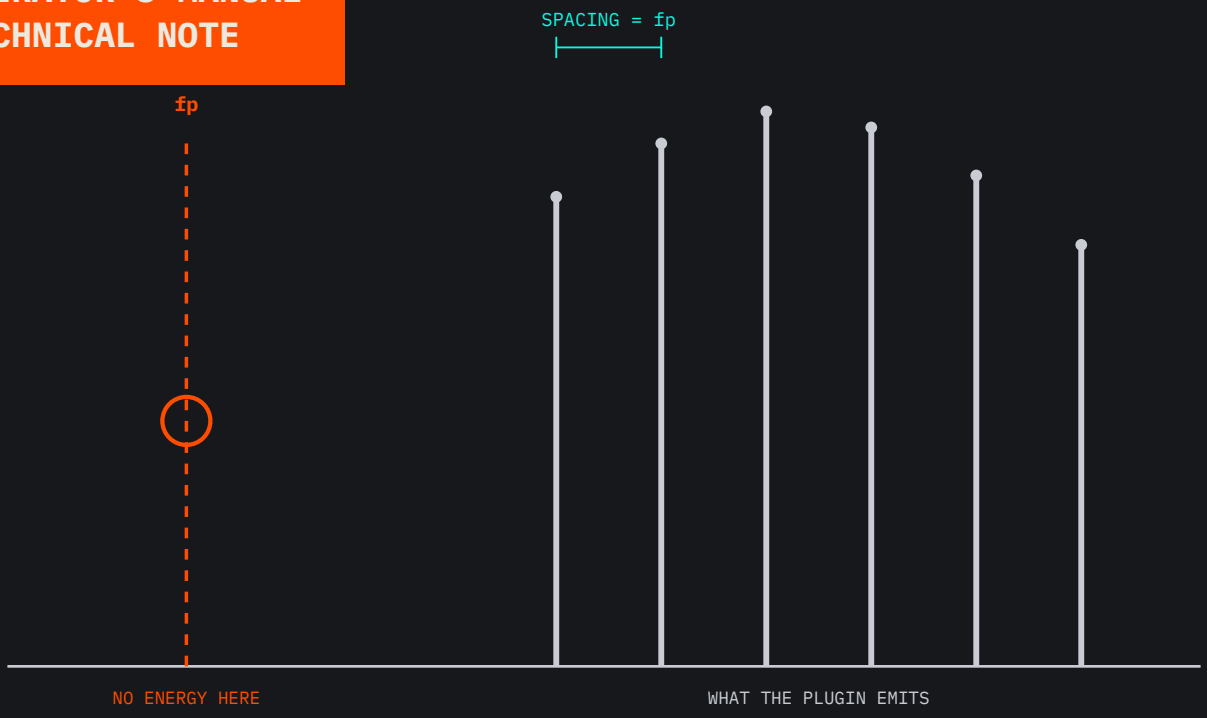
OPERATOR'S MANUAL
TECHNICAL NOTE



THIRDEAR

AUDITORY DISTORTION PRODUCT INSTRUMENT

For ThirdEar 1.4.0

> *How to set up and listen · the panel, the controls and the presets · the bench, for measuring your own ear · how the ear makes a note the plugin never plays · how the phase of the tones sets that note's level · one equation for spontaneous emissions, compression and distortion · why the two modes drive two different mechanisms · Maryanne Amacher · safe listening levels*

DATABURN
AUDIO RESEARCH
DATABURN.ORG

PSYCHOACOUSTIC SYNTHESIS

BEFORE YOU PLAY

ThirdEar plays pure tones. Your ear turns them into a note that is not in the signal: a distortion product, made in the cochlea, the hearing organ of the inner ear. This page says how to set up and how to hear it. The pages after it cover the panel, the controls, the presets and the bench; the second half of this note explains the science.

QUICK START

1. Put ThirdEar on an instrument track with nothing after it. Use headphones for the first session, and turn your system volume down before you start.
2. Open the preset list (the name at the top of the panel) and load Tune Passion Chorale. Hold one key. The tune plays from the key you hold; F#4 plays it in its own key.
3. Raise the volume slowly until you hear a low melody under the high, bright tones. Stop there. That is the level to play at: the lowest at which the phantom is audible.
4. The melody is made by your ear. Under the note name, the line SIGNAL AT fp reads "register empty": ThirdEar outputs nothing at the melody's pitch. Click A · YOUR EAR to switch to B, a real tone at that pitch, and click again to go back.

SET UP THE SIGNAL CHAIN

- Put nothing that distorts after ThirdEar. A saturator, clipper, limiter, amp simulator or a clipping master bus makes a real tone at the phantom's frequency, and then you hear your speakers, not your ear. The meter on the panel measures ThirdEar's own output and cannot see what comes after it. EQ, a volume change or reverb add no new frequencies and are safe.
- Keep the phantom's register empty. A distortion product is easy to mask: a bass line, a pad or noise near its pitch hides it.
- Headphones give each ear a known level and keep the room out, and SPLIT needs them. Speakers work too, and the phantom can seem to come from inside your head, but the room changes the level at each ear and each ear hears both channels.

SET THE LEVEL

Play at the lowest level at which the phantom is audible. The product grows faster than the tones that make it, about 2 dB for every dB in QDT mode [13], so a small rise in level makes the phantom much stronger. Set the system volume first, then use Level (default -15 dB) for small changes. Nothing on the panel raises Level in one action: loading a preset, the reset in the settings menu and a double-click on LEVEL all keep the lower of the two values. FIND LEVEL on the bench measures your own threshold (page 7). The exposure limits are on page 16.

WHAT TO LISTEN FOR

A low tone under the high tones. In both modes it is the note you play: in QDT mode far below the comb, in CDT mode just under the pair. It can seem to sit inside your head rather than in the speakers. To check that it is your ear:

- Turn Level down slowly. The phantom fades before the high tones do, because it falls faster than they do.
- Click A · YOUR EAR. B plays a real tone at the phantom's pitch to compare with.
- Click CANCEL and add a real tone at the phantom's frequency. At the right level and phase the phantom disappears.
- Set PHASE to ALT. The tones keep their frequencies, but the low note moves up an octave. SCR makes it weaker.

HEARING

The effect needs a healthy cochlea. With hearing loss in the register of the tones, the phantom may be weak or absent. If your ears ring or feel full after a session, the level was too high: stop for the day.

THE PANEL

The panel at its default size, holding C3 in QDT mode. The numbers match the table under it. A control the current mode does not use is drawn dimmed.

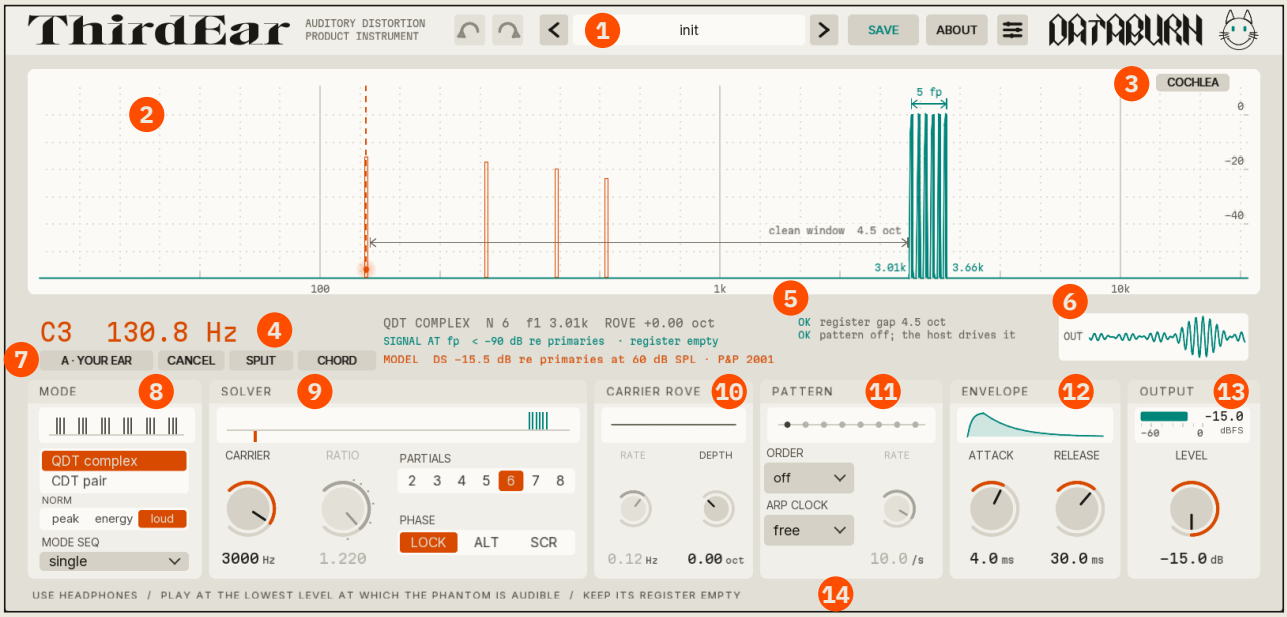


fig. 1: ThirdEar 1.4.0, light theme in the light editions of this note, dark in the dark.

NO.	PART	WHAT IT SHOWS OR DOES
1	Preset bar	Undo and redo; the previous and next preset, and the preset list under the name; SAVE a preset of your own; ABOUT; the settings cog (dark or light, or reset every control). The cat at the right also switches dark and light.
2	Screen	The output spectrum, measured. Teal lines are the tones ThirdEar plays. The dashed orange line marks fp, the phantom's frequency, where the signal holds nothing. Outlined orange bars are the products the model predicts your ear makes there and at its harmonics. "clean window" is the gap from fp to the lowest tone, in octaves.
3	COCHLEA	Switches the screen to the cochlea unrolled, low frequencies at the left: where the tones are felt, where the product is made, the place where it is heard, and, dashed, the emission that travels back out to the ear canal. That emission is what a probe microphone measures, not what you hear. Greenwood's map [9].
4	Readout	The note and fp; the mode line (mode, partials, the lowest tone f1, the rove); the meter; and MODEL, the product the model predicts re the tones at 60 dB SPL, with its source. The meter, SIGNAL AT fp, measures ThirdEar's output at the phantom's frequency, re the loudest tone: "register empty" is more than 60 dB down, and "< -90" under what it can resolve. Its frame, 171 ms to 1 s, is sized from the nearest tone; when a tone is too close even for the longest it says "unresolved" rather than guess. With the cancel tone on it reads that tone, and under B the real note.
5	Conditions	Up to four checks, OK or !, from the documented conditions for the effect: the gap between phantom and tones, the pattern's speed, the ratio, a level held under the clip or set above Level by Norm, chord voices, SPLIT on headphones, and Level against your threshold once FIND LEVEL has found one.
6	OUT	The output waveform.
7	Chips	A/B, CANCEL (opens the bench), SPLIT and CHORD.
8	MODE	A picture of what the mode plays, the mode switch, NORM and MODE SEQ.
9	SOLVER	CARRIER, RATIO, PARTIALS and PHASE (Start phase), with a picture of the comb or pair.
10	CARRIER ROVE	RATE and DEPTH: Rove rate and Rove depth in the host.
11	PATTERN	ORDER (Pattern), ARP CLOCK and RATE (Pattern rate). The dots show the steps.
12	ENVELOPE	ATTACK and RELEASE.
13	OUTPUT	The output meter and LEVEL.
14	Footer	The standing safety line. While the mouse rests on a control, the right end of the footer explains it in one line.

GESTURES

Drag a knob up or down to turn it; hold Shift for fine steps, at 1/8 of the speed. The mouse wheel steps it. Double-click a knob to reset it to its default; LEVEL resets only downward. Click a list to open it, and click it again or anywhere else to close it. Esc closes the About card, the bench and any open list. The theme is kept in %APPDATA%\databurn\thirdear\settings.txt and survives a restart; presets you save are files in %APPDATA%\databurn\thirdear\presets.

CONTROLS
1 OF 2

ThirdEar has twenty-one parameters. The table gives each one's name in the host, where it sits on the panel, its range and default, and the mode it works in; the panel dims a control the current mode does not use. "re" means "relative to". To hear the instrument straight away, load a Tune preset and hold one key (page 2).

CONTROL	ON THE PANEL	RANGE, DEFAULT	MODE	WHAT IT SETS
Mode	MODE	QDT complex or CDT pair; default QDT complex	both	Which kind of product you play. QDT complex plays a comb of tones spaced by the note's frequency; the ear's quadratic difference tone f2 − f1 lands on the note. CDT pair plays two tones whose cubic difference tone 2f1 − f2 is the note (page 13). In CDT a pair whose lower tone would fall under 20 Hz is not played: at the default ratio the bottom 12 notes, C-1 to B-1, are silent.
Carrier	SOLVER, CARRIER	800 Hz to 8 kHz; default 3000 Hz	QDT	Where the comb sits. Its lowest tone is the note's harmonic nearest this frequency, and the others are the next harmonics up. The phantom does not move with it. Changes glide over about 1.5 s, so faster automation is smoothed.
Partials	SOLVER, PARTIALS	2 to 8; default 6	QDT	How many tones in the comb. Each doubling of the number of adjacent pairs adds about 3 dB to the phantom in listeners, and 6 dB in the simple model; page 9 explains the gap. A tone above the sample rate's limit is not played.
Ratio	SOLVER, RATIO	1.05 to 1.30; default 1.22	CDT	f2/f1, the ratio of the pair's two tones. It sets how far under the pair the phantom sits and the chord the three make. The knob catches at the six chord ratios of page 14. The panel draws the heard CDT strongest at 1.10 and falling in a straight line to 20 dB weaker at 1.30 [3]; the two ends are measured, the line between them is assumed.
Rove rate	CARRIER ROVE, RATE	0.02 to 2 Hz; default 0.12 Hz	QDT	How fast the carrier wanders. The comb moves to the nearest harmonics of the note as it goes, so the tones move and the phantom stays where it is. Moving tones keep the ear from adapting to a steady sound. Does nothing while Rove depth is 0.
Rove depth	CARRIER ROVE, DEPTH	0 to 1 oct; default 0	QDT	How far the carrier wanders either side of Carrier, in octaves. A fast, deep rove leaves a faint real trace at fp.
Attack	ENVELOPE, ATTACK	0.5 to 100 ms; default 4 ms	both	How fast a note rises. In the model of page 8 it does not change the phantom's level.
Release	ENVELOPE, RELEASE	2 to 500 ms; default 30 ms	both	How fast a note fades once its key is up. It is a time constant: the tail lasts several times this value.
Level	OUTPUT, LEVEL	-60 to -6 dB; default -15 dB	both	Output level. Play at the lowest level at which the phantom is audible (page 2). The QDT grows about 2 dB per dB of the tones [13] and the CDT about 1 dB per dB [1], so Level changes the balance between the modes as well as the loudness. The synth is held under the safety clip, because a clipped signal puts a real tone at the phantom's pitch. No single action raises Level: a preset, a reset and a double-click keep the lower value.
Norm	MODE, NORM	peak, energy or loud; default loud	both	How each tone's level is set. Peak (1/N) holds the summed peak; the phantom falls about 7 dB from two partials to eight in the model. Energy (1/√N) holds the summed power; it rises about 5 dB. Loud holds the tones' A-weighted power (A-weighting: a standard curve that discounts low and very high frequencies the way the ear does) at the default comb's, so a CDT pair at C4 plays level with the comb, not 2 dB under it. Loud adds at most 6 dB, fading out below 40 Hz (none at 20) and above 4 kHz (none from 8 kHz). When a note sits above Level, the conditions strip says by how much.
Pattern	PATTERN, ORDER	off, 11 orders, 4 tunes; default off	both	Plays the held keys as a pattern (up, down, pingpong, pingpong+, converge, diverge, thumb, pinky, walk, random or played), or one of four tunes from the first key held, so that key sets the tune's key. Off plays the keys directly, so a host arpeggiator works. Remembers up to 8 held keys. In CDT the steps are legato: the pair glides to the next note over 6 ms, and a repeated note holds as one note.
Pattern rate	PATTERN, RATE	2 to 20 /s; default 10 /s	both	Steps per second while Arp clock is free; ignored when it is a division. An arpeggio wants 5 to 15 steps a second and a tune 1 to 7 notes a second. The conditions strip checks both, and warns when a step holds fewer than five cycles of the phantom.
Start phase	SOLVER, PHASE	locked, alternating or scrambled; default locked	QDT	Each tone's phase at note-on; it does nothing to a note already sounding. Locked starts every tone at its peak together (cosine phase), which makes the strongest phantom. Alternating shifts every other tone by a quarter turn: the odd components cancel and the pitch rises an octave. Scrambled uses random phases and costs about 8 dB in the model (page 9).

CONTROLS
2 OF 2

CONTROL	ON THE PANEL	RANGE, DEFAULT	MODE	WHAT IT SETS
Mode seq	MODE, MODE SEQ	single, or 10 N:M patterns; default single	both	Alternates the two modes across successive notes. N:M plays N notes in the mode the Mode switch sets, then M notes in the other mode, and repeats. Single turns it off.
Arp clock	PATTERN, ARP CLOCK	free, or 9 divisions, 1/1 to 1/32; default free	both	What clocks the pattern. Free uses Pattern rate. A division sets the step rate from the host tempo; the steps start with the first key, not on the beat. At 120 bpm 1/16 is 8 steps a second and 1/8 is 4.
Cancel	bench, CANCEL TONE	off or on; default off	both	A real sine at the phantom's own frequency, locked in phase to the tones: Goldstein's cancellation method [1]. Turn Cancel level and Cancel phase until the phantom disappears (page 7). A QDT comb is harmonics of the note, so the brain can still hear the note's pitch from them after the null (the missing-fundamental pitch).
Cancel level	bench, LEVEL	-60 to -5 dB; default -20 dB	both	The cancel tone's level, re one tone of the comb. Where the phantom disappears, this is the level of a real tone that matches the product your ear makes. Headphones colour it. The product's phase can change with level, so null again after changing Level.
Cancel phase	bench, PHASE	0 to 360 deg; default 0 deg	both	The cancel tone's phase, re a sine at fp that starts with the note. Every adjacent pair of tones differs by exactly fp, so the product's phase is locked to the tone, and the null holds for the whole note and while the carrier roves (the release tests check both).
A/B	A · YOUR EAR	your ear or real tone; default your ear	both	A: the tones alone, and the note is made in your ear. B, REAL TONE: the tones fade out and a real sine at the phantom's pitch plays at Cancel level; the meter shows it.
Split	SPLIT	both ears or split; default both ears	both	Sends every other tone to each ear; headphones only. A product needs its tones in one cochlea, so nothing is made at fp. In QDT each ear's comb is spaced 2 fp and its phantom moves an octave up, 6 dB weaker at six partials. The brain can still combine the two ears' tones into a pitch at the note [10]. In CDT each ear gets one tone and there is no CDT. Turn Cancel off first.
Chord	CHORD	off, just or equal; default off	both	With Pattern off, each held key sounds its own phantom, up to 4. A chord always plays QDT combs, whatever Mode says: a CDT pair's tones would land on or near the chord's other notes and be heard as real notes. Tuned just to the lowest key, the combs' cross products land on the chord's root; tuned equal they scatter (stray product -42 and -4 dB for C major in the model). A new lowest key retunes the others (C E G, C let go: E and G move 14 cents). Each voice adds level.

Velocity, pitch bend, the mod wheel and sustain are not used. A host that sends per-note tuning (CLAP note expressions) retunes each note. With Chord and Pattern off, one note sounds at a time: letting go of the sounding key brings back the most recent key still held.

In CDT mode Carrier, Partials, Rove rate, Rove depth and Start phase do nothing; in QDT mode Ratio does nothing. Rove rate does nothing while Rove depth is 0, and Pattern rate nothing while Pattern is off or Arp clock is a division. The panel dims each of them while that is so.

SIGNAL PATH, ONE VOICE

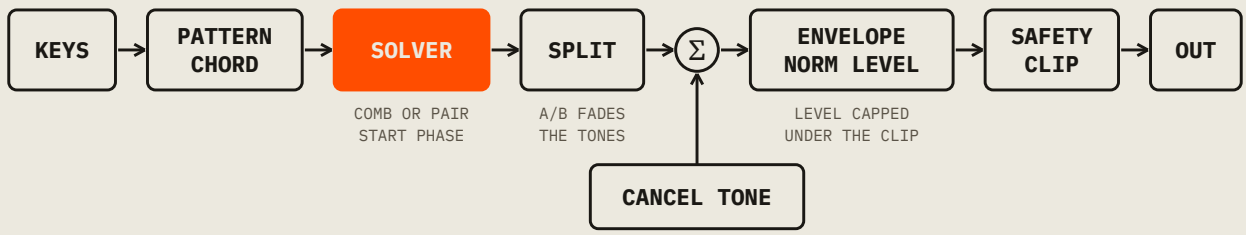


fig. 2: Each voice: the solver places the tones for the note, SPLIT sends every other tone to each ear, the cancel tone joins both ears, and envelope, Norm and Level set the gain. Level is held under the safety clip, so the clip never shapes the tones; only the synth reaches it.

PRESETS

27 factory presets come first in the preset list, then your own. The 16 experiments each reproduce a measurement or a documented sound, with the paper or the key to play in the name. Loading a preset never raises Level. PLAY names the key to hold; "any" means any key in the middle of the keyboard.

EXPERIMENTS

PRESET	MODE	PLAY	WHAT IT REPRODUCES OR PLAYS
P&P exp 2 one pair	QDT	G2	Pressnitzer and Patterson's experiment 2 [2]: one pair of tones near 1.5 kHz, whose product is the reference level, -22.5 dB re the tones at 60 dB SPL. Level sits at -30 so each tone is as loud as in exp 1.
P&P exp 1 cosine comb	QDT	G2	Experiment 1: the first eight tones of the paper's comb, in cosine phase. More pairs than exp 2, so a stronger phantom.
P&P exp 3 alternating	QDT	G2	Experiment 3: seven tones, every other one a quarter turn on. The odd components cancel and the pitch rises an octave.
Notebook 01 scrambled	QDT	any	Six tones in random phase: about 8 dB weaker than locked (page 9).
Kendall 3:4:5 major M3	CDT	any	Kendall, Haworth and Cadiz, table 1 [3]: a pair at 1.25; CDT, f1 and f2 stand as 3:4:5, a major triad (M3).
Kendall 4:5:6 major m3	CDT	any	Kendall, Haworth and Cadiz, table 1 [3]: a pair at 1.20; CDT, f1 and f2 stand as 4:5:6, a major triad (m3).
Kendall 5:6:7 diminished	CDT	any	Kendall, Haworth and Cadiz, table 1 [3]: a pair at 1.166; CDT, f1 and f2 stand as 5:6:7, a near diminished triad.
Kendall 6:7:8 non-tertian	CDT	any	Kendall, Haworth and Cadiz, table 1 [3]: a pair at 1.1428; CDT, f1 and f2 stand as 6:7:8, a non-tertian triad.
Kendall 8:9:10 whole-tone	CDT	any	Kendall, Haworth and Cadiz, table 1 [3]: a pair at 1.111; CDT, f1 and f2 stand as 8:9:10, a whole-tone cluster. Under 1.14, so rough.
Mauermann 1.22 DPOAE peak	CDT	any	A pair at 1.22, where the emission measured in the ear canal is largest [4]. The tone you hear is weaker here than at 1.10.
Goldstein 1.10 reference	CDT	any	A pair at 1.10, Goldstein's reference ratio [1], where the heard CDT is strongest, about -15 dB re the tones.
Tartini 1754 third tone, G3	QDT	G3	Two tones, 784 and 980 Hz, a third apart like a violin double stop. Tartini's third tone [14] sounds at G3 (the QDT); the CDT falls on 588 Hz.
Kendall 1a harmonics 5-8, C4	QDT	C4	Kendall's Example 1a [3]: harmonics 5 to 8 of middle C, 1308 to 2093 Hz. In this low register the brain's pitch from the harmonics (residue pitch) mixes with the QDT.
Umngqokolo 9:8 pair	CDT	any	A pair a whole tone apart (1.125), after Xhosa umngqokolo overtone singing and the umrhubhe mouth bow; CDT, f1 and f2 stand as 7:8:9. Rough, under 1.14. A model, not yet checked by ear.
Quintina 4:3:2 C2 F2 C3	QDT, chord	C2 F2 C3	Chord just. Harmonic 4 of C2's phantom, 3 of F2's and 2 of C3's all fall on C4: after the Sardinian quintina, a fifth voice no singer sings. A model, not yet checked by ear.
After Schorpp high pairs, G5	CDT	G5 to G7	Pairs at 1.20 with f1 from 980 to 3920 Hz, after Estelle Schorpp's In My Ears (2024). The phantom stays high, 0.32 octave under f1. The ratio is ours, not hers.

SOUNDS

PRESET	MODE	PLAY	WHAT IT REPRODUCES OR PLAYS
Wide comb	QDT	any	8 tones near 4.5 kHz with a slow rove.
Arp up	QDT	a chord	Pattern up at 10 steps a second.
Pingpong slow	QDT	a chord	Pattern pingpong at 5 steps a second.
Scrambled drift	QDT	any	Scrambled phase, new on every note, and a wider, faster rove.

AFTER AMACHER

PRESET	MODE	PLAY	WHAT IT REPRODUCES OR PLAYS
After Amacher 77 Hz	QDT	D#2	Two tones under a slow, deep rove. Amacher said a low D-sharp "of 77hz will sound from many different intervals" [11].
After Amacher into a third	CDT	any	A pair a major third apart (1.25) with a long attack and release, after her pair that "turns into a shape" as it closes on a third [11].
After Amacher ostinato	QDT	a chord	Pattern up at 6 steps a second, after the four-note ostinato of Head Rhythm 1 [12].

TUNES

PRESET	MODE	PLAY	WHAT IT REPRODUCES OR PLAYS
Tune Passion Chorale	QDT	F#4	The Passion Chorale in Bach's setting, 3 steps a second.
Tune Gymnopedie 1	QDT	A	Satie's Gymnopedie No. 1, 3.9 steps a second, with a slow rove against adaptation.
Tune Ave maris stella	QDT	C	The chant Ave maris stella, 4 steps a second.
Tune I IV V I 3:4:5	CDT	any	A cadence in CDT at 1.25, so each note is heard as a major triad, 2 s a chord.

THE BENCH

CANCEL, under the note, opens the bench. It measures your own ear by Goldstein's cancellation method [1]: a real tone at the phantom's frequency, whose level and phase you turn until the phantom is gone. At that point the tone is equal and opposite to the product your ear makes, so its settings measure that product. The bench shows three numbers: THE MODEL, the product predicted re the tones at 60 dB SPL, with its source; YOUR SETTING, the cancel tone's level and phase; and MEASURED AT fp, what the meter reads in the signal there, which is your cancel tone. The bench keeps what you measure, so the instrument can run small listening experiments as well as play music.

THE ACTIONS

BUTTON	WHAT IT DOES	WHAT IT ANSWERS
CANCEL TONE, LEVEL, PHASE	The cancel tone and its level and phase (Cancel, Cancel level and Cancel phase in the host). Shift-drag turns them in fine steps.	How strong is the phantom in my ear?
RECORD	Needs a held note and the cancel tone on. Writes the null to readings.csv as one row of kind null, and lists this note's earlier nulls. When nulls under LOCK and ALT exist for the same note and settings, it prints the difference: your alternating-phase suppression.	Does my null hold from day to day? How much does alternating phase remove for me?
GROWTH -6 dB, -3 dB	Keeps the first null, then lowers Level by 6 dB, or 3 dB if the phantom is lost at 6. Null again and press SECOND NULL. The growth is 1 + (change in Cancel level) / (change in Level), in dB per dB; both rows go to readings.csv.	Does my phantom follow the square law? A QDT should grow 2 dB per dB [13], a CDT about 1 [1].
FIND LEVEL	First plays B, the real tone at the phantom's pitch, so you know what to listen for. START: Level falls 1.5 dB a second until you press GONE, then rises until you press BACK, and again: 4 turns. Their mean is your threshold for this mode, kept in settings.txt. Level is set 3 dB above it, never higher than where it started, and the conditions strip shows Level against it. STOP ends the walk.	What is the lowest level at which I hear the phantom, on these headphones at this volume? It is not dB SPL.
COPY STIMULUS	Puts the sounding stimulus on the clipboard as CSV: each tone's frequency, level re full scale and start phase, the cancel tone if it is on, and each product the model predicts with its source.	Can a notebook rebuild exactly what I heard and check it against the model?
HEADPHONES	Names the headphones. The name is kept and written with every reading, since headphones colour the reading.	Which readings can be compared?

READINGS.CSV

One row per reading, under a header row, in %APPDATA%\databurn\thirdear\readings.csv beside the presets folder. Esc or a click outside the card closes the bench.

COLUMN	MEANING
date	date and time, local, ISO order
headphones	the name given with HEADPHONES
kind	null, or growth for the two rows of a growth test
mode	QDT or CDT
note	the note held, as C3
fp_hz	the phantom's frequency
ratio	Ratio
partials	Partials

COLUMN	MEANING
phase	Start phase: locked, alternating or scrambled
norm	Norm: peak, energy or loud
level_db	the Level knob
cancel_level_db	the null's level, re one tone
cancel_phase_deg	the null's phase
model_db	the model's figure, re the tones at 60 dB SPL
slope_db_per_db	growth, on a growth test's second row; 0 otherwise

OPEN QUESTIONS THE BENCH SERVES

- Alternating phase. Pressnitzer and Patterson found the odd components cancelled fully for one listener and partly for the other [2]. Which is typical? RECORD under LOCK and under ALT.
- Growth. Most ears' QDT follows the square law, but some listeners' does not [15]. How common is that? The growth test measures one ear.
- Your own map. The phantom is not equally strong at every pitch, and ears differ (page 12). Can a listener map their own without clinical equipment? Nulls recorded across the keyboard are a first map.
- Register. How audible is the CDT with its tones at 500 Hz, 2 kHz or 4 kHz? FIND LEVEL and nulls in CDT at low and high notes give one listener's answer.

These are open questions, not results. Readings from one listener on one pair of headphones answer them for that listener only.

THE PHANTOM

In QDT mode, the default, ThirdEar plays a comb of tones, called the primaries (6 by default), spaced by exactly the frequency of the note you play, written fp below. The lowest tone is the harmonic of fp nearest the Carrier frequency, and the others sit above it. In CDT mode it plays a pair of tones instead, placed so that their cubic difference tone lands on fp (page 13). Either way nothing is played at fp itself, unless you turn on Cancel or A/B; a fast, deep Rove leaves a faint trace (page 4). Any energy at fp would put the note in the signal as a plain sine wave, and there would be no illusion.

The listener hears the note anyway, because the ear is not a linear device. A compressive, asymmetric nonlinearity applied to a sum of tones produces sum and difference frequencies that were not in the input. The two that matter here are the quadratic difference tone (QDT) $f_2 - f_1$ and the cubic difference tone (CDT) $2f_1 - f_2$. Model the nonlinearity with the first terms of a power series, $y = x + a_2x^2 - a_3x^3$. Because the comb is exactly harmonic, the even term's difference tone for every adjacent pair lands on the same frequency, and they sum there. The odd term's products fall at $(m + i + j - k)\cdot fp$ for partial indices up to $N - 1$, so for a comb whose lowest harmonic number m exceeds the partial count N (fp below carrier/(N + 0.5); here m = 27 against 6 partials) it cannot reach fp at all, and removing it leaves the phantom unchanged. Above that boundary it can: at C5 the comb is harmonics 6 to 11, and at these coefficients the cubic contribution sits 13 to 42 dB under the quadratic one across the instrument's range of partial counts and harmonic numbers. The gap is smallest with two or three partials on the lowest harmonics: for a pair at harmonics 2 and 3, the CDT and the QDT are the same frequency. Removing the even term erases the phantom everywhere. Page 13 runs the same test on a model that has only the cubic term.

A 110 Hz PHANTOM FROM A 3 kHz COMB

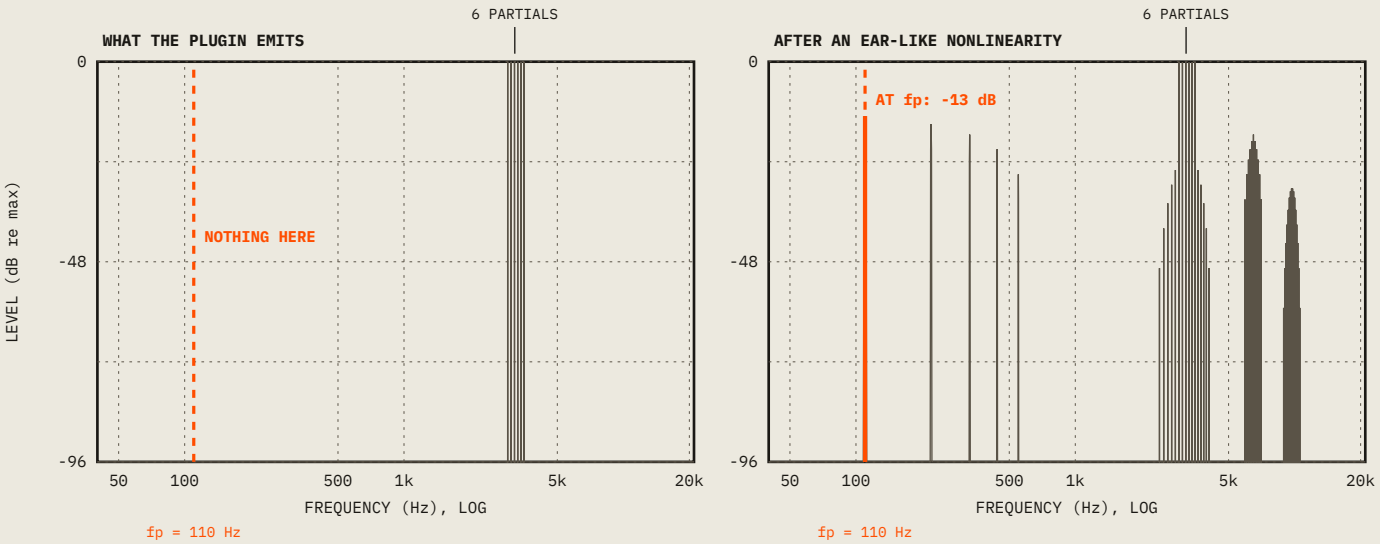


fig. 3: The same second of audio, before and after a memoryless compressive nonlinearity. The phantom register is empty to the numerical floor on the left and carries a component on the right. Nothing was added to the signal: the nonlinearity is the only difference between the two panels.

	LEVEL AT fp	READING
emitted signal	-180 dB	the numerical floor: there is no energy at the phantom frequency, by construction.
after the ear	-13 dB	generated inside the listener, from partials more than four octaves above it.

PHASE

The distortion product is not generated once. Every adjacent pair in the comb contributes its own difference tone at the same frequency, and what the listener hears is the vector sum of all of them, a sum in which each contribution has a size and a phase, so contributions can cancel. Pressnitzer and Patterson measured this directly [2]: with the primaries in cosine phase (all starting at their peak together) the contributions add, and each doubling of the number of pairs adds about 3 dB. Scramble the phases and they partially cancel.

This is the reason the instrument resets every oscillator phase together at note-on, and the reason Phase is a control on the panel (Start phase in the host). Two hundred random-phase trials against the phase-locked stimulus put the cost of scrambling at 8 dB of phantom level.

WHAT PHASE SCRAMBLING COSTS THE PHANTOM

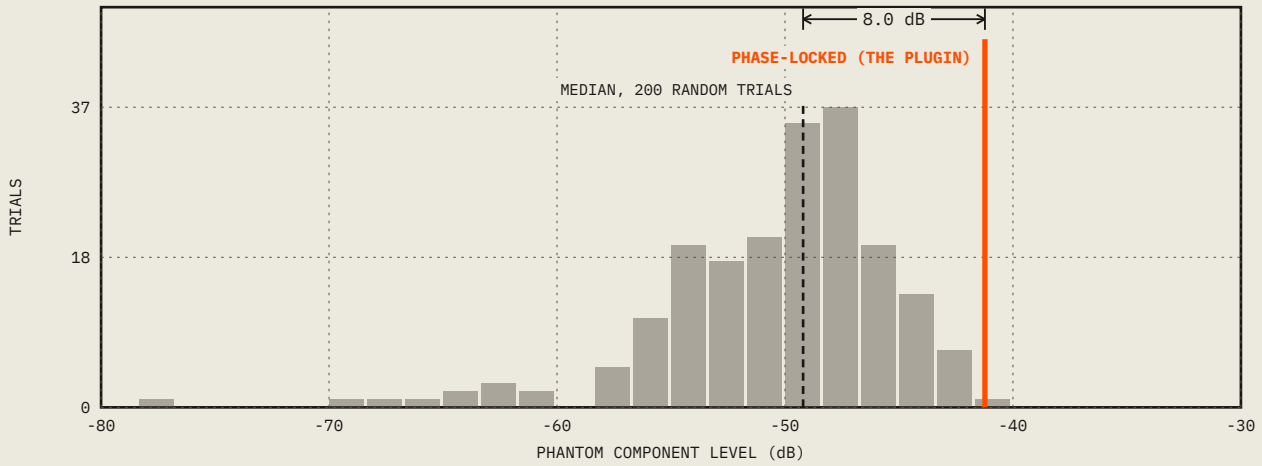


fig. 4: Two hundred trials with the six partials given uniformly random phases, against the phase-locked stimulus the plugin actually emits. Scrambling never helps and typically costs about 8 dB. The distribution has a long left tail: some phase draws cancel the phantom almost entirely.

PHASE SETTING	WHAT IT DOES	MEASURED
Locked	all partials in cosine phase; contributions add	-41.2 dB, the reference
Alternating	every other partial shifted by π/2. Pressnitzer and Patterson's experiment 3 [2]: the odd components of the distortion spectrum cancel, leaving a spectrum built on 2fp, so the percept rises an octave. With six partials, one of the five pairs has no partner to cancel.	-55.2 dB, 14 dB down (five pairs: 14 dB in closed form)
Scrambled	uniformly random phases per note	-49.2 dB median, 8 dB down

LIMITS OF THIS MODEL

The nonlinearity of page 8 is memoryless, so every pair of primaries contributes its difference tone with the phase difference of that pair and nothing else. It gets the alternating-phase experiment right for exactly that reason: a quarter turn on every other partial makes adjacent pairs differ by plus and minus a quarter turn, and their difference tones cancel in twos. (A half turn does nothing, in this model or any other: the comb with every other partial flipped is the locked comb delayed by half a period of fp, and the first release of this instrument shipped that by mistake.) What the model gets wrong is the buildup. Doubling the pair count adds a full 6.0 dB here against about 3 dB reported in listeners. One plausible reason is that the travelling wave reaches each place at a different time, so a real cochlea gives each pair its own phase and the contributions only partly reinforce; Pressnitzer and Patterson also report the product's phase nearly constant across partial counts, so the mechanism is not settled. Whichever it is, it is an ingredient this model lacks.

THE HOPF OSCILLATOR

The power series on page 8 accounts for the phantom and nothing else. It cannot explain why the ear amplifies quiet sounds far more than loud ones, and it cannot explain why a healthy ear emits sound of its own [6]. Those are usually presented as three separate properties of the cochlear amplifier: active, nonlinear, sharply tuned. One equation produces all three. It is the normal form of a Hopf bifurcation: the simplest equation for a system that switches from silent to oscillating as one setting passes zero.

$$\dot{z} = (\mu + i\omega)z - |z|^2 z + F(t)$$

the normal form of a supercritical Hopf bifurcation [5]

z is the displacement of one section of the cochlear partition, ω its natural frequency, and μ the control parameter that decides whether the section is damped or self-oscillating. Below zero it is silent and needs a stimulus. Above zero it oscillates with no input at all, which is a spontaneous otoacoustic emission: a sound the ear makes rather than receives. Many adult ears have one, and women's more often than men's [7]. Just above zero the amplitude grows as the square root of μ: it rises smoothly from nothing rather than jumping, which is the mark of the supercritical case. (The plots write μ as mu.)

ONSET OF A SPONTANEOUS EMISSION

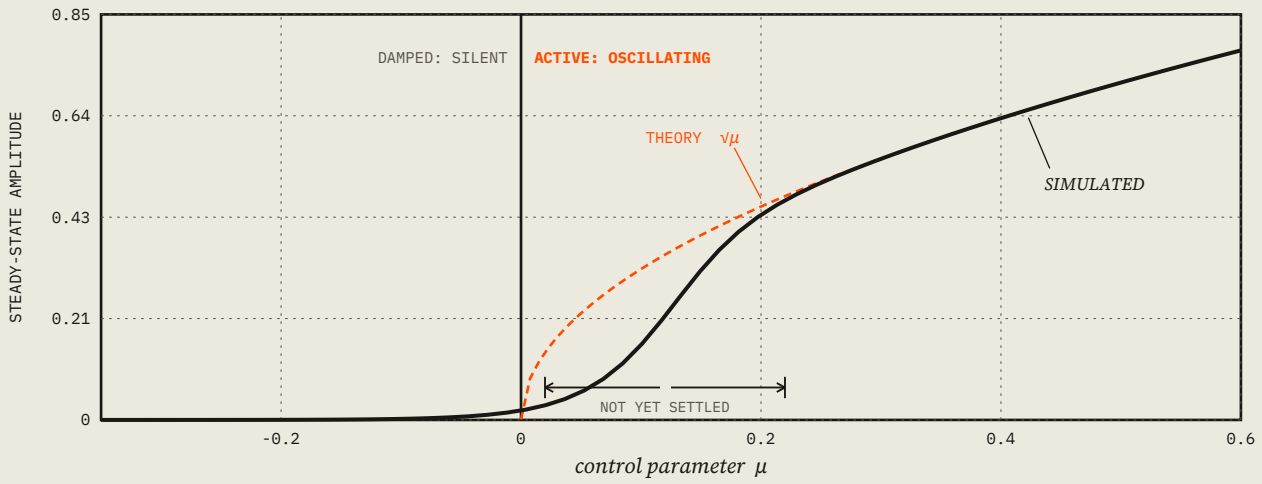


fig. 5: Sixty oscillators integrated from rest with no input, one per value of the control parameter. Below threshold every one decays to silence; above it they settle onto a limit cycle, a steady oscillation, whose amplitude follows the square root of the distance past threshold. The simulated points sit below the theory just past threshold, and that gap is not a plotting artefact: the time constant there diverges as one over the control parameter, so a fixed-length run has not finished growing. It is the same critical slowing down measured on the next page, showing up here as a visible bias.

COMPRESSION

Hold the same oscillator exactly at the critical point and drive it. On resonance the fast rotation drops out and what is left is the amplitude equation $dr/dt = \mu r - r^3 + F$. A passive resonator answers proportionally: double the drive, double the response. The critical oscillator answers as the cube root of the drive [8]. A cube-root law maps 120 dB of acoustic range into 40 dB of response, and it is why quiet sounds receive enormous gain and loud ones almost none. The power law is the settled part; its biological implementation is still under investigation.

RESPONSE AGAINST DRIVE, LOG-LOG

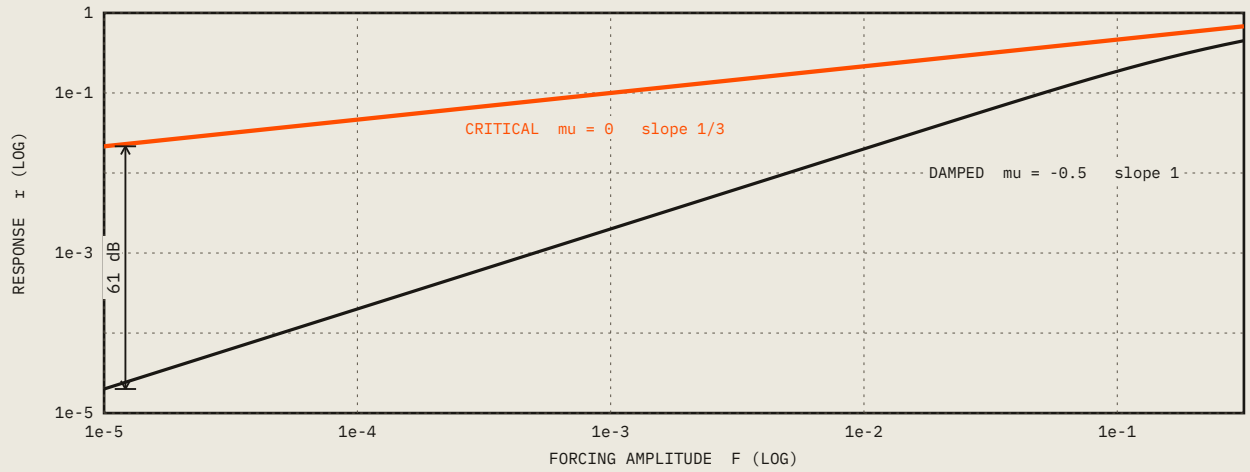


fig. 6: Steady state solved exactly rather than integrated, for the reason in the caution below. Fitted log-log slopes: 0.333 at the critical point against a theoretical one third, and 0.987 for the damped oscillator against a theoretical one. At the quietest drive the critical oscillator answers 61 dB louder than the passive one.

CAUTION

Settling time at the critical point diverges as F^(-2/3), which is critical slowing down. At the quietest drive on this plot the time constant is 718 model-seconds against 0.72 at the loudest, a factor of 1000. A fixed-length integration silently under-reports the quiet end; the first attempt at this figure measured a slope of 0.69 for exactly that reason. The slowness is not a numerical nuisance either. It follows from operating at the critical point, and it may be one reason faint sounds take a moment to become audible.

SPONTANEOUS EMISSIONS

A real cochlea is not one oscillator but a graded array of them, and no two are identical. That irregularity is called cochlear roughness, and it is why emission patterns differ from ear to ear. Here are sixty oscillators with smoothly graded frequencies, small random scatter in both frequency and control parameter, weak coupling to their neighbours, and no input whatsoever.

SIXTY OSCILLATORS, NO INPUT

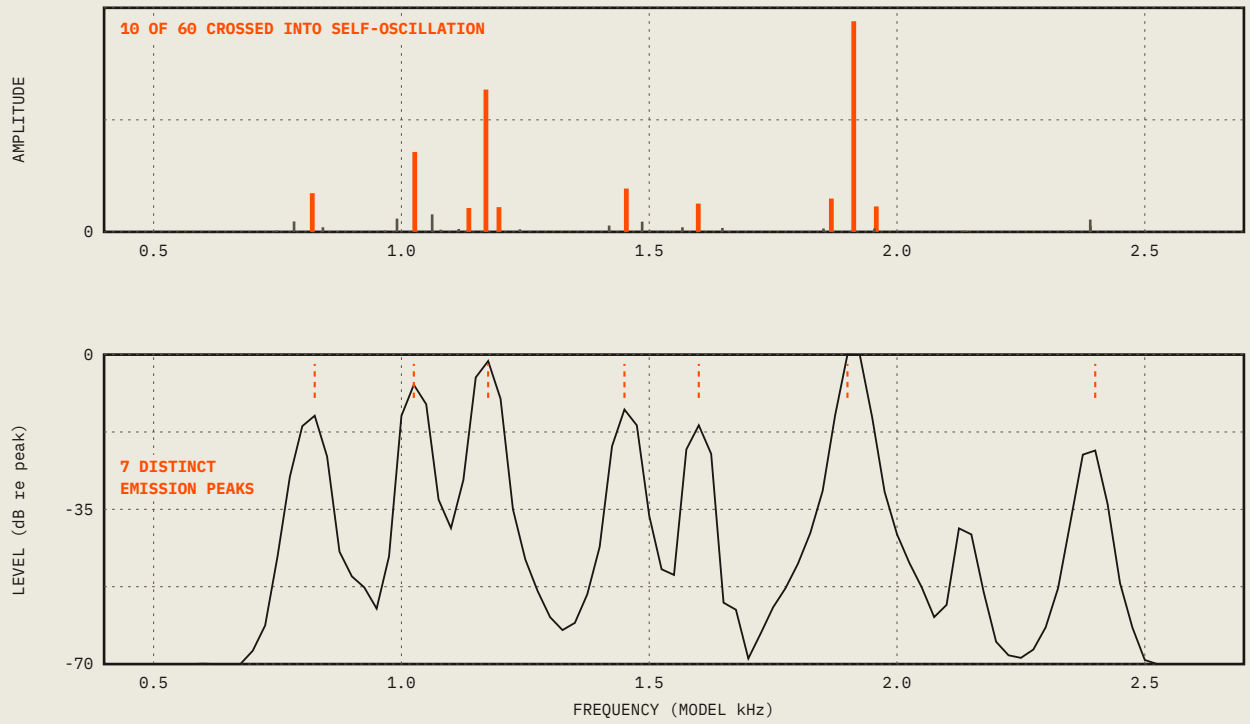


fig. 7: Top: the oscillators that crossed into self-oscillation, in accent. Bottom: the sound in the ear canal, which is the sum of everything. It holds a handful of discrete, stable tones, irregularly spaced but never too close together, which is the observed character of real spontaneous emissions. The floor on that spacing is the coupling: rerunning the identical bank with the coupling removed gives 11 peaks with a minimum spacing of 0.075, against 7 peaks at 0.150 here. Neighbours that fall too close together entrain and emit as one.

What this establishes is narrow: a heterogeneous coupled-oscillator field with no input produces a small number of discrete, stable, irregularly spaced emissions, and the coupling sets the minimum spacing. It does not establish where real emissions come from. Distortion-product emissions, the ones evoked by two tones, vary by as much as 20 dB from one frequency to the next, and that ripple is attributed to interference between two sources in the cochlea rather than to neighbouring oscillators locking together [4]. For the player the consequence is the same either way: the phantom is not equally strong at every pitch, and where it is strong differs between listeners. Each listener hears the instrument through the irregularities of their own cochlea.

TWO MECHANISMS

Look again at the nonlinear term in the normal form. $|z|^2 z$ is purely cubic. Drive a critical oscillator with two tones and it must produce intermodulation at $2f_1 - f_2$, the cubic difference tone, which is the strongest distortion product at moderate levels and the one this instrument's CDT mode places deliberately. So the model that explains why some ears emit spontaneous tones also explains why every healthy ear hears a cubic difference tone. Both come from the same term.

TWO TONES IN, DISTORTION PRODUCTS OUT

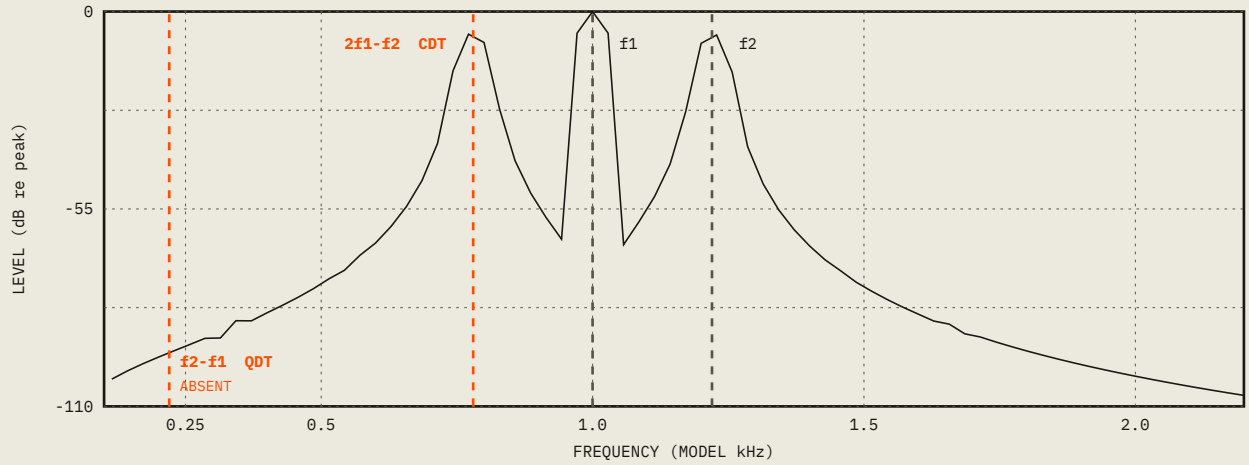


fig. 8: A single near-critical oscillator (mu = -0.02) whose characteristic frequency is the CDT place, driven by two tones at the clinical ratio of 1.22. The cubic difference tone arrives at -6.3 dB. The quadratic difference tone does not arrive at all: it sits at -91 dB, in the numerical floor. The level is not a prediction: it is set by the drive amplitude, the detuning and the oscillator answering at its own place. Goldstein's -15 dB re primaries [1], which the panel draws, is the listener's figure; the result here is the parity, not the number.

The missing QDT is the most useful result in this document. The normal form's only nonlinearity is cubic, so it can generate odd-order intermodulation and nothing else. Real ears produce strong quadratic difference tones, which means they carry an asymmetric nonlinearity this model does not have: the asymmetry of the hair cell's own transduction curve, which is exactly what the a_2x^2 term on page 8 was standing in for.

WHY THERE ARE TWO MODES

The two panel modes drive two different mechanisms. QDT Complex relies on transducer asymmetry, which needs many partials and rewards phase locking. CDT Pair relies on the critical oscillator, which needs only two tones and holds a nearly constant level relative to its primaries. They have different growth laws, different ratio sensitivities and different audibility floors [1]. For a pair the split is exact: the quadratic term makes only f2 − f1 and the cubic term makes only 2f1 − f2. The ear has both mechanisms; the Mode switch chooses which one to drive.

THE RATIO WINDOW

In CDT mode the phantom sits at $2f_1 - f_2$, so the interval between the two primaries decides the pitch of the tone the listener's ear generates, and the three-note chord that results. Kendall, Haworth and Cadiz [3] tabulated the ratios at which that chord is recognisable. They are simple integer relationships, and the Ratio control carries a soft detent at each one.

WHERE THE RATIO CONTROL IS USEFUL

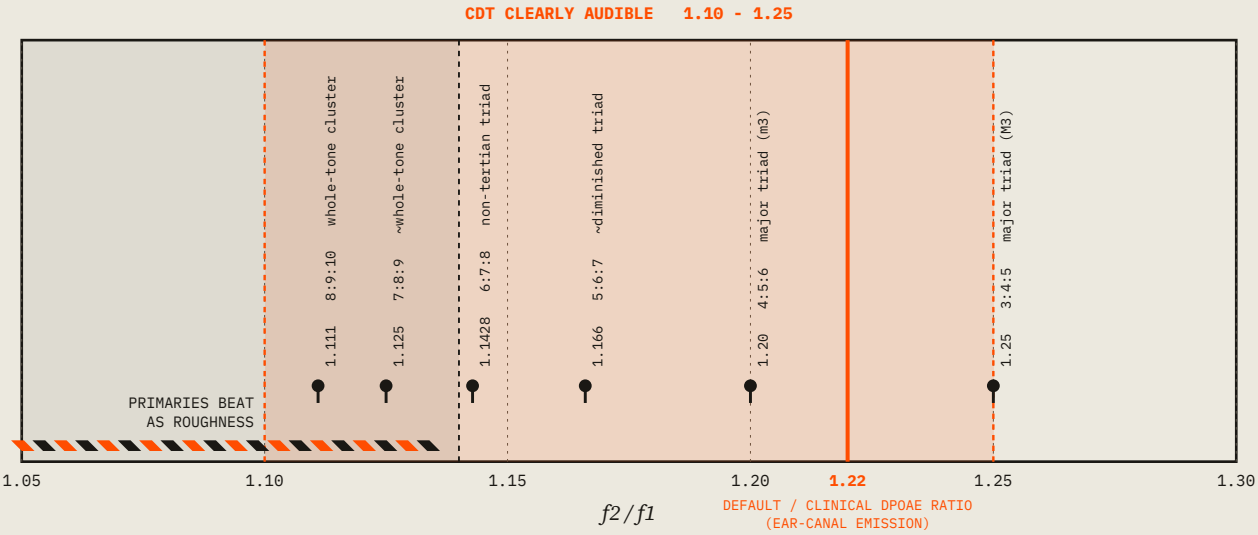


fig. 9: The whole useful travel of the Ratio control, on one axis. Below 1.14 the two primaries are close enough to beat against each other and the roughness competes with the phantom; above 1.25 the distortion product drops sharply. The detents are the integer chords. The default, 1.22, is the ratio at which the emission measured in the ear canal (the DPOAE used in hearing tests) is largest [4]. The tone you hear is strongest lower, near 1.10, and the panel draws it 12 dB weaker at 1.22 than at 1.10.

f2/f1	CDT : f1 : f2	CHORD	NOTE
1.111	8:9:10	whole-tone cluster	in the rough zone, under 1.14
1.125	7:8:9	~whole-tone cluster	in the rough zone, under 1.14
1.1428	6:7:8	non-tertian triad	
1.166	5:6:7	~diminished triad	
1.20	4:5:6	major triad (m3)	inside the clearly-audible window
1.25	3:4:5	major triad (M3)	top of the window; drops sharply above

MARYANNE AMACHER

Maryanne Amacher (Kane, Pennsylvania, 1938 - Rhinebeck, New York, 2009) studied with George Rochberg and Karlheinz Stockhausen, did graduate work in acoustics and computer science at Illinois, and spent the rest of her life building music out of the room, the building and the listener rather than out of instruments. City Links (from 1967) piped live sound between distant places; Music for Sound-Joined Rooms (from 1980) staged a piece across the architecture of a house; the Mini Sound Series (from 1985) serialised that idea like television. She made the storm for John Cage's Lecture on the Weather and Close Up for his Empty Words, and scores for Merce Cunningham, Torse among them.



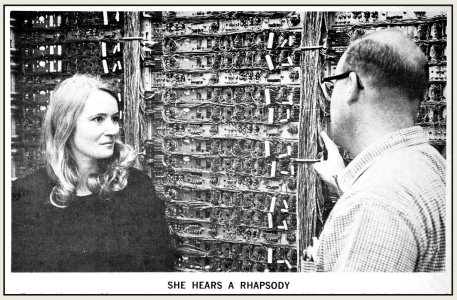
MARYANNE AMACHER AT THE DESK, 1970s
PHOTOGRAPHER NOT YET CREDITED

From the 1970s on she composed with what she called ear tones: the distortion products a healthy cochlea generates when it is driven by two or more tones at once. She made them her musical material, systematically and for decades, and she was explicit that the listener was part of the instrument. The 1999 Tzadik record carries the method in its subtitle. The effect itself is old: the violinist Giuseppe Tartini described it in 1754 and called it the third sound [14].

She played her pieces loud, through speakers. Her note on the sleeve asks for a high level:

> *"When played at the right sound level, which is quite high and exciting, the tones in this music will cause your ears to act as neurophonic instruments that emit sounds that will seem to be issuing directly from your head."*

That is history, not advice. This instrument asks for the lowest level at which the phantom is audible (page 16 says why), and the effect is audible well below the level she asked for.



CITY-LINKS, WBFO BUFFALO, 26-27 MAY 1967:
28 HOURS ON FIVE TELEPHONE LINES.
CAMPUS PAPER CLIPPING, CREDIT NOT YET ESTABLISHED

SOUND CHARACTERS

(MAKING THE THIRD EAR)

Maryanne Amacher. Tzadik TZ 7043, 1999. Opens with "Head Rhythm 1" and "Plaything 2". Followed by Sound Characters 2 (Making Sonic Spaces), Tzadik, 2008.

ThirdEar adds what researchers measured after her record. Kendall, Haworth and Cadiz [3] set out in 2014 to reproduce her method with a synthesis engine and a listener model, and their ratio table is on the panel as detents. Pressnitzer and Patterson's vector-sum account [2] is why every oscillator here starts in phase, and their alternating-phase control is a switch on the front. The meter under the note exists because her claim, that the melody is in the listener and not in the signal, is a claim a spectrum analyser can check, and the cochlea view draws it on Greenwood's map of the membrane [9]. The three After Amacher presets follow her documented technique, not her pieces: a low D-sharp at 77 Hz that she said sounds "from many different intervals", a pair closing on a third until it "turns into a shape" [11], and an ostinato after Head Rhythm 1 [12].

LISTENING

WORK	WHERE TO START	WHY
Sound Characters (1999)	track 1, at a moderate level; raise it only until the low line appears	the interlocking high patterns; the low line under them is produced in the ear
Sound Characters 2 (2008)	any track	the same method in rooms rather than heads
Kendall, Haworth, Cadiz (2014)	the audio examples that accompany the paper	the CDT triads the Kendall presets reproduce
Estelle Schorpp, In My Ears (for Maryanne), LINE_155, 2024	synaptic choral I; SuperCollider, pairs of close pure tones, 1 to 4 kHz	the other mode: a close pair puts 2f1-f2 a third of an octave under f1, so the phantom stays high

EXPOSURE

This instrument works by driving a nonlinearity in the listener. It emits narrowband tones at moderate to high level, held for as long as a note is held, and the effect is strongest exactly where hearing is most sensitive. The effect is an illusion, but the stimulus is not: it is ordinary acoustic energy delivered to a listener's ear, and the ordinary exposure limits apply to it. They are listed below with their sources.

LIMIT	VALUE	SOURCE
Occupational exposure	85 dBA, 8 h time-weighted average. The 3 dB exchange rate means every additional 3 dB halves the allowable duration.	NIOSH 1998, pub. 98-126
Recreational allowance	80 dBA over 40 h per week for adults (mode 1); 75 dBA for sensitive listeners, children among them (mode 2).	WHO-ITU H.870, modes 1 and 2
Rest breaks	Take them: narrowband tones concentrate energy on a small region of the cochlear partition. No published standard sets an interval for this; the limits above are dose-based.	lab practice, not a standard
A-weighting	For tones between 800 Hz and 8 kHz the correction stays within -1.1 to +1.3 dB, so dB SPL and dBA agree. CDT tones and Rove read lower in dBA (C3 -11, C2 -19, 16 kHz -7 dB): dB SPL is the safe dose.	IEC 61672 weighting curve

WARNING

Work at the lowest level at which the phantom is audible. Raising the primaries raises the quadratic product about 2 dB for every dB [13], so the illusion gets stronger long before the stimulus gets loud. A phantom that fades during a long session is the mechanism adapting to a steady stimulus: raise Rove depth (QDT mode), or stop for a few minutes. Do not turn the level up to bring it back. If your ears ring or feel full afterwards, the level was too high; stop for the day. If you have tinnitus or sensitive hearing, keep sessions short and quiet.

WHAT THE PANEL DOES ABOUT IT

- Level starts at -15 dB and stops at -6 dB. Nothing on the panel raises it in one action: a preset, the reset in the settings menu and a double-click on LEVEL keep the lower of the two values.
- The synth is held under the safety clip. The conditions strip says when it holds a note under Level, and when Norm sets a note above it.
- Under Norm loud the boost stops at 6 dB and fades out between 4 and 8 kHz, so high tones, which A-weighting rates quiet, are not made louder in dB SPL.
- FIND LEVEL on the bench finds your own threshold and sets Level a little above it, never higher than where you started.

SOURCES

[1] Goldstein, J. L. "Auditory nonlinearity." Journal of the Acoustical Society of America 41, 1967. Growth laws and audibility floors for both products; the ratio-sensitivity difference that separates the two modes; the cancellation method.

[2] Pressnitzer, D. and Patterson, R. D. "Distortion products and the perceived pitch of harmonic complex tones." Physiological and Psychophysical Bases of Auditory Function, 2001. The vector-sum account, the 3 dB per pair doubling, and the alternating-phase experiment behind the Start phase control.

[3] Kendall, G., Haworth, C. and Cadiz, R. "Sound synthesis with auditory distortion products." Computer Music Journal 38(4), 2014, pp. 5-23. Table 1 is the source of the six ratio detents; Example 1a.

[4] Mauermann, M., Uppenkamp, S., van Hengel, P. W. J. and Kollmeier, B. "Evidence for the distortion product frequency place as a source of DPOAE fine structure in humans. I." Journal of the Acoustical Society of America 106(6), 1999, 3473-3483. The fine-structure ripple of distortion-product emissions, and the largest emission near a ratio of 1.2 to 1.225; it argues a two-source origin, not the mechanism of the oscillator model.

[5] Hudspeth, A. J., Julicher, F. and Martin, P. "A critique of the critical cochlea: Hopf, a bifurcation, is better than none." Journal of Neurophysiology 104(3), 2010, 1219-1229. The normal form used in the model pages.

[6] Kemp, D. T. "Stimulated acoustic emissions from within the human auditory system." JASA 64, 1978. The original observation that a healthy ear emits sound.

[7] Penner, M. J. and Zhang, T. "Prevalence of spontaneous otoacoustic emissions in adults revisited." Hearing Research 103, 1997, 28-34. How often adult ears have a spontaneous emission.

[8] Eguiluz, V. M., Ospeck, M., Choe, Y., Hudspeth, A. J. and Magnasco, M. O. "Essential nonlinearities in hearing." Physical Review Letters 84, 2000, 5232. The one-third power law and the cubic difference tone from a critical oscillator.

[9] Greenwood, D. D. "A cochlear frequency-position function for several species, 29 years later." JASA 87(6), 1990, 2592-2605. The cochlea view.

[10] Houtsma, A. J. M. and Goldstein, J. L. "The central origin of the pitch of complex tones." JASA 51, 1972. What Split leaves to the brain.

[11] Oteri, F. J. "Extremities: Maryanne Amacher in conversation with Frank J. Oteri." NewMusicBox, New Music USA, 2004. The 77 Hz D-sharp and the third.

[12] Kahrs, N. Review of A. Cimini, Wild Sound: Maryanne Amacher and the Tenses of Audible Life (Oxford University Press, 2021). Music Theory Online 28(4), 2022. Head Rhythm 1's ostinato.

[13] Zwicker, E. "Different behaviour of quadratic and cubic difference tones." Hearing Research 1, 1979, 283-292. The quadratic difference tone's growth of about 2 dB per dB.

[14] Tartini, G. Trattato di musica secondo la vera scienza dell'armonia. Padua, 1754. The third sound.

[15] Fastl, H. and Zwicker, E. Psychoacoustics: Facts and Models, 3rd ed. Springer, 2007. The level law of the quadratic difference tone, and listeners for whom it does not hold.

COMPANION MATERIAL

The model figures are recomputed from notebook 01, which you can read at databurn.org/thirdear/01-phantom-in-silico.html, and notebook 03 on the Hopf oscillator (not yet published). When this note is built, every model figure, every control range and default, the presets and the bench's file format are recomputed or read from the plugin's source and its table of measured constants, and the build stops if they disagree. The constants' sources are the ones listed above.