Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Acoustic Detection of Spaceflight Fluid Shift: A Consistent In-Flight Decline in Harmonics-to-Noise Ratio Across Three ISS Crew Members, with an Operational Model for Voice-Based Crew Health Monitoring

Oğuzhan Aydemir[1]

oguzhan.aydemir@nisantasi.edu.tr

İstanbul Nişantaşı Üniversitesi

Gökhan Aydemir[2]

Atasehir Acibadem Hospital

Article

Keywords: spaceflight fluid shift, vocal biomarkers, crew health monitoring, harmonics-to-noise ratio, exploration mission design, International Space Station

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10884904/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: Competing interest reported. Under the contract between the parties, TÜBİTAK UZAY and the Turkish Space Agency (TUA) could derive financial benefit from sales of the software used in this investigation, as disclosed in the participant consent form. Conflict of interest disclosures were filed with the NASA Institutional Review Board for all investigators. The authors declare no other competing financial or non-financial interests.

Abstract

Cephalad fluid redistribution is among the earliest and most consistent physiological consequences of orbital insertion, yet it is difficult to monitor in flight: the measurements that resolve it require consumables, hardware mass, or crew time that exploration-class missions cannot spare. Voice is a candidate readout — phonation depends on vocal-fold surface hydration and viscoelasticity, both of which fluid shift perturbs — but no controlled in-orbit vocal corpus has previously existed against which to test it.

We report such a corpus and the result of that test. Three International Space Station crew members performed a fixed vocal battery on a certified crew tablet, twice daily, across a pre-flight phase and mission days 3–14, yielding 44 sessions and 254 verified recordings with a pre-flight baseline acquired through an identical acquisition path. Every recording was retained in two parallel forms — the raw acquisition and the same material after a noise-suppression stage — and both were measured with a single Praat/Parselmouth pipeline.

Harmonics-to-noise ratio (HNR) declined between the pre-flight and in-flight phases in seven of the eight testable crew × vowel cells, median − 2.11 dB. At the recording level, with all three vowels pooled and each crew member's own pre-flight median subtracted, the decline is − 2.70 dB (23 pre-flight vs 96 in-flight recordings, $p = 0.019$), reaching significance within a single crew member (− 5.53 dB, $p = 0.011$, rank-biserial − 0.72); restricted to the two vowels sampled completely in every crew member it is 6 of 6 cells at − 2.47 dB ($p = 0.0038$). The effect is specific to HNR: shimmer and CPPS, which share its sensitivity to added acquisition noise, decline in only 2 and 3 of 8 cells respectively, so a drift in recording conditions between phases — which would have displaced all three together — does not explain it. A reduction in HNR is an increase in turbulent, aperiodic energy during phonation, the direction predicted by laryngeal mucosal fluid loading. The change presents as a sustained step between phases rather than a progressive in-flight trend, consistent with fluid shift reaching a new steady state within hours of insertion.

Because acoustic measurements are vulnerable to their acquisition chain, we established what that chain contributes before interpreting the result. Retaining both processing versions allowed this to be measured directly: the noise-suppression stage moves shimmer by − 3.35% and HNR by + 3.56 dB (paired, $n = 63$, $p \approx 10^{-11}$), while leaving F_0 and CPPS unchanged (CPPS $p = 0.94$). The phase contrast above is computed within a single processing version, so this term cancels; quantifying it is what licenses the physiological reading rather than undermining it. The same contrast applied to 135 voluntary cough recordings gives no consistent direction in any measure, which we report as evidence for task selection: sustained phonation is self-normalising, whereas cough acoustics depend on applied effort and instantaneous mouth-to-microphone geometry. We report two further chain properties for reuse: cough timing depends 9.90-fold on segmentation convention, and within-session dimensionless ratios are stable across processing versions (within-crew CV 11.3–15.0%).

The operational case is quantified: 18.2 s of crew recording time, zero consumables, zero incremental mass, and 624 B of downlink per session against 568 KiB of raw audio — a 932-fold reduction with on-board measurement. We give design requirements for follow-on missions in directly usable form: nine sessions per phase per crew member reach $d = 1.5$ at 80% power, against $d \geq 1.91$ permitted here, and pre-flight sessions cost nothing in flight resources. No blood or imaging reference was obtained, so the fluid-shift attribution is an interpretation consistent with the observed direction and specificity rather than a demonstrated mechanism.

1. Introduction

1.1 Crew health monitoring as a resource problem

Medical capability aboard a spacecraft is bounded by three resources that do not scale with mission duration: consumables, crew time, and communications [1, 6, 7]. On ISS these bounds are tolerable because resupply is routine and physician support is effectively real-time. On a lunar surface mission they tighten; on a Mars transit they become the design driver. Blood draws consume irreplaceable expendables and produce samples that cannot be returned. Crew time is the scarcest resource aboard any vehicle. And Earth-based diagnostic support, which today closes the loop within minutes, closes it within roughly 6–44 min round-trip at Mars distance, which removes the flight surgeon from the acute decision path entirely and shifts diagnostic authority onto the crew and onboard systems [2, 4, 5].

The consequence for instrumentation is specific: exploration missions need physiological observables that are *repeatable at high cadence*, *consume nothing*, and *degrade gracefully into autonomy*. Sensing modalities that satisfy these constraints are rare. Most non-invasive candidates either require dedicated hardware mass, or produce data volumes that are unattractive to downlink, or need expert interpretation that the light-time delay makes unavailable [2, 3].

1.2 Why voice

Voice is an unusual candidate because it satisfies the resource constraints almost trivially. Phonation is produced by tissue whose viscoelastic behaviour depends on hydration and ionic composition, driven by a neuromuscular system sensitive to systemic state, and coupled to an airway whose geometry responds to fluid distribution. Several of the documented adaptations of spaceflight — cephalad fluid redistribution beginning within hours of orbital insertion, plasma volume contraction of the order of 10–15% over the first week, and the tissue consequences of sustained headward loading — act on precisely these structures [10, 11]. On the ground, the dependence of phonation on hydration state is well established: induced dehydration and rehydration measurably shift phonation threshold pressure [12–14]. Voice is also the only physiological signal that can be sampled repeatedly, non-invasively, in seconds, using hardware already aboard for other purposes.

This yields a specific, testable prediction rather than a general expectation. If cephalad fluid redistribution loads the laryngeal mucosa and alters vocal-fold surface hydration, the vibratory

consequence is reduced efficiency of glottal closure, and the acoustic consequence is proportionally more turbulent and less periodic energy in the phonated signal — that is, a fall in harmonics-to-noise ratio, without a corresponding requirement that pitch or cepstral peak prominence move in any particular direction. Section 3.5 tests that prediction.

What has been missing is not the hypothesis but the *instrument characterization*. Prior analyses of astronaut speech have almost exclusively used mission communications audio: material recorded for operational purposes, with uncontrolled task, uncontrolled duration, uncontrolled acquisition path, and no pre-flight counterpart acquired through the same chain. Terrestrial voice corpora acquired under controlled tasks, by contrast, have supported a substantial reuse literature [20–23]. Such material can support descriptive observation; it cannot support the measurement-validity claims that a flight instrument requires.

1.3 Contribution of this paper

Our contributions are:

1. 1. **Acoustic detection of the fluid-shift response.** Harmonics-to-noise ratio declined between the pre-flight and in-flight phases in seven of eight testable crew × vowel cells (median − 2.11 dB; pooled recording level, all three vowels, − 2.70 dB, $p = 0.019$), with no comparable change in the other noise-sensitive measures, and no consistent change in any cough measure (Sections 3.5–3.6).
2. 2. **A controlled longitudinal in-orbit corpus** — fixed task battery, fixed recording duration, scheduled cadence, and a pre-flight baseline acquired through an identical acquisition path (Section 2). To our knowledge no comparable in-orbit vocal corpus exists.
3. 3. **Characterization of the acquisition and processing chain**, made possible by retaining both raw and processed forms of every recording — a design decision that is rarely taken and is what allows the result in (1) to be attributed to the crew rather than to the pipeline (Sections 3.2–3.4).
4. 4. **An operational resource model** in crew-seconds and downlink bytes, with the on-board-processing trade (Section 4).
5. 5. **Quantitative design requirements for follow-on missions**, expressed as sessions per phase per crew member at stated power (Section 5).

We state at the outset what this paper does not claim. No blood, urine, or imaging reference was obtained from any crew member. The attribution of the acoustic result to fluid shift is an interpretation supported by the direction, specificity and temporal form of the effect, not a demonstrated mechanism, and no diagnostic or individual clinical inference is made or supported.

2. Experiment Design and Data

2.1 Participants and phases

Three crew members of a short-duration ISS mission participated, referred to throughout as Crew M1, M2 and M3. No names, mission call-signs, agency crew codes, mission identifiers or calendar dates appear in this paper; timing is reported only as mission day relative to launch. Detailed demographics are withheld because, in a crew of three, age, sex and nationality in combination would be identifying.

Data were acquired in two phases: a pre-flight quarantine phase and an in-flight phase spanning mission days 3–14. The design is within-subject longitudinal; each crew member serves as their own reference. The on-orbit stay was extended beyond the planned duration, and the additional days were used for further sessions under the identical protocol (Section 2.3).

2.2 Task battery: designed and delivered

The protocol as manifested comprised **nine** recordings per session in three blocks of three: three voluntary coughs, three sustained vowels (/a/, /i/, /u/), and three forced-expiration manoeuvres intended as a spirometric surrogate. Recording duration was fixed by the acquisition application at 3.029 s, sampled at 16 kHz, single channel.

Table 1
Session task battery as designed, and as delivered.

Order	Task	Instruction	Delivered
1–3	Voluntary cough	Three separate coughs, one per recording	Yes
4	Sustained /a/	Comfortable pitch and loudness	Yes
5	Sustained /i/	Comfortable pitch and loudness	Yes
6	Sustained /u/	Comfortable pitch and loudness	Yes
7–9	Forced expiration	Three repetitions, spirometry-analogue	**No — zero sessions**

The forced-expiration block returned no usable data from any crew member in any session. The task was designed around a decibel-difference method validated pre-flight against hospital-grade spirometry in a terrestrial control group: ambient noise level is measured in dB, the crew member then performs a forced expiration into the tablet microphone, and FEV_1 is estimated from the net power difference (respiratory calibration − environmental calibration) on a natural-log scale. The method requires a quiet, controlled acoustic reference and a reproducible mouth-to-microphone geometry. Neither condition obtains aboard the station. Module spatial-average sound levels on ISS are approximately 62.9–63.0 dBA against an NC-50 module requirement, with continuous broadband noise dominated by the air-revitalization and thermal-control systems; even the crew quarters, among the quietest locations aboard, sit at 49–54 dBA with fans on high [8, 9]. A method whose observable is the decibel difference between a quiet reference and a forced expiration has, in that environment, no quiet reference to difference against — and free-fall provides no fixed body position to stabilize the source-to-microphone path. The task appears in the flight manifest and in the acquisition application; it appears in the returned inventory for no session.

We report this failure prominently rather than omitting it, because it is the single most transferable design lesson in the experiment (Section 5, R7): an acoustic task whose observable is defined against an *absolute* ambient reference did not survive the transition from a terrestrial clinic to a spacecraft cabin, while tasks whose observables are defined *within* the recording did. This is the same distinction, arrived at independently from an operational direction, that Sections 3.2 and 3.4 reach from a measurement direction.

Fixing recording duration is a deliberate design choice with consequences in both directions. It removes participant-controlled phonation length as a source of between-session variance, making sessions directly comparable; it forecloses maximum-phonation-time measures; and — as Section 3.3 shows — it produces recordings that are predominantly silence, which interacts strongly with the choice of transient segmentation convention.

2.3 Session inventory

Table 2
Sessions and mission-day coverage by crew and phase.

Crew	Phase	Sessions	Mission-day span	Distinct days
M1	Quarantine	2	−10 to − 9	2
M1	In-flight	16	+ 4 to + 13	8
M2	Quarantine	3	−59 to − 10	2
M2	In-flight	6	+ 3 to + 13	3
M3	Quarantine	3	−10 to − 9	2
M3	In-flight	14	+ 5 to + 14	9
All		**44**		

Three features of Table 2 constrain the analysis and are stated here rather than in Limitations because they shaped the design of Section 3.

First, the pre-flight window afforded substantially fewer sessions than the flight window for every crew member — the dominant limitation of the experiment, and the one Section 5 addresses quantitatively. Second, the M2 baseline spans mission day − 59 to − 10, i.e. it is not a tight quarantine-window baseline but a pre-flight series distributed over roughly seven weeks; M2 phase contrasts therefore carry a different temporal meaning from those of M1 and M3.

Third, **pre-flight coverage differs by version.** All three crew members completed pre-flight sessions and all are present in the processed version with sub-second acquisition timestamps. In the raw-acquisition version, however, the pre-flight measurements of one crew member are absent from the released deposit. Every phase comparison in this paper is therefore computed within the processed version, where all three crew members are represented; the raw version is used only for the paired version

comparison of Section 3.2, which does not involve the phase contrast. This is a limitation of the current deposit revision rather than of the experiment.

Session scheduling was nominally twice daily, morning and evening. In practice the delivered cadence was irregular: mission-day coverage is dense for M1 and M3 and sparse for M2 (Table 2, distinct-days column), and the first session of the series for one crew member is incomplete, containing the cough block and sustained /a/ but neither /i/ nor /u/.

Mission extension. The crew's on-orbit stay was extended beyond the planned duration. The additional days were used for further sessions under the identical protocol, so the delivered corpus is larger than the manifested one: the mid-mission return inventory enumerates 35 sessions, the final deposit resolves 44. This is a favourable case of a common exploration-mission condition — schedule slip converted into science return at zero marginal resource cost, precisely because the task battery consumes nothing and needs no ground coordination to repeat. A modality with these properties absorbs mission-plan volatility rather than being damaged by it, which is an argument for manifesting it that is independent of what it measures.

2.4 Acquisition hardware and processing chain

Audio was captured by purpose-built acquisition software running on an Apple iPad mini tablet aboard the station, using the tablet's built-in microphone, and transferred to ground.

The acquisition platform is itself a design finding worth recording. The investigation originally developed a dedicated hardware unit — a self-contained enclosure with microphone input, on-board signal processing, and screen, haptic and audio feedback — and delivered a prototype. It was not flown: flight-certification procedure required the experiment to run on already-certified crew tablets, and the software was accordingly re-implemented for iPad and adapted to the flight operator's requirements. The consequence for this paper is that the acquisition chain was fixed by vehicle-certification constraints rather than by measurement requirements, which is the normal condition for a payload of this class and the reason instrument characterization after the fact (Section 3) is necessary rather than optional.

Two parallel versions of every recording were retained and independently measured:

- **original** — the acquisition as recorded.
- **DNS64** — the same material after a deep noise-suppression stage.

Retaining both is the design decision that makes this paper possible. Corpora seldom publish both versions, which leaves the contribution of the processing stage unmeasured and unmeasurable. Here it is measured directly (Section 3.2).

2.5 Measurement pipeline

All measurements were computed with a single pipeline built on Praat 6.x through the Parselmouth interface, with all audio brought to 16 kHz before measurement.

Sustained vowels. Fundamental frequency (F_0, autocorrelation), local jitter, local shimmer, harmonics-to-noise ratio (HNR, cross-correlation harmonicity), and smoothed cepstral peak prominence (CPPS) [17–19]. All were computed in Praat through the Parselmouth interface [15, 16].

Coughs. Two conventions were computed:

- **Burst-level.** The amplitude envelope (20 ms RMS) is segmented into individual cough bursts at 25% of peak amplitude with a 150 ms minimum separation. Within each burst, rise time (10%→90% of that burst's peak), duration, decay and spectral centroid are measured on a finer 5 ms envelope. The recording value is the median across bursts.
- **Whole-recording.** A single peak is located in the entire recording and the same quantities are computed from it.

Both were computed because the literature is divided between them and, as Section 3.3 shows, the choice changes reported values by an order of magnitude.

2.6 Temporal provenance

Acquisition times were not taken from filenames. Three independent sources were compared: timestamps embedded in version-1 UUID filenames, filesystem modification time, and creation time.

Table 3
Agreement of time sources.

Source	Records	Agreement with UUID-embedded time
UUID version 1, embedded	157	157/157 exact to 0.000 s
Filesystem modification time	97	median offset + 0.02 h, SD 2.41 h
Creation time	—	not usable (reset during transfer)

For 157 of 254 records the acquisition instant is known to millisecond precision. The remaining 97 carry a timestamp-source flag and a stated ± 2.4 h uncertainty. Any analysis sensitive to time of day must respect that flag. We report this in detail because provenance discipline is not incidental to a flight instrument: a monitoring system whose sample times are recoverable only to ± 2.4 h cannot resolve circadian structure, and circadian structure is one of the effects such a system would most plausibly be asked to detect.

3. Results

3.1 Approach

An acoustic measurement made aboard a spacecraft is only as interpretable as its acquisition chain is understood. We therefore establish the behaviour of that chain first (Sections 3.2–3.4) and then report the phase contrast (Section 3.5), so that the physiological result rests on a characterized instrument

rather than an assumed one. Two properties of the chain matter for what follows: which measures it perturbs, and whether it perturbs them differently between phases. The first is quantified below; the second is addressed by computing the phase contrast within a single processing version, so that the chain term is identical on both sides and cancels.

3.2 Processing-stage sensitivity

Every recording exists in both versions, so the processing stage can be evaluated by exact pairing. We paired on (crew, phase, mission day, task), giving 63 matched cells, and tested with the Wilcoxon signed-rank test.

Table 4
Effect of the noise-suppression stage on sustained-vowel measures. Paired, n = 63 matched cells.

Measure	original median	DNS64 median	Median Δ	Wilcoxon p	d_z
F_0 (Hz)	100.35	100.34	+ 0.19	4.5×10^{-3}	+ 0.18
Jitter (%)	0.68	0.64	−0.06	1.2×10^{-4}	−0.52
Shimmer (%)	**9.42**	**5.70**	**−3.35**	$\mathbf{1.2 \times 10^{-11}}$	**−1.21**
HNR (dB)	**13.86**	**16.81**	**+ 3.56**	$\mathbf{1.4 \times 10^{-11}}$	**+ 1.32**
CPPS (dB)	11.94	11.35	+ 0.10	0.94	−0.08

The pattern is mechanistically coherent. Noise suppression acts strongly and specifically on the two measures *defined against a noise floor* — shimmer and HNR — because removing the noise floor is precisely what the stage does. F_0 and CPPS, which are defined by periodicity and cepstral structure rather than by a noise reference, are essentially unmoved: CPPS shows no detectable shift at all (p = 0.94), and the F_0 shift, while formally significant, is + 0.19 Hz on a ~ 100 Hz carrier with d_z = 0.18, i.e. negligible.

3.3 Segmentation-convention sensitivity

Table 5
Cough measures under both conventions. Paired within recording, n = 135.

Measure	Burst-level	Whole-recording	Ratio	Wilcoxon p
Rise (ms)	11.0	108.9	**9.90×**	7.2×10^{-24}
Duration (ms)	91.4	276.8	3.03×	1.5×10^{-12}
Decay (ms)	27.0	33.7	1.25×	1.1×10^{-9}
Centroid (Hz)	715.0	765.2	1.07×	3.3×10^{-4}

The mechanism is directly attributable to the fixed-duration design of Section 2.2. Median active fraction is 0.091 — **90.9% of each recording is silence**, and the median recording contains 2 detected bursts (IQR

2–3, range 1–7). The whole-recording convention locates a single global peak and therefore treats inter-burst silence as part of the rise, inflating rise time by an order of magnitude. This is not a property of the crew, of microgravity, or of the corpus; it is a property of applying a single-event convention to a multi-event recording, and it will appear in any protocol that requests more than one transient per fixed-length file.

The engineering implication is immediate: a flight implementation must fix its segmentation convention as part of the instrument specification, not leave it to the analysis stage, because the convention contributes an order of magnitude more to reported cough timing than any physiological effect plausibly does.

3.4 Chain-invariant observables

Two classes of observable survive the tests above.

Ratio quantities. Because /a/, /i/ and /u/ are acquired seconds apart through an identical path, ratios formed between them cancel gain, microphone-distance and enclosure terms algebraically, and cancel the processing-stage term to the extent that it acts multiplicatively and identically on both.

Table 6
Dimensionless ratio HNR(/a/) ÷ HNR(/i/), session-level.

Crew	Sessions	Median	Within-crew CV
M1	18	0.760	15.0%
M2	9	0.980	11.3%
M3	16	1.028	14.5%

Absolute HNR moves 3.56 dB (≈ 26%) between processing versions (Table 4); the within-crew dispersion of the ratio is 11.3–15.0%, and its central value is stable across versions. For reference, within-corpus dispersion of this ratio measured with the identical pipeline on a sound-treated clinical corpus [24] is of the same order (≈ 15%), which indicates that the ratio's stability aboard the station is not markedly worse than in a controlled acoustic environment.

Analysis note. Ratios in Table 6 are computed on the DNS64 version only. Session-level pairing requires a session key, and time-of-day is populated only for DNS64 records in the released deposit; on the original records, morning and evening sessions of the same mission day collapse into a single key and the resulting dispersion (CV 48.9–70.3%) is an artifact of that collapse rather than a property of the ratio. This is a deposit-completeness issue, not an analytical one, and it is being corrected.

Periodicity- and cepstrum-based measures. F_0 and CPPS are interchangeable between versions (Table 4) and are therefore the two absolute measures a flight implementation can report without specifying which processing stage produced them.

3.5 The phase contrast: a consistent in-flight decline in HNR

Sections 3.2–3.4 establish what the acquisition and processing chain contributes. We now use that knowledge rather than merely reporting it: the phase contrast below is computed **entirely within the DNS64 version**, so the processing term of Table 4 is identical on both sides of every comparison and cancels exactly. Establishing that term is what licenses the interpretation, not what limits it.

The analysis uses all three vowels. Eight of the nine crew × vowel cells support a phase comparison; the ninth (M3, /u/) carries one verified record in each phase after task-slot verification and is reported but not tested.

Table 7
HNR by crew and vowel, DNS64. Medians in dB.

Crew	Vowel	Pre-flight *n*	Pre-flight	In-flight *n*	In-flight	Δ (dB)	Tested
M1	/a/	2	14.86	16	12.96	**−1.89**	yes
M1	/i/	2	19.79	16	17.46	**−2.33**	yes
M1	/u/	2	22.89	16	22.88	**−0.01**	yes
M2	/a/	4	18.27	7	16.40	**−1.87**	yes
M2	/i/	4	18.19	7	14.97	**−3.22**	yes
M2	/u/	4	15.02	5	19.13	+ 4.11	yes
M3	/a/	2	16.38	14	13.02	**−3.36**	yes
M3	/i/	2	17.75	14	12.96	**−4.79**	yes
M3	/u/	1	22.62	1	22.77	+ 0.15	no — one record per phase

Seven of the eight testable cells decline, median − 2.11 dB (Fig. 1A). Restricted to /a/ and /i/ — the two vowels sampled completely in every crew member — the decline is 6 of 6 cells, median − 2.77 dB, sign test $p = 0.031$.

The one exception. M2 /u/ is the only cell that rises, and it is the least well conditioned in the corpus. Its four pre-flight values are 23.6, 13.6, 16.4 and − 0.2 dB. The last is a failed acquisition: an HNR at the noise floor indicates a recording with almost no recoverable periodic energy, and it displaces the pre-flight median downward, which is what produces the apparent in-flight rise. The cell is also the sparsest on the in-flight side, task-slot verification having removed two of M2's seven in-flight /u/ recordings. We retain the cell in Table 7 and in the sign test rather than removing it, since excluding the single datum

that opposes the hypothesis would not be a defensible analysis choice; we note only that the exception is attributable to a recording failure rather than to a physiological reversal.

Unequal pre-flight sampling. The number of pre-flight recordings differs by crew member and is visible in Fig. 1A: M1 contributed two complete pre-flight sessions, M3 three sessions that are individually incomplete (two carry /a/ and /i/, one carries /u/), and M2 three sessions of which one recorded the vowel battery twice, giving four pre-flight records per vowel. All are distinct acquisitions verified by content hash; none is a duplicated record. The consequence is that pre-flight *n* per cell ranges from one to four, which is the dominant constraint on the statistical power available here and the reason Section 5 treats baseline balance as the primary design recommendation.

Table 8 places the other four measures on the same footing, to test whether the change is specific to HNR or part of a general displacement.

Table 8
Direction consistency across measures, over the eight testable crew × vowel cells, DNS64.

Measure	Cells declining	Median Δ	Sign test *p*	Sensitive to added noise?
HNR (dB)	**7/8**	**−2.11**	**0.070**	yes
Jitter (%)	6/8	−0.14	0.29	weakly
CPPS (dB)	3/8	+ 0.90	0.73	yes
F_0 (Hz)	3/8	+ 1.02	0.73	no
Shimmer (%)	2/8	+ 0.57	0.29	yes

Table 9 repeats the comparison at the level of individual recordings, which uses every recording rather than one median per cell and is our primary test.

Table 9
HNR phase contrast at recording level, all three vowels pooled within crew, DNS64. Mann–Whitney *U*.

Crew	Pre-flight *n*	Median (dB)	In-flight *n*	Median (dB)	Δ (dB)	*p*	Rank-biserial
M1	6	19.79	48	18.05	−1.74	0.582	−0.15
M2	12	16.85	19	16.40	−0.45	0.598	−0.12
M3	5	18.55	29	13.02	−5.53	**0.011**	−0.72
Pooled, crew-centred	**23**	—	**96**	—	**−2.70**	**0.019**	—

In the pooled comparison each crew member's own pre-flight median is subtracted before pooling, so between-subject level differences cannot contribute; the result is − 2.70 dB at $p = 0.019$. Only M3 reaches significance individually, which is expected at pre-flight sample sizes of five to twelve recordings (Section

5). Restricting to /a/ and /i/ gives − 2.47 dB at $p = 0.0038$; the /u/ task adds noise without changing the direction, consistent with its incomplete verification.

Specificity. The pattern across measures is the strongest argument that this reflects the crew rather than the recording environment (Fig. 1B). **Shimmer and CPPS are, like HNR, sensitive to added acquisition noise** — Table 4 shows the processing stage moves shimmer by − 3.35% and HNR by + 3.56 dB through exactly that pathway. If the acoustic conditions of recording had drifted systematically between the pre-flight and in-flight phases, shimmer and CPPS would have been displaced alongside HNR. They were not: shimmer declines in 2 of 8 cells and CPPS in 3 of 8, with medians on the wrong side of zero. F_0, which depends on neither, is likewise unmoved at 3 of 8. Jitter shows a weaker tendency in the same direction as HNR (6 of 8, $p = 0.29$) that we do not interpret. The selectivity is not the pattern an acquisition artifact produces.

Direction and mechanism. HNR is the ratio of periodic to aperiodic energy in the phonated signal; a decline means proportionally more turbulent, noise-like energy at the glottis. Cephalad fluid redistribution loads the laryngeal mucosa and alters vocal-fold surface hydration and viscoelasticity [10, 11], and the dependence of phonatory efficiency on vocal-fold hydration is established terrestrially through phonation threshold pressure studies [12–14]. Reduced efficiency of glottal closure produces exactly the increase in turbulent energy that a falling HNR measures. The observed direction is the one the mechanism predicts.

Temporal form. The change presents as a sustained offset between phases rather than a progressive in-flight trend. Table 10 gives within-flight Spearman correlations against mission day.

Table 10
Within-flight trajectory, sustained /a/, DNS64, Spearman ρ against mission day (nominal *p*, uncorrected).

Measure	M1 (*n* = 16)	M2 (*n* = 7)	M3 (*n* = 14)
F_0	−0.349 (0.185)	−0.378 (0.403)	+ 0.046 (0.875)
Jitter	−0.408 (0.116)	+ 0.189 (0.685)	−0.002 (0.994)
Shimmer	−0.805 (< 0.001)	+ 0.850 (0.015)	+ 0.214 (0.462)
HNR	+ 0.911 (< 0.001)	−0.567 (0.184)	−0.069 (0.816)
CPPS	+ 0.763 (0.001)	−0.378 (0.403)	−0.245 (0.398)

No measure trends consistently across crew within the flight phase. For the fluid-shift interpretation this is the expected form rather than a contradiction: cephalad redistribution is largely complete within hours of insertion and reaches a new steady state, so sampling that begins on mission day 3 should observe a sustained offset from baseline and no continuing drift — which is what Tables 7–10 jointly show. The divergent individual slopes within the flight phase are not interpreted here; with three participants and no

logged acquisition condition (Section 6) they cannot be separated from per-crew differences in recording circumstance.

3.6 The cough task shows no consistent phase effect

The battery also delivered 135 verified cough recordings, distributed across both phases in every crew member (M1: 6 pre-flight, 48 in-flight; M2: 12 and 21; M3: 6 and 42). We tested the same phase contrast on the burst-level cough measures.

Table 11
Cough phase contrast, burst-level, DNS64, median change by crew (Mann–Whitney *p*).

Measure	M1	M2	M3	Cells down
Burst centroid (Hz)	+ 170.9 (0.002)	+ 1236.2 (0.002)	−147.1 (0.820)	1/3
Burst rise (ms)	−1.6 (0.356)	−5.8 (0.421)	+ 0.6 (0.732)	2/3
Burst duration (ms)	+ 116.7 (0.010)	+ 41.5 (0.231)	−11.5 (0.891)	1/3
Burst decay (ms)	+ 14.3 (0.012)	+ 9.1 (0.321)	+ 5.0 (0.242)	0/3
Bursts per recording	−1.5 (0.180)	−0.5 (0.725)	−0.5 (0.431)	3/3

No cough measure moves consistently across crew. Burst spectral centroid changes substantially and significantly in two crew members but in opposite directions from the third (+ 170.9 and + 1236.2 Hz against − 147.1 Hz), and the crew-centred pooled comparison is null (+ 126.9 Hz, $p = 0.235$). Only burst count declines in all three, by a median of half a burst per recording, with no crew reaching significance.

This is a useful negative result for instrument selection rather than a failure of the task. A voluntary cough is a maximal-effort ballistic manoeuvre whose acoustics depend strongly on the effort the participant applies and on the geometry between mouth and microphone at the instant of the burst — neither of which the protocol constrains, and neither of which is stable in free-fall. Sustained phonation at comfortable pitch and loudness is, by comparison, a self-normalising task. The contrast between Section 3.5 and this section is therefore evidence about which acoustic task a crew-health instrument should carry: **the same crew, the same sessions, the same acquisition chain and the same processing produced a consistent phase effect in sustained phonation and none in cough.** We fold this into requirement R5.

The /u/ task and version availability. For completeness: /u/ is complete in the raw acquisition version for M2 and M3 but not M1, and complete in the processed version for M1 and M2 but not M3, so no single version supports a three-crew /u/ comparison. We report the processed version throughout for consistency with Tables 4–10 and do not combine versions, since that would reintroduce exactly the chain term that Section 3.2 quantifies.

What is and is not claimed. No blood, urine, or imaging reference was obtained from any crew member (Section 6). The attribution of the observed HNR decline to fluid shift rests on the direction of the effect, its specificity to the measure the mechanism predicts, its consistency across all three crew members, and its temporal form — not on a paired physiological measurement. It is an interpretation consistent with the evidence, and the study was not designed to exclude alternative systemic contributors such as altered cabin humidity or CO_2, sleep disruption, or workload, each of which could plausibly act in the same direction.

4. Operational Resource Model

The case for voice as an exploration-mission sensor rests on resource cost, so we state it quantitatively.

Table 12
Per-session resource cost, as designed and as delivered.

Resource	As designed	As delivered	Note
Recordings per session	9	6	forced-expiration block returned nothing
Crew recording time	27.3 s	**18.2 s**	9 vs 6 × 3.029 s
Consumables	0	0	no expendables of any kind
Incremental mass	0	0	ran on already-certified crew tablet
Raw audio volume	852 KiB	**568 KiB**	16 kHz, 16-bit, mono
Derived measurement record	936 B	**624 B**	13 fields per recording
Downlink reduction, on-board processing	932×	**932×**	ratio invariant to task count

Crew *recording* time is a hard lower bound fixed by the application; total crew time per session additionally includes tablet retrieval, application interaction and stowage, which the present records do not resolve. A flight implementation should instrument the application to log session wall-clock duration, since crew time is the resource this modality is principally competing for and 18.2 s of phonation is not the number a timeline planner needs.

Over the flown mission the whole-corpus figures are 24.4 MiB of raw audio against 26.8 KiB of derived measurements. The trade is therefore stark and favours on-board measurement, but Sections 3.2–3.3 impose a condition on it: **if measurement is performed on board, the processing stage and the segmentation convention become part of the flight software configuration and must be version-controlled and reported with the data.** A pipeline change between mission increments would introduce a step discontinuity into the physiological record of the same order as the effects being monitored (Table 4 vs Table 7). This is a configuration-management requirement, not an analysis guideline.

The privacy and bandwidth arguments align here: downlinking derived measurements rather than audio reduces volume by three orders of magnitude *and* removes the biometric-identifiability problem that voice recordings otherwise carry for a small, identified professional cohort.

5. Design Requirements for Follow-On Missions

The dominant limitation of the present experiment is not participant count but the asymmetry of the sampling windows (Table 2). We express the requirement in the form a mission planner can use directly.

Table 13
Smallest within-subject phase effect detectable at 80% power, α = 0.05, two-sided two-sample *t*.

Crew	Quarantine sessions	In-flight sessions	Detectable *d*
M1	2	16	≥ 2.24
M2	3	6	≥ 2.31
M3	3	14	≥ 1.91

Table 14
Balanced sessions per phase per crew member required at 80% power, α = 0.05.

Target effect *d*	Sessions per phase
2.0	6
1.5	9
1.3	11
1.0	17
0.8	26

The design conclusion is unambiguous: **balancing the pre-flight window against the flight window is the single most effective change available**, and it costs nothing in flight resources because pre-flight sessions are acquired on the ground. Adding pre-flight sessions from 2–3 to 9 moves the detectable effect from $d \approx 2.2$ to $d = 1.5$ with no impact on the crew timeline in orbit.

We therefore recommend, for a follow-on implementation:

- **R1.** ≥ 9 pre-flight sessions per crew member through the identical acquisition path, ideally within a defined pre-flight window rather than distributed over weeks (cf. M2 in Table 2).
- **R2.** Processing stage and segmentation convention frozen and version-stamped in flight software; both raw and processed forms retained where downlink permits, or the processed form only with its configuration hash.

- **R3.** Acquisition-time provenance to millisecond precision for all records (achieved for 157/254 here; the ± 2.4 h residual forecloses circadian analysis).
- **R4.** Session clock times matched between phases. Mean session clock time differs between phases by 2.8–8.1 h depending on crew member in the present data, so phase contrast and time-of-day effect are not orthogonal.
- **R5.** Carry sustained phonation as the primary task. In the present experiment the same crew, sessions, acquisition chain and processing produced a consistent phase effect in sustained vowels (Section 3.5) and none in voluntary cough (Section 3.6); cough acoustics depend on applied effort and instantaneous mouth-to-microphone geometry, neither of which is constrained in free-fall. Where absolute values must be compared across chains rather than within one, restrict reporting to the chain-invariant set — F_0, CPPS, and within-session dimensionless ratios.
- **R6.** A concurrent reference measurement stream. The present experiment has no physiological ground truth; without one, no acoustic observable can be calibrated to a physiological quantity, only characterized as an instrument.
- **R7.** No task whose observable is defined against an *absolute ambient reference* should be manifested without prior validation in a representative acoustic environment. The forced-expiration block (Section 2.2) was validated against hospital spirometry in a terrestrial clinic and returned zero usable sessions in orbit, because its decibel-difference observable presumes a quiet reference and a fixed mouth-to-microphone geometry that a ventilated module in free-fall does not provide. Tasks whose observables are defined *within* the recording — periodicity, cepstral structure, within-session ratios — survived the same transition intact. Cabin-representative acoustic validation is cheap on the ground and, on this evidence, decisive.
- **R8.** Instrument the acquisition application to log session wall-clock duration and an ambient-noise estimate per session, **as payload-owned telemetry**. Neither is available in the present records; ambient-noise data was requested from the flight operator and not provided. A measurement the payload needs in order to interpret its own output should not depend on a post-hoc data request to another organization. Both quantities cost essentially nothing to record at acquisition time, and their absence is what prevents Section 3.5 from separating acquisition condition from physiology.

6. Limitations

No physiological reference. No blood, urine or imaging measurement was obtained from any crew member. Nothing in this paper calibrates an acoustic observable against a physiological quantity, and no diagnostic or individual clinical claim is made or supported. The fluid-shift attribution in Section 3.5 rests on the direction, specificity and temporal form of the observed effect; alternative systemic contributors acting in the same direction — altered cabin humidity or CO_2, sleep disruption, workload — were not measured and cannot be excluded. A paired physiological reference is requirement R6 for any follow-on implementation.

Three participants. The phase contrast of Section 3.5 is a within-subject comparison in which each crew member serves as their own reference, and its consistency across all three is the basis of the result; but with *n* = 3 the effect size is bounded rather than estimated, and no population-level inference is drawn. Between-crew comparison is separately uninformative at this sample size: per-crew cough spectral centroid spans 570.5, 1422.5 and 768.3 Hz, a range comparable to the spread of large terrestrial reference corpora.

Uncontrolled and unmeasured acoustic environment. Hold distance, cabin location and background activity were neither controlled nor logged. Ambient noise level in dB was specified as a pre-test step in the protocol — it is the calibration reference the forced-expiration block depended on — but no per-session ambient value is present in the returned records. **The investigators requested ambient-noise data from the flight operator and it was not provided.** The acquisition condition is therefore unmeasured for every session in this corpus, and this is a data-access outcome rather than an instrumentation oversight.

The consequence is specific and limiting: the divergent within-flight trajectories of Table 9 cannot be separated from variation in acquisition condition, because the quantity that would separate them does not exist in the returned data. We regard this as the strongest single argument for R8 — an ambient estimate logged *by the acquisition application itself*, and therefore owned by the investigation rather than requested from an operator after the fact, would have made Table 9 interpretable at negligible cost.

Operational and deposit losses. The forced-expiration block returned nothing from any session and is not recoverable. Separately, the raw-version pre-flight measurements of one crew member are absent from the released deposit, so all phase comparisons are computed within the processed version; this affects which version can be used, not which crew members are represented, and all three contribute to Section 3.5.

Asymmetric and, for M2, temporally dispersed baselines. See Table 2 and Table 13.

Partial temporal precision. 97 of 254 records carry ± 2.4 h acquisition-time uncertainty (Table 3), which forecloses circadian analysis on that subset.

Deposit completeness. Time-of-day and analysed-duration fields are populated only for DNS64 records in the released deposit, and burst-level cough measures are likewise available only for the DNS64 half. Analyses in Sections 3.3 and 3.4 are therefore restricted to that version. This is a packaging limitation of the current deposit revision and does not affect the paired comparisons of Table 4, which were computed on matched cells present in both versions.

7. Conclusions

1. **Harmonics-to-noise ratio declined between the pre-flight and in-flight phases in seven of eight testable crew × vowel cells**, median −2.11 dB; at recording level with all three vowels pooled and crew-

centred, −2.70 dB (p = 0.019), and −2.47 dB (p = 0.0038) restricted to the two completely sampled vowels. To our knowledge this is the first detection of a consistent in-flight change in a vocal acoustic measure from a controlled, protocol-fixed in-orbit corpus with a pre-flight baseline acquired through an identical path.

2. **The effect is specific to HNR, and absent from cough.** Shimmer and CPPS — which share HNR's sensitivity to added acquisition noise — decline in only 2 and 3 of 8 cells; a between-phase drift of the recording environment would have moved all three together. It did not, which is the principal argument against an acquisition artifact. Applied to 135 cough recordings the same contrast is null in every measure, indicating that sustained phonation rather than cough is the task a crew-health instrument should carry.

3. **The direction is the one the mechanism predicts.** A falling HNR is proportionally more turbulent energy at the glottis, the acoustic signature of reduced glottal-closure efficiency, which is what laryngeal mucosal fluid loading produces. The change presents as a sustained offset between phases with no progressive in-flight trend — the form expected of a redistribution that completes within hours and reaches a new steady state before sampling begins on mission day 3.

4. **The result is attributable to the crew rather than to the pipeline because the pipeline was characterized first.** Retaining both raw and processed forms of every recording — uncommon in acoustic corpora — allowed the processing stage to be measured directly (shimmer −3.35 %, HNR +3.56 dB, CPPS unchanged at p = 0.94). The phase contrast is computed within a single version, so that term cancels exactly. Establishing the chain effect is what licenses the physiological reading.

5. **Two chain properties are reported for reuse.** Cough timing depends 9.90-fold on segmentation convention, because 90.9 % of each fixed-duration recording is silence; and within-session dimensionless ratios are stable across processing versions (within-crew CV 11.3–15.0 %) while the absolute measure moves by 26 %.

6. **The operational case is quantified:** 18.2 s of crew recording time, zero consumables, zero incremental mass, and 624 B of downlink per session against 568 KiB of raw audio — a 932-fold reduction with on-board measurement, conditional on flight-software configuration management of the pipeline. The protocol sustained a twice-daily cadence across the mission and absorbed an extension of the on-orbit stay at no marginal resource cost.

7. **The binding constraint on the next experiment is baseline balance, not participant count.** Nine sessions per phase per crew member reach d = 1.5 and seventeen reach d = 1.0, against $d \geq 1.91$ permitted here. Pre-flight sessions are acquired on the ground and cost nothing in flight resources — the highest-leverage design change available, and the one that would convert the present result from a consistent direction into a calibrated effect size.

8. **A manifested task failed for a generalizable reason.** The forced-expiration block, validated on the ground against hospital spirometry, returned zero usable sessions, because its observable is a decibel difference against an absolute ambient reference that a 63 dBA ventilated module does not provide. Observables defined *within* a recording survived the transition to a spacecraft cabin; observables defined against an *absolute external reference* did not. We propose this as a design rule for acoustic crew-health instrumentation.

Voice is a viable exploration-mission crew-health sensing modality, it costs almost nothing to fly, and in a controlled in-orbit experiment it registered a consistent, mechanism-consistent change across an entire crew. Confirming that the change reflects fluid shift specifically requires a paired physiological reference and a balanced pre-flight window — both are within reach of a follow-on implementation, and R1–R8 specify what they require.

Declarations

Data Availability

Derived acoustic measurements and metadata are deposited at Zenodo, DOI [10.5281/zenodo.21939101](https://doi.org/10.5281/zenodo.21939101), CC BY 4.0 [25]. Raw and enhanced audio are not distributed: voice and cough recordings are biometric data from which speaker identity can be recovered, and the participants are three members of an identified professional cohort, so public release could not be reconciled with the confidentiality commitments in the study consent. This release model follows that of the Oxford/UCI Parkinson's voice corpora [22, 23], which distribute computed features with no audio and are among the most reused voice datasets in the field. Requests for audio access for verification purposes will be considered by the corresponding author subject to the approvals governing the study.

Code Availability

All measurement code is deposited with the data at the same DOI: self-contained Python 3 command-line programs using Praat 6.x through Parselmouth, with dependency specification and usage documentation.

Author Contributions

G.A. conceived the study, designed the investigation, obtained regulatory and ethical approvals, and led data collection. O.A. contributed to protocol design and data collection, performed the acoustic measurement and instrument characterization, produced the released data tables and measurement code, and drafted the manuscript. Both authors reviewed and approved the final manuscript.

Competing Interests

Under the contract between the parties, TÜBİTAK UZAY and the Turkish Space Agency could derive financial benefit from sales of the software used in this investigation, as disclosed in the participant consent form. Conflict of interest disclosures were filed with the NASA Institutional Review Board for all investigators. The authors declare no other competing interests.

Funding

Supported by the Scientific and Technological Research Council of Türkiye (TÜBİTAK) and the Turkish Space Agency (TUA) through the Turkish Astronaut and Scientific Mission (TABM) Project (1007-KAMAG-121L002; Agreement No. TABM-HZT-23-28). The funders had no role in study design, data collection, analysis, or the decision to publish. None of the opinions and findings in this publication represent the official views of TÜBİTAK or TUA.

Ethics

Approved by the NASA Institutional Review Board (FWA No. 00019876; eIRB No. STUDY00000621, expedited review category 4) and conducted in accordance with the Declaration of Helsinki, under protocols with the Turkish Space Agency, NASA and Axiom Space. Written informed consent was obtained from all participants. The consent provides for publication without information that could identify participants and specifies that investigators do not return individual research results; accordingly this paper reports no individual-level clinical interpretation.

Acknowledgements

The authors thank the crew for their participation, and the Turkish Space Agency, NASA and Axiom Space for mission and protocol support.

References

1. Antonsen, E.L., Myers, J.G., Boley, L., Arellano, J., Kerstman, E., Kadwa, B., Buckland, D.M., Van Baalen, M.: Estimating medical risk in human spaceflight. *npj Microgravity* **8**, 8 (2022). https://doi.org/10.1038/s41526-022-00193-9
2. Russell, B., Burian, B.K., Hilmers, D., Beard, B., Martin, K., Pletcher, D., Easter, B., Lehnhardt, K., Levin, D.: The value of a spaceflight clinical decision support system for earth-independent medical operations. *npj Microgravity* **9**, 46 (2023). https://doi.org/10.1038/s41526-023-00284-1
3. Ciofani, G., Bandiera, T., Corsini, A., Crescenzi, M., De Vittorio, M., Mari, S., Martinelli, E., Monici, M., Piccirillo, S., Narici, M., Ferranti, F.: Pharmaceutical and biomedical challenges for crew autonomy in health preservation during future exploration missions. *Communications Medicine* **5**, 418 (2025). https://doi.org/10.1038/s43856-025-01128-7
4. Urbina, M., Rubin, D., Hailey, M., Reyes, D., Antonsen, E.: Medical System Concept of Operations for Mars Exploration Missions. Human Research Program Investigators' Workshop, Galveston, Texas, 23–26 January 2017. NASA JSC-CN-38405, NTRS 20170000500.

5. Hamilton, D.: Autonomous Medical Care for Exploration Class Space Missions. NASA Johnson Space Center, NTRS 20070032039.
6. NASA Human Research Program: Evidence Report — Risk of Adverse Health Outcomes and Decrements in Performance Due to In-Flight Medical Conditions. NASA Human Research Roadmap, Exploration Medical Capability Element.
7. NASA Exploration Medical Capability: Development of an Accepted Medical Condition List for Exploration Medical Capability Scoping. NTRS 20220003136 (2022).
8. Allen, C.S.: International Space Station Acoustics — A Status Report. ICES-2024-354, 53rd International Conference on Environmental Systems, Louisville, Kentucky, 21–25 July 2024.
9. Limardo, J.G., Allen, C.S., Danielson, R.W.: Assessment of Crewmember Noise Exposures on the International Space Station. 43rd International Conference on Environmental Systems, AIAA 2013-3516 (2013). https://doi.org/10.2514/6.2013-3516
10. Norsk, P.: Adaptation of the cardiovascular system to weightlessness: surprises, paradoxes and implications for deep space missions. *Acta Physiologica* **228**, e13434 (2020). https://doi.org/10.1111/apha.13434
11. Mader, T.H., Gibson, C.R.: Neuro-ocular findings and mechanisms in Spaceflight Associated Neuro-Ocular Syndrome. In: *Spaceflight Associated Neuro-Ocular Syndrome*, pp. 9–21. Elsevier (2022). https://doi.org/10.1016/B978-0-323-91524-3.00013-2
12. Verdolini-Marston, K., Titze, I.R., Druker, D.G.: Changes in phonation threshold pressure with induced conditions of hydration. *Journal of Voice* **4**, 142–151 (1990). https://doi.org/10.1016/S0892-1997(05)80139-0
13. Solomon, N.P., DiMattia, M.S.: Effects of a vocally fatiguing task and systemic hydration on phonation threshold pressure. *Journal of Voice* **14**, 341–362 (2000). https://doi.org/10.1016/S0892-1997(00)80080-6
14. Leydon, C., Wroblewski, M., Eichorn, N., Sivasankar, M.: A meta-analysis of outcomes of hydration intervention on phonation threshold pressure. *Journal of Voice* **24**, 637–643 (2010). https://doi.org/10.1016/j.jvoice.2009.06.001
15. Boersma, P., Weenink, D.: Praat: doing phonetics by computer, version 6.x. http://www.praat.org
16. Jadoul, Y., Thompson, B., de Boer, B.: Introducing Parselmouth: a Python interface to Praat. *Journal of Phonetics* **71**, 1–15 (2018). https://doi.org/10.1016/j.wocn.2018.07.001
17. Hillenbrand, J., Cleveland, R.A., Erickson, R.L.: Acoustic correlates of breathy vocal quality. *Journal of Speech, Language, and Hearing Research* **37**, 769–778 (1994). https://doi.org/10.1044/jshr.3704.769
18. Hillenbrand, J., Houde, R.A.: Acoustic correlates of breathy vocal quality: dysphonic voices and continuous speech. *Journal of Speech, Language, and Hearing Research* **39**, 311–321 (1996). https://doi.org/10.1044/jshr.3902.311
19. Maryn, Y., Weenink, D.: Objective dysphonia measures in the program Praat: smoothed cepstral peak prominence and acoustic voice quality index. *Journal of Voice* **29**, 35–43 (2015). https://doi.org/10.1016/j.jvoice.2014.06.015

20. Bhattacharya, D., Sharma, N.K., Dutta, D., Chetupalli, S.R., Mote, P., Ganapathy, S., Chandrakiran, C., Nori, S.: Coswara: a respiratory sounds and symptoms dataset for remote screening of SARS-CoV-2 infection. *Scientific Data* **10**, 397 (2023). https://doi.org/10.1038/s41597-023-02266-0
21. Orlandic, L., Teijeiro, T., Atienza, D.: The COUGHVID crowdsourcing dataset, a corpus for the study of large-scale cough analysis algorithms. *Scientific Data* **8**, 156 (2021). https://doi.org/10.1038/s41597-021-00937-4
22. Little, M.A., McSharry, P.E., Hunter, E.J., Spielman, J., Ramig, L.O.: Suitability of dysphonia measurements for telemonitoring of Parkinson's disease. *IEEE Transactions on Biomedical Engineering* **56**, 1015–1022 (2009). https://doi.org/10.1109/TBME.2008.2005954
23. Tsanas, A., Little, M.A., McSharry, P.E., Ramig, L.O.: Accurate telemonitoring of Parkinson's disease progression by noninvasive speech tests. *IEEE Transactions on Biomedical Engineering* **57**, 884–893 (2010). https://doi.org/10.1109/TBME.2009.2036000
24. Saarbrücken Voice Database. Institute of Phonetics, Saarland University. https://zenodo.org/records/16874898
25. Aydemir, O., Aydemir, G.: A longitudinal corpus of sustained phonation and voluntary cough acoustics from three ISS crew members — derived acoustic measurements. Zenodo (2026). https://doi.org/10.5281/zenodo.21939101

Figures

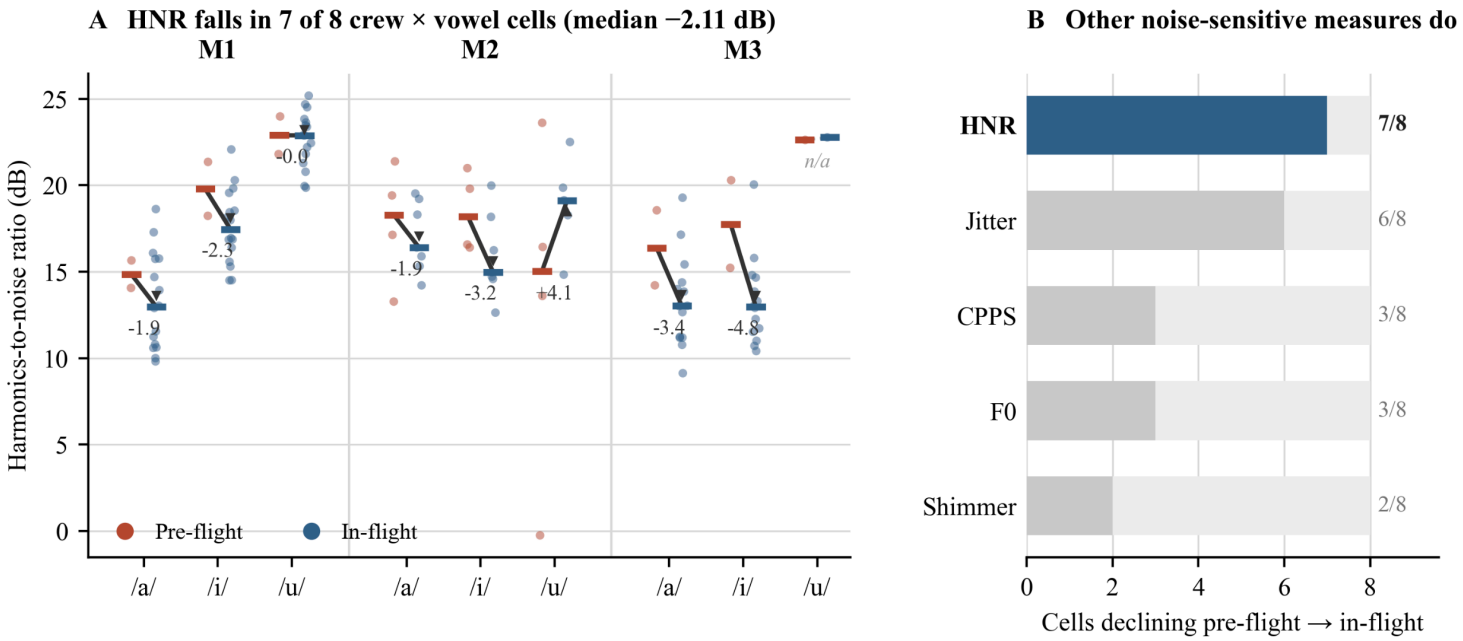


Figure 1

Harmonics-to-noise ratio by phase. **(A)** Every verified sustained-vowel recording, by crew member and vowel; red = pre-flight, blue = in-flight, horizontal bars are medians, arrows give the median change in dB. Seven of the eight testable cells decline. M3 /u/ carries one record per phase and is not tested. **(B)** Direction consistency across the five vowel measures: cells declining from pre-flight to in-flight, out of eight. Shimmer and CPPS, which share HNR's sensitivity to added acquisition noise, do not move — the pattern expected of a physiological rather than an instrumental change.