

# THE SCIENCE BEHIND SOUNDBREATH

How Personalized Audio Works —  
What the Research Shows,  
What We Built On It

*Research Companion Paper*

February 2026

---

*“Your brain already knows how to do this.  
We are just sending it a signal it recognizes.”*

This document explains the published research that SoundBreath is built on.

Where evidence is strong, we say so. Where we have innovated beyond published findings, we label it.

Contact: [theashzero@gmail.com](mailto:theashzero@gmail.com)

## 1. What This Paper Is For

You are using — or considering — SoundBreath. You may have questions: Is there real science behind this? Why does personalized audio matter? What is actually happening when I put on headphones and listen?

This paper answers those questions by walking you through the published research behind SoundBreath, what we have added, and where current knowledge stops.

Every claim is labeled with its evidence strength. Every study we reference has a PubMed identifier (PMID) so you can look it up yourself. Where our contribution goes beyond published research, we tell you.

## Contents

<b>1</b>	<b>What This Paper Is For</b>	<b>1</b>
<b>2</b>	<b>Why Generic Audio Does Not Work the Same for Everyone</b>	<b>3</b>
<b>3</b>	<b>Your Voice Tells More Than Your Words</b>	<b>3</b>
3.1	What SoundBreath Extracts . . . . .	3
3.2	Two Readings, One Gap . . . . .	4
<b>4</b>	<b>Your Constitutional Baseline — What “Normal” Looks Like for You</b>	<b>4</b>
4.1	The Gap — Where You Are vs. Where You Belong . . . . .	4
<b>5</b>	<b>How Personalized Audio Actually Works</b>	<b>5</b>
5.1	Frequency-Following Response . . . . .	5
5.2	Why Multiple Frequencies Matter . . . . .	5
5.3	Your Voice as the Carrier . . . . .	5
<b>6</b>	<b>Beyond Sound — Why Multi-Sensory Matters</b>	<b>6</b>
6.1	The Visual Layer . . . . .	6
6.2	The Breath Layer . . . . .	6
6.3	The Olfactory Dimension . . . . .	7
<b>7</b>	<b>What Happens Over 30 Days</b>	<b>7</b>
7.1	Day 1: Where Are You? . . . . .	7
7.2	Day 15: How Have You Moved? . . . . .	7

7.3	Day 30: What Has Become Visible? . . . . .	8
<b>8</b>	<b>What We Know, What We Hypothesize, and What We Do Not Know</b>	<b>8</b>
8.1	What We Know (Tier 1–2: Strong to Moderate Evidence) . . . . .	8
8.2	What We Hypothesize (Tier 3–4: Experimental to Our Innovation) . . . . .	8
8.3	What We Do Not Know . . . . .	9
<b>9</b>	<b>How to Use This Information</b>	<b>9</b>
	<b>References</b>	<b>10</b>

## 2. Why Generic Audio Does Not Work the Same for Everyone

You have probably tried a meditation app. You pressed play. A calm voice told you to breathe. Maybe it helped. Maybe it did nothing. Maybe you felt vaguely irritated and concluded that “this stuff doesn’t work for me.”

There is a clear reason for that.

A meta-analysis of binaural auditory beats — sound designed to influence brainwave patterns — found a medium effect size ( $g = 0.45$ ) across studies of cognition, anxiety, and pain perception [PMID: 30073406]. In plain language: on average, across large groups of people, binaural audio produces a measurable effect. Not huge, not negligible. Medium.

But “on average” hides something important. An average effect of 0.45 means some people respond strongly, some respond moderately, some do not respond at all, and a few respond in the wrong direction. The average tells you nothing about which category *you* fall into.

This is not a failure of the audio. It is a failure of the approach. The same 10 Hz alpha track cannot possibly work the same way for a person whose brain is already running high alpha (they do not need more) and a person whose brain is alpha-deficient (they need exactly this). One person gets redundancy. The other gets precisely what was missing.

SoundBreath was built on a simple premise: if you match the audio to the individual, the response should be more consistent than if you play the same track for everyone. The rest of this paper explains how that matching works and what it is built on.

## 3. Your Voice Tells More Than Your Words

When SoundBreath asks you to record a voice sample, it is not listening to what you say. It is measuring the *music* of your voice — the pitch, rhythm, energy, and spectral patterns that shift with your emotional and physiological state in ways you cannot consciously control.

This is not interpretation; it is measurement grounded in published research that has established specific correlations between acoustic features and internal states. Jitter (micro-variations in pitch), standard deviation of fundamental frequency (how much your pitch moves), speech speed, and pause durations all correlate significantly with emotional state as measured by clinical rating scales [PMID: 38089742].

Think of it like a blood pressure reading. Your blood pressure reveals cardiovascular state whether or not you feel anything. Similarly, your voice reveals nervous system state whether or not you are aware of it. When you say “I’m fine” but your voice carries the acoustic signature of stress — elevated pitch variability, compressed energy contours, altered harmonic ratios — the measurement captures what the words miss.

### 3.1 What SoundBreath Extracts

The system extracts several features from your voice. None of these require you to say anything specific. The content of your words is not recorded, stored, or analysed.

**Fundamental frequency (F0):** The base pitch of your voice, measured in Hertz. This varies with arousal and emotional state. A nervous system under stress produces different F0 patterns than one at rest.

**Spectral centroid:** Where the energy is concentrated across the frequency spectrum. A bright, tense voice has a high spectral centroid. A low, relaxed voice has a lower one.

**Harmonic-to-noise ratio:** How “clean” versus “breathy” or “rough” your voice sounds. This correlates with vocal cord tension, which in turn reflects autonomic state.

**Energy contours:** How loudly and softly you speak across the duration of the sample. The pattern of energy rise and fall carries information about activation level.

These features are combined to create an acoustic fingerprint of your current state — not a label like “anxious” or “calm,” but a multi-dimensional measurement of where your nervous system is actually sitting.

### 3.2 Two Readings, One Gap

SoundBreath captures voice at two moments. The first is your resting voice — how you sound when nothing in particular is happening. The second is your activated voice — how you sound when prompted to describe something that stirs a specific emotional response.

The difference between these two readings is informative. It reveals not just where your nervous system sits at rest, but how it *moves* when activated. Two people might sound identical at rest but diverge completely under activation. That divergence pattern is unique to you and directly shapes the audio that the system generates.

## 4. Your Constitutional Baseline — What “Normal” Looks Like for You

Not everyone’s nervous system is configured the same way. Some people are naturally more oscillatory — their system moves through states quickly, responds rapidly, and tends toward variability. Others are naturally more steady — their system resists perturbation, recovers slowly, and tends toward consistency. Neither is better. They are different configurations with different strengths and different vulnerabilities.

SoundBreath maps your constitutional baseline through a structured assessment. This is not a personality test. It measures stable physiological tendencies that have demonstrated genetic correlates in peer-reviewed research.

A genome-wide study of 262 individuals found 52 single-nucleotide polymorphisms (SNPs) that significantly differentiated constitutional types, with principal component analysis correctly classifying individuals irrespective of ancestry [PMID: 26511157]. A separate study showed that drug metabolism rates (CYP2C19 polymorphisms) correlate with constitutional phenotype — meaning that constitutional type is reflected in measurable biochemistry, not just self-report [PMID: 20015960].

### 4.1 The Gap — Where You Are vs. Where You Belong

Your constitutional baseline is where your system naturally tends. Your current state is where it actually sits — shaped by stress, sleep, habits, and the accumulated weight of how life has asked you to respond.

The distance between these two — the gap — is the target: how far your system has drifted from its natural configuration, and in which direction.

Think of it like a thermostat. Your constitutional baseline is the temperature your house is set to. Your current state is the actual temperature in the room. The gap tells the system whether to heat or cool — and by how much. A house set to 22°C that reads 18°C needs warming. The same house reading 26°C needs cooling. Same house, opposite interventions.

This is what makes the audio personal. Your session is not “relaxation music.” It is a calibrated signal intended to nudge your specific nervous system configuration in a specific direction.

## 5. How Personalized Audio Actually Works

### 5.1 Frequency-Following Response

When your brain is exposed to a rhythmic auditory signal — a steady pulse at a specific frequency — neural oscillations tend to synchronize with that external rhythm. This is called the frequency-following response (FFR). It is not something you choose to do; it occurs below conscious attention.

Research has documented this effect across multiple frequency bands. A meta-analysis found that binaural audio produces measurable changes in EEG (brainwave) activity, with effects on anxiety, cognition, and pain perception [PMID: 30073406]. Multi-band stimulation — using multiple frequencies simultaneously rather than a single tone — co-modulates perception more effectively than single-frequency stimulation [PMID: 25225394].

**Honest caveat:** The reliability of auditory entrainment varies across studies. Some experiments show clear brainwave synchronization; others show inconsistent or weak effects. A systematic review raised concerns about small sample sizes, inadequate controls, and heterogeneous methods across studies [Ingendoh et al., 2023]. We do not treat auditory entrainment as a guaranteed mechanism. We treat it as a well-supported hypothesis with moderate evidence — strong enough to build on, honest enough to acknowledge the gaps.

### 5.2 Why Multiple Frequencies Matter

Your brain does not operate at one frequency. At any given moment, there is a mix of activity across delta (slow, deep), theta, alpha, beta, and gamma (fast, active) bands. Your emotional and cognitive state is reflected in the *ratio* between these bands, not in any single one.

Research confirms this: EEG classification using frequency band ratios (beta/gamma, theta/alpha, theta/beta) significantly outperforms single-band analysis for distinguishing emotional states [PMID: 34455110].

SoundBreath generates audio using multiple frequency bands whose relative emphasis reflects your individual profile. The mix is adjusted to address imbalances rather than applying a fixed, one-size-fits-all pattern. The audio is not a single dial. It is an equalizer, adjusted to match what your system needs.

### 5.3 Your Voice as the Carrier

Here is where our innovation begins.

Your fundamental frequency (F0) — the base pitch of your speaking voice — is unique to you.

It reflects the physical dimensions of your vocal apparatus and, by extension, something about your neural-acoustic signature.

Published research shows that the brain responds more strongly to familiar, personally relevant signals. In a study of 66 patients with disorders of consciousness, responses to the patient's own name spoken by a familiar voice were significantly stronger than responses to neutral stimuli, and this response predicted recovery outcomes [PMID: 25880206]. The brain preferentially processes signals it recognizes.

We extend this principle to the audio's frequency architecture. The carrier frequency — the tonal foundation of your session — is derived from your own voice. This means the audio resonates with your individual acoustic signature.

**Evidence tier:** The familiar-voice advantage in neural processing is well-established (Tier 1–2). Our extension from content-level familiarity (recognizing a known voice) to structural-level familiarity (tuning the carrier to your vocal harmonics) is our innovation (Tier 4). We believe it has a sound rationale, but it has not been independently tested. We are transparent about this distinction.

## 6. Beyond Sound — Why Multi-Sensory Matters

SoundBreath is primarily an audio experience. But the science behind it extends to a broader principle: the brain responds more strongly when multiple sensory channels are engaged simultaneously.

A Cochrane review — the gold standard of evidence synthesis — confirmed that multimodal sensory stimulation outperforms single-modality stimulation [PMID: 11969582]. A separate randomized controlled trial demonstrated that combined sensory stimulation improves arousal and consciousness scores versus control conditions [PMID: 23473238].

This is why SoundBreath combines audio with synchronized visual geometry and breath rhythm rather than delivering audio alone. Each channel reinforces the others.

### 6.1 The Visual Layer

The geometric patterns you see during a session are not decorative. They employ optical illusion principles — peripheral drift, contrast patterns, moiré interference — that engage sustained visual attention without requiring conscious effort. When the visual rhythm matches the audio rhythm, the brain tends to process both inputs as a single unified experience rather than two separate streams — a form of cross-modal binding.

### 6.2 The Breath Layer

The visual geometry expands and contracts at a calculated rate. When you match your breath to this rhythm, you create a closed sensory-motor loop: auditory input, visual input, and respiratory motor output all aligning to the same frequency. This is not relaxation breathing. It is synchronized engagement, and changes in heart rate variability and vagal tone of this kind have been confirmed in a systematic review and meta-analysis of 223 studies on voluntary slow breathing, which found increases in vagally-mediated HRV during sessions, immediately after, and following multi-session interventions [PMID: 35623448].

## 6.3 The Olfactory Dimension

There is one sensory channel with a unique property. The olfactory pathway — your sense of smell — is the only sensory system that reaches the brain's cortex without first passing through the thalamus, the brain's central relay station [PMID: 24409147, PMID: 25698978]. Every other sense — hearing, vision, touch — is filtered through the thalamus before it reaches cortical processing.

This means smell has a more direct path to the limbic system (amygdala, entorhinal cortex) than any other sense. Specific aromatic compounds have demonstrated measurable neural effects:

Rosemary (specifically the compound 1,8-cineole) enhances cognitive performance in healthy subjects, with task accuracy significantly correlated with plasma levels of the absorbed compound — suggesting a pharmacological mechanism, not just a pleasant smell [PMID: 23983963].

Linalool (from lavender) produces anxiolytic effects through GABAergic transmission — the same receptor system targeted by benzodiazepine medications. The effect is triggered by olfactory input and confirmed by the fact that it is blocked by flumazenil, a benzodiazepine antagonist [PMID: 30405369].

This research informs the broader SoundBreath ecosystem, where olfactory calibration complements the audio-visual experience.

## 7. What Happens Over 30 Days

A single session is a measurement. A trajectory is information.

SoundBreath maps your state at three points: Day 1, Day 15, and Day 30. This creates a trajectory — a path through your personal state-space that can reveal patterns a single snapshot would miss.

### 7.1 Day 1: Where Are You?

Your constitutional baseline is assessed. Your voice samples are captured and analysed. The gap between baseline and current state is calculated. Your personalized audio and visual intervention is generated based on your specific profile.

This is not a generic assignment. Key parameters of the session are derived from your measurements. Two people completing Day 1 on the same afternoon receive different sessions.

### 7.2 Day 15: How Have You Moved?

New voice samples are compared against your Day 1 baseline. Has the gap narrowed? Has it shifted direction? Has something new emerged?

Your intervention is regenerated. The audio evolves with you. If your system has shifted, the audio shifts to match — still targeting the gap, but the gap itself may have changed shape.

### 7.3 Day 30: What Has Become Visible?

Final voice samples complete the trajectory. The full 30-day path is analysed: where you started, where you moved, what patterns emerged.

Success is not a number on a scale. Success is: have previously inaccessible states become available? Has the gap between where your system sits and where it naturally belongs narrowed? Has the primary pattern — the thing that was actually driving the surface-level experience — become visible to you?

A user whose surface label stayed the same from Day 1 to Day 30 might appear unchanged. But if the underlying composition shifted — if what was hidden became clear — something significant occurred. The system revealed its own architecture to you. That is access, not improvement.

## 8. What We Know, What We Hypothesize, and What We Do Not Know

We label every element of the system with its evidence strength. Here is the honest summary.

### 8.1 What We Know (Tier 1–2: Strong to Moderate Evidence)

**Your voice carries measurable information about your emotional and physiological state.** Multiple studies, clinical rating scale correlations, acoustic feature extraction is established science [PMID: 38089742].

**Constitutional types have genetic correlates.** Genome-wide studies, SNP clusters, drug metabolism correlations. The constitutional framework has biological validation [PMID: 26511157, PMID: 20015960].

**Binaural and isochronic audio can modulate brainwave activity.** Meta-analysis with medium effect size ( $g = 0.45$ ). Effects are real but variable across individuals [PMID: 30073406].

**Multimodal sensory stimulation outperforms single-modality stimulation.** Cochrane review and RCTs [PMID: 11969582, PMID: 23473238].

**Multi-band frequency ratios distinguish emotional states better than single bands.** EEG classification research [PMID: 34455110].

**The olfactory pathway reaches the cortex without thalamic relay.** Established neuroanatomy [PMID: 24409147].

**Specific aromatic compounds have measurable neural effects.** Rosemary-cognition correlation with plasma levels [PMID: 23983963]. Linalool anxiolytic via GABAergic pathway [PMID: 30405369].

### 8.2 What We Hypothesize (Tier 3–4: Experimental to Our Innovation)

**Matching audio frequencies to the individual's specific gap produces a more consistent response than generic audio.** This is the core hypothesis. It has strong logical support and builds on Tier 1–2 evidence, but has not been independently validated in a controlled trial. We are building the trajectory data to test it.

**Voice-derived carrier frequencies create stronger neural engagement than arbitrary carriers.** We extrapolate from the established finding that the brain responds preferentially to familiar signals. The extension from content-level familiarity to spectral-level tuning is our hypothesis, not a published finding.

**Targeting underrepresented aspects of a person's profile, rather than reinforcing what is already dominant, may produce more useful state shifts than random or redundant stimulation.** This idea is intuitively appealing and aligns with general neurostimulation principles, but has not been directly validated at the audio entrainment level.

**Tracking the gap trajectory over 30 days reveals patterns that single measurements miss.** We believe this strongly based on the logic of longitudinal measurement, but the way we interpret those trajectories is our own working model, not a published standard.

### 8.3 What We Do Not Know

We do not know the exact ceiling of audio-based state modulation. There are limits to what sound can do, and those limits vary by individual.

We do not know whether the personalization advantage holds for all constitutional types equally, or whether some profiles respond more than others. The trajectory data will reveal this over time.

We do not know how much of the subjective experience (time distortion, delayed trigger responses, altered emotional availability) is attributable to the audio personalization specifically versus the broader multi-sensory engagement plus daily practice. Disentangling these requires controlled studies that we intend to pursue.

We commit to reporting honestly what works, what does not, and for whom.

## 9. How to Use This Information

This paper is not required reading. You do not need to understand the science to benefit from the experience. But if you wanted to know what the system is built on — and whether it is built on anything real — this is the answer.

The research exists. The citations are verifiable. The evidence tiers are honest. Where we have gone beyond published findings, we have told you.

Your 30-day journey is not generic relaxation content. It is a calibrated intervention built on your specific measurements, evolving with your trajectory, targeting the distance between where you are and where your system naturally belongs.

Whether that matters to you is something you will discover by doing it.

---

*The goal is not to tell you what to feel.*

*The goal is to send your brain a signal it already knows —  
and let it remember what it forgot it could do.*

## References

All references below are indexed in PubMed. Enter the PMID number at <https://pubmed.ncbi.nlm.nih.gov/> to read the original study.

### Voice Biomarkers

Seifpanahi MS, et al. (2023). The association between depression severity, prosody, and voice acoustic features in women with depression. *TheScientificWorldJournal*, 2023:9928446.  
PMID: 38089742

### Constitutional / Genetic Profiling

Govindaraj P, et al. (2015). Genome-wide analysis correlates constitutional phenotypes with genetic markers. *Scientific Reports*, 5:15786.  
PMID: 26511157

Ghodke Y, et al. (2011). Constitutional phenotypes and CYP2C19 gene polymorphism associated with metabolic variability. *Evid Based Complement Alternat Med*, 2011:249528.  
PMID: 20015960

### Auditory Entrainment

Garcia-Argibay M, et al. (2019). Efficacy of binaural auditory beats in cognition, anxiety, and pain perception: a meta-analysis. *Psychological Research*, 83(2):357–372.  
PMID: 30073406

Ingendoh RM, et al. (2023). Binaural beats to entrain the brain? A systematic review of effects on EEG power and ERP phase. *PLoS ONE*, 18(5):e0286023. Cited for limitations and inconsistent results across entrainment studies.

Orozco Perez HD, et al. (2025). Frequency-following response modeling. *Scientific Reports*.  
DOI: 10.1038/s41598-025-88517-z

### Multi-Band EEG / Emotional States

Beta/gamma energy ratios classify emotions better than single bands.  
PMID: 34455110

Multi-band entrainment co-modulates perception.  
PMID: 25225394

### Multi-Sensory Stimulation

[Cochrane Review] (2002). Multimodal outperforms unimodal sensory stimulation.  
PMID: 11969582

[RCT] (2013). Multimodal sensory stimulation improves arousal scores.  
PMID: 23473238

### **Olfactory Pathway / Aromatic Compounds**

Olfactory direct projection anatomy — reaches cortex without thalamic relay.  
PMID: 24409147

Olfactory primary cortex access without first-order thalamic relay.  
PMID: 25698978

Moss M, Oliver L. (2012). Plasma 1,8-cineole correlates with cognitive performance following rosemary aroma exposure. *Ther Adv Psychopharmacol*, 2(3):103–13.  
PMID: 23983963

Harada H, et al. (2018). Linalool odor-induced anxiolytic effects in mice. *Front Behav Neurosci*, 12:241.  
PMID: 30405369

### **Respiratory / Heart Rate Variability**

Laborde S, et al. (2022). Effects of voluntary slow breathing on heart rate and heart rate variability: A systematic review and a meta-analysis. *Neurosci Biobehav Rev*, 138:104711.  
PMID: 35623448

### **Familiar Voice / Neural Processing**

Wang F, et al. (2015). Cerebral response to subject's own name showed high prognostic value in traumatic vegetative state. *BMC Medicine*, 13:83.  
PMID: 25880206