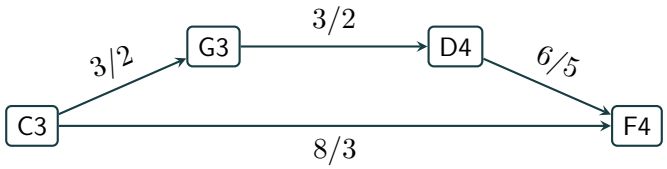

Figure 2: Measure 10, at $q = 38$. Each arrow asks for a pure frequency ratio. Starting at C3, the upper route requires F4 at $27/10$ times its frequency. The lower route gives $8/3$. Dividing the two answers gives $81/80$, the defect that the four comparisons must absorb.

We retain all anchors, holds, registers and the 40-cent band. Vary only the palette. In the continuous case each written occurrence can use any comma exponent x_i in $[-1, 1]$. In the n-subdivision case, n is a positive integer, and $x_i \in \{-n, \dots, n\}/n$. In both cases the objective is the maximum over pair/window comparisons of the absolute cent deviation from the prescribed ratio. This is an interval error, distinct from the pitch-to-grid allowance.

Proposition 5.2. *The minimum is $c/4$ in the continuous palette and $c\lceil n/4\rceil/n$ in the n-subdivision palette.*

Proof. Consider the four-edge cycle at $q = 38$ obtained by following the indirect route and reversing the direct one. The realized frequency ratios telescope. The signed cent errors relative to the directed targets therefore sum to $-c$, and some absolute error is at least $c/4$. The base pitches C3, G3, D4, F4 are $132, 198, 297, 352$ Hz. The first two fifths and the direct edge are exact at base, and the base ratio on D4–F4 divided by its target is $80/81$. Thus on the n-subdivision grid every cycle error is a multiple of c/n. This yields the rounded lower bound.

A fixed quarter-comma meantone scale supplies an upper bound for the whole excerpt. Tune along F–C–G–D–A–E–B–F♯, using a fifth of ratio

$$\frac{3}{2}\kappa^{-1/4} = 5^{1/4}$$

and restoring octaves as required. Relative to the C reference this gives the offsets x in Table 2. The anchors are on C and have $x = 0$. Each occurrence has a single fixed frequency over its full duration. The endpoints of the entire envelope $[-1, 1]$ satisfy the band for each written pitch. Specifically, if m is the number of semitones above C4 and $g = b_i\kappa^{\pm 1}$, the endpoint test is $2^{-2} \le g^{60}/2^{5m} \le 2^2$. The intervening pitches satisfy the band by monotonicity.

For any comparison write $d = \log_\kappa(b_j/(b_i r_{ij}))$, so the signed interval error in commas is $d + x_j - x_i$. Computing registered ratios from the score gives $d \in \{-1, 0, 1\}$ for all 148 comparisons. The table gives $4(d + x_j - x_i) \in \{-1, 0, 1\}$ in all those comparisons. Of these, 59 are zero and 89

have absolute value one. These exact integer checks over the stated score certify feasibility at $c/4$.

For the grid, take $y_i = \lfloor nx_i \rfloor / n$. Since nd is an integer, $nd + ny_j - ny_i$ is either the floor or the ceiling of $n(d + x_j - x_i)$. This follows by subtracting the two fractional parts of nx_i, nx_j, whose difference has magnitude less than one. The resulting integer has absolute value at most $\lceil n/4 \rceil$. The anchors are unchanged, and every y_i lies in the band-certified envelope. The same occurrence keeps its rounded value through all its windows. This gives the matching upper bound for every n. □

The corresponding minima for full, half, third, quarter and fifth steps are 21.506, 10.753, 7.169, 5.377 and 8.603 cents. Fifths of a comma omit the quarter-comma positions, so the optimum need not improve monotonically as the number of subdivisions increases. Every multiple of four attains the continuous optimum. Even allowing arbitrary positive frequencies or extra attacks could not improve on $c/4$ for these targets, since the lower bound uses only a single window.

Quarter-comma meantone and the minimax distribution of the syntonic comma are historical tuning ideas [13, 8]. Earlier models also consider optimization of interval error [11]. The score-specific finding is that one tuning per pitch class meets the lower bound throughout this excerpt, satisfying the temporal and absolute constraints. On this objective the ability to vary the pitches occurrence by occurrence therefore gives no improvement here. Changing the interval targets, anchors, palette or excerpt would require a new upper-bound check. These mathematical and synthetic examples establish feasible decisions and exact compromises for the models. They leave rehearsal utility and listening judgments open.

6 What can a rhythm tell us?

Write the seven pulse numbers $5, 6, 9, 11, 12, 18, 23$ on a circle numbered from 0 to 23. They are the observed attacks of one part, V_1. The pulse grid, its origin and the name of the part are known. Exactly one of its eight attacks was omitted. Can we recover it? First decide using those seven marks alone. Then suppose you are told which version came immediately before and that one move was made. The marks stay the same, but the answer changes.

Recover a missing attack

The source has three named parts. The fixed part V_0 attacks at

$$A = \{0, 1, 3, 4, 10, 15, 21, 22\} \subset \mathbb{Z}/24\mathbb{Z}.$$

The two moving parts start at $A + 8$ and $A + 16$, with addition modulo 24. Each occupies one pulse in each residue class modulo 8. A move swaps their two pulse assignments in one such class. A mask $m = (m_0, \ldots, m_7) \in \{0, 1\}^8$ selects which swaps to make from the starting assignment, with case identifier $\mathrm{UID}(m) = \sum_{r=0}^{7} 6^r m_r$. In each part, sort the eight pulses, take consecutive differences and include the wraparound gap $24 + s_0 - s_7$. An assignment belongs to this bank when every such gap in every part is at most 6. The opening example uses profile 0, with unit durations in the moving parts.

Let $R = \{5, 6, 9, 11, 12, 18, 23\}$. These seven pulses cover residues 1 through 7 modulo 8, and agree with $A + 8$ in every one. Thus all seven corresponding mask bits are fixed to zero. Residue 0 can contribute either 8 or 16. Both completions pass the gap test. The reader can check the entire computation using the companion supports below. In every row V_0 remains A.

UID	V_1 attacks	V_2 attacks	Maximal gaps (V_0, V_1, V_2)
0	$5, 6, 8, 9, 11, 12, 18, 23$	$2, 7, 13, 14, 16, 17, 19, 20$	$(6, 6, 6)$
1	$5, 6, 9, 11, 12, 16, 18, 23$	$2, 7, 8, 13, 14, 17, 19, 20$	$(6, 6, 6)$
6	$5, 6, 8, 11, 12, 17, 18, 23$	$2, 7, 9, 13, 14, 16, 19, 20$	$(6, 6, 6)$

With a known predecessor of UID 0, requiring exactly one move eliminates UID 0 and leaves UID 1, so the missing attack was at 16. Allowing zero or one move brings back UID 0 and the alternative missing attack at 8. UID 6 appears in the table for the next experiment. It cannot explain R, since it has no attack at 9. These decisions use only R and the stated predecessor condition.

Here is the general reason the exact-one-move condition helps in this model. A support S of size eight produces the observations $S \setminus \{s\}$ for $s \in S$. Two distinct supports of size eight share a seven-element observation exactly when their intersection has size seven, or equivalently when their symmetric difference has size two. Indeed, a common observation must be contained in their intersection, and a seven-element intersection is itself obtained by one deletion from either support. Since different mask coordinates use disjoint residue classes,

$$|V_1(m) \triangle V_1(m')| = 2\,\#\{r : m_r \neq m'_r\}.$$

Consequently two distinct bank states can explain the same observation exactly when the mask of one is a single swap away from the other. Two distinct neighbors of a fixed predecessor differ in two coordinates, so they cannot explain the same observation. This proves uniqueness conditional on existence. When zero moves are allowed, each edge from the predecessor to a neighbor again supplies a shared observation. These are statements about the permitted histories of symbolic attacks.

Make a move or put two cycles in succession

Retain the admitted states of odd mask parity from the component containing UID 0. They include UIDs 1 and 6. Every single swap changes parity. Therefore the retained vocabulary has no atomic step between different states, and each seven-attack observation has at most one explanation within it. Now write the row for UID 1 over pulses 0 through 23 and the row for UID 6 over pulses 24 through 47. Across the seam the gaps

$$24 + \min V_i(6) - \max V_i(1), \qquad i = 0, 1, 2,$$

are respectively $2, 6, 6$. Here $V_i(u)$ denotes the support for UID u. Each satisfies the gap bound. In profile 0 the moving parts have unit durations and the fixed part takes its duration from $D = (1, 2, 3, 1, 3, 1, 2, 3)$ indexed by the pulse residue modulo 8. The left cycle ends in V_0, V_1, V_2 at pulses $22, 23, 20$, so its three outgoing durations are $D_6 = 2, 1, 1$. None exceeds the corresponding seam gap. Thus the concatenation is also legal with monophony required, even though the states cannot be linked by moves inside the odd vocabulary.

Choose at random

Direct uniform sampling from N admitted states gives each state probability $1/N$ and can jump without using an edge. A walk which chooses uniformly from the $d(x)$ neighbors of the current state x respects atomic moves. On a connected, nontrivial undirected graph this walk has stationary probabilities proportional to $d(x)$, since the probability flux in either direction along an edge is the same. A walk with a uniform stationary law can instead assign each edge xy the transition probability $1/\max(d(x), d(y))$ and place the remaining mass on staying at x. These are the Metropolis–Hastings probabilities for a uniform target and uniform neighbor proposals [6]. The edge probabilities are symmetric, and each row sums to one.

The 256 masks can be checked by the gap calculation above. Join admitted masks which differ in one bit and explore from UID 0. This gives 30 admitted assignments in all and 15 in that component. Within it eleven vertices have degree 4 and four have degree 3, giving 28 edges. Every edge has at least one degree-4 endpoint. The uncorrected stationary probabilities are $4/56$ and $3/56$, while the corrected law gives $1/15$ to each. A degree-3 vertex has a $1/4$ chance of

staying under the correction. With duration profile 1, $(1, 1, 1, 1, 1, 1, 1, 2)$ for both moving parts, and monophony required, the source component is an eight-state cube. Every vertex has three neighbors and the correction gives no holding probability. The uniform law is stationary, yet a walk from a fixed vertex alternates between the two mask parities. Neither walk leaves its component. Random choice, a local change, and a legal seam are three different operations.

Verification and scope. From the P75 project root, run

```
python -B -X utf8 scripts/build_rhythm_choices_data.py -check
```

to check the retained finite catalog, including all 108 specified duration profiles. The hand example and the uniqueness proof above require only the labeled symbolic grid. They do not use a recording, a pitch, or the location of the attack omitted by the laboratory.

Conclusions

A chord description is useful to an arranger when it retains what the next decision depends on. The shared E in Figure 1 is a statement about three sets. Turning it into a held sound requires putting the same registered pitch in the same part through the changes. The named parts of Section 3 make such assignments explicit. In Section 4 the player's identity is already known and the written E still leaves a choice of just frequencies. The missing pulse in Section 6 has an analogous dependence on the previous event, with a different observation and different rules.

Two comparisons mark the limits of the musical conclusions. Autocorrelation already permits the two endpoint gains of Section 3. In Bruckner, a single pitch-class tuning attains the continuous all-pairs optimum on the transcribed excerpt. The six marked attacks solve a different problem with exact bass-relative targets. Each result can guide a specific experiment in the Studio. The exercises use constructed material and a specified reading of one score. Listening preferences, rehearsal outcomes and the ease with which a new visitor learns the material remain empirical questions. A synthetic realization lets a listener inspect the decision. The proofs and finite checks establish the stated consequences of the model.

AI assistance

OpenAI GPT models assisted with manuscript preparation, criticism of the proofs and computational checks. Three fresh model instances reviewed the first complete draft against its source material. The checks reported here concern the stated mathematics and finite examples. Human peer review and trials with musicians remain future work.

A Finding an experiment in the Studio

Start with one part, one decision and one comparison. This paper contains the inputs needed to check the principal examples without the application. The Studio lets you change a condition or choose another result from the same catalog. The complete local edition uses Italian labels. From the P75 project root run

```
python -B -X utf8 scripts/serve_studio_lab.py -port 8791
```

and open `http://127.0.0.1:8791/laboratorio`. Each case in the table also opens at `/studio?case=` followed by its identifier. In cases with sound, try reference A before choosing B. In the symbolic cases, inspect the positions or moves. *Tutti i casi* opens the catalog where offered, *Annulla* undoes the latest change, and *Testo grande* enlarges the text.

Area	Case and local label	First decision
Chord twins and voices	`c12-3`, Three voices (*Tre voci*)	Swap V_1 and V_3. Compare a tight leap bound with the return to the initial assignment.
	`c12-6`, Six voices (*Sei voci*)	With a three-semitone leap bound, compare all Q-twins with the transposition/inversion branch of the starting chord within that same Q fibre.
	`c24`, Quarter tones (*Quarti di tono*)	Change the middle-chord family and compare on the same scale. Browse ties before tightening the bound.
	`st19`, Symbolic paths (*Cammini simbolici*)	Lock channel k and inspect the moves left among the twelve positions.
Held notes	`holds`, Tuning and holds (*Accordatura e tenute*)	Protect D in V_1 in C–Am–F–Dm–G–C. Protect all eight requested holds, then make the final G in V_3 optional.
Tuning cycles	`cycles`, Tuning cycles (*Cicli d'accordatura*)	In C–F–Dm–G–C, compare preserving all holds and restarting G in V_3. Compare final with initial frequencies.
Register	`register`, Register limit (*Limite del registro*)	Protect the final G in V_3. Lower the upper endpoint to 990 Hz and check the B that lies at the excluded endpoint.
Motion	`motion`, Motion by voice (*Moto per voce*)	In C–F–Dm–G–C, test a 4/3 leap bound. Make the final G in V_3 optional and allow one broken hold. Compare the largest leap with total motion.
Rhythm	`rhythm`, Rhythm and observations (*Ritmi e osservazioni*)	Keep the seven observed attacks fixed and vary what you know about the predecessor. Check Section 6.

Two English lessons. At `/music-lab/` on the same local server, *Keep an E sounding* (the passing-E study) asks whether one frequency can survive both changes of partner pitch. Try first to keep the written pitches of P, then to keep E held throughout. Compare the two required frequencies, 330 and 334.125 Hz, and follow which part has to make a fresh attack. In *Tune a choir passage* (*Locus iste*, measures 1–11), follow the soprano's entrance over the bass in measure 7. Then compare the direct C–F ratio 8/3 with 27/10, the product along the three-edge route in measure 10. This second comparison answers the all-pairs question. The measure 7 comparison uses ratios to the bass. Sections 4 and 5 give the calculations and their assumptions. In the earlier English pilot, the same two lessons are available in the extracted `Tuning-Lab.zip` by opening `index.html` directly. The ZIP contains these two lessons. The nine cases above run in the local Studio.

Follow a result to its evidence. Use the exact values beside the outcome before consulting the sources. For instance, in the tuning-cycle case compare the first and last frequency of the same part. Under the hold-all policy the final value is the initial value times 80/81. The graph and WAV represent the selected realization. In a register case with no solution, the separate available pitches show individual possibilities, and the missing complete path is the result. The P75 project guide `docs/STUDIO_GUIDE.md` identifies the native cases, session formats and source reports.

The register and motion controls let you test two endpoint conventions. The register is half-open. For an upper limit of 990 Hz the B at 990 is unavailable. In the motion experiment the register is fixed at $[528, 1056)$ Hz and a leap bound includes equality. In C–F–Dm–G–C the penultimate B can be 8800/9 or 990 Hz. The highest returning note is 792 Hz, so its part must make a leap of ratio at least $(8800/9)/792 = 100/81$. With permission to give up the final G

in V_3, the last chord change can take $(1760/3, 7040/9, 8800/9)$ to $(528, 660, 792)$ Hz, and the preceding changes fit this same bound. Protecting that G instead forces B at 990 to return to at most 660 Hz, requiring the larger bound 3/2.

Replaying the printed examples. With Python 3.11 or later, run the following from the project root. The first command checks Figure 1, the six-part graphs at $t = 1$ and the missing pulse. The second reads the score transcription and checks the tuning of the passing E and both Bruckner optima.

```
python -B paper/R1/examples.py
python -B reports/MC4_MUSIC_STUDY_20261117/paper/replay.py
```

The score file `score.json` is in the same directory as the second script. The proofs above cover all positive integers t and all positive grid subdivisions n in their respective statements. Finite checks use the particular inputs supplied with each experiment.

References

[1] Oleksiy Babanskyy. Sharp denominator and mass bounds for nonnegative convolution-square fibres on cyclic groups of order $4p^a$, 2026. Paper and companion, version 0.2.1, worked C_{12} example.

[2] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30:479–513, 1983. Theorem 3.4.

[3] Anton Bruckner. Locus iste, WAB 23, 2016. Edited by Alexander Reuter. IMSLP score 426169, p. 1, systems 1–2.

[4] J. Douthett and P. Steinbach. Parsimonious graphs: A study in parsimony, contextual transformations, and modes of limited transposition. *Journal of Music Theory*, 42:241–263, 1998. Sections 2 and 5.

[5] Jakob Hansen. Comma pumps, impossible musical figures, and holonomy, 2026. Author's exposition, 5 August. Section on comma pumps.

[6] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970. Publisher's version.

[7] J. R. Hughes. Generalizing the orbifold model for voice leading. *Mathematics*, 10(6):939, 2022. Sections 2 and 4.

[8] D. C. Keenan. A method for optimally distributing any comma, 1999. Author's exposition, revised in 2001. Example with the syntonic comma and pure octaves.

[9] João Marques-Silva, Federico Heras, Mikoláš Janota, Alessandro Previti, and Anton Belov. On computing minimal correction subsets. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pages 615–621, 2013. Section 2.

[10] A. Milne, W. Sethares, and J. Plamondon. Tuning continua and keyboard layouts. *Journal of Mathematics and Music*, 2:1–19, 2008. Section 2.

[11] L. Polansky, D. Rockmore, M. K. Johnson, D. Repetto, and W. Pan. A mathematical model for optimal tuning systems. *Perspectives of New Music*, 47(1):69–110, 2009. Author's publication record.

[12] S. Rosenzweig, H. Cuesta, C. Weiß, F. Scherbaum, E. Gómez, and M. Müller. Dagstuhl ChoirSet: A multitrack dataset for MIR research on choral singing. *Transactions of the International Society for Music Information Retrieval*, 3:98–110, 2020. Figure 2.

[13] K. P. Scholtz. Algorithms for mapping diatonic keyboard tunings and temperaments. *Music Theory Online*, 4(4), 1998. Sections 6.1–6.2 and Figure 12.

[14] K. Stange, C. Wick, and H. Hinrichsen. Playing music in just intonation: A dynamically adapting tuning scheme, 2018. Author preprint, arXiv:1706.04338v3, section 4.

[15] T. Zaslavsky. Biased graphs. I. bias, balance, and gains. *Journal of Combinatorial Theory, Series B*, 47:32–52, 1989. Section 5.