MOTION CLOSES THE RANK GAP: BAYES-OPTIMAL MOBILE SOUND-FIELD ESTIMATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Recovering a room's sound field from microphone recordings is basic to room acoustics and spatial audio. Most methods improve the estimator, with better priors, kernels and learned physics models. Unfortunately, a static array of m microphones resolves at most m independent directions of a K-mode field, so for $K > m$, absent sparsity, no estimator, noise level or recording length recovers the rest. In contrast, we treat the loss as one of acquisition. Moving the same array makes all K directions identifiable at equal observation time. Under a Bayes-optimal model its risk then falls inversely with observation time, with a constant we derive in closed form as a trace. That constant is governed by a contraction factor, which approaches one in some rooms as K grows. Therefore we prove which rooms keep it bounded away from one for every K. Facets in all four axis directions, each at least one grid spacing longer than the aperture, suffice, and the length threshold is sharp. Simulations with no fitted parameter match the predicted constant, and at the longer observation times the moving array beats the best static position across the room under independent estimators as well as our own. On a real 16-microphone array in a furnished room, under a pre-registered protocol, 16 surveyed positions resolve more independent directions than any static position can, and the held-out error falls from one position to 16 and ends below the best-predicted static position.

1 INTRODUCTION

An array cannot recover spatial information it never measures, and absent sparsity no estimator repairs that. A compact static array of m microphones, held at one position, observes a K-mode field through a Gram matrix of rank at most m. For $K > m$ and coefficients not assumed sparse, a $(K - m)$-dimensional subspace stays unidentifiable at any noise level and any observation time. Recording for longer only adds samples of the directions the array already measures. Moving the same array through the room adds the missing ones, at equal microphone-seconds (Figure 1). Under the declared Gaussian prior, the risk of the Bayes-optimal estimator then has an exact limit whose constant is a geometric quantity of the room and the array.

Measurement found this pathology before theory formalized it. Han & Kim (2026) report that a real 16-microphone array in a small room resolves only a few effective modal directions, and defer the trajectory formulation and its analysis to follow-up work. That work optimizes estimation inside the identifiable subspace. This paper asks when that subspace is the whole problem. Proposition 1 states the ceiling no static array escapes. Theorem 1 fixes the rate and the finite horizon that motion restores. Theorem 2 identifies the geometry that bounds the contraction factor uniformly in K.

The insight behind these results is that the gap is one of acquisition, not of estimation. Motion changes which subspace the array identifies, while a sharper estimator changes only how well it fills that subspace in, and this paper contributes the following.

- a rank obstruction for any static array with $m < K$, exact at zero noise and unlimited time (Proposition 1).
- the moving-array rate in trace form, $T\,\mathbb{E}\|\hat{c} - c\|^2 \to \sigma^2 \operatorname{tr}((I_K - B)^{-1})\,|D|/m$, with matched upper and lower constants and a concentration horizon $\widetilde{O}(K/(\rho_\star \Delta))$ (Theorem 1).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.



Figure 1: One array, two ways to record the same room. Held at one position, its microphones resolve only a few independent directions of the sound field, because a compact array hears nearly the same pressure at every microphone. Recorded at 16 surveyed positions for the same observation time, four of them shown, the same microphones resolve 17 directions, past the dotted ceiling of 16 that no static position of 16 microphones can pass, against three at one position (squares). This is the rank obstruction of Proposition 1 and its remedy, measured on real hardware. No field is drawn, because this room has no ground truth.

- the coverage identity $\sup_K \gamma_B = 1 - \kappa_{\min}/m$ and a sharp geometric certificate, in which four axis facets of length at least ns bound γ_B uniformly in K (Theorem 2).
- validated numerics for γ_B and $\mathrm{tr}((I_K - B)^{-1})$, in-model checks of the rate and the floor, the static rank pathology on public data, and a comparison in which, at the longer observation times, the moving array beats the best static position across the room under independent estimators too (Appendices C–F).
- a pre-registered measurement in a real room, in which 16 surveyed positions of a 16-microphone array resolve more independent directions than any static position of that array can, and its held-out reconstruction error falls from one position to 16 (Section 7.3).

2 Related work

Reconstruction from a static array. Room sound field reconstruction normally means estimating a field from microphones that stay where they were put. Work in this line varies the basis and the prior, from harmonic and spatio-temporal representations to Bayesian inference in rooms of general shape (Betlehem & Abhayapala, 2005; Fernandez-Grande, 2016; Tervo et al., 2013; Antonello et al., 2017; Mignot et al., 2013; Vu & Lissek, 2020; Schmid et al., 2021). The static geometry does not vary, and placing the array fixes what the measurement can resolve. Rafaely (2008) states that ceiling in array terms, tying spatial resolution to the harmonic order of the array. The recent answer to it is more hardware, and Xu & Jin (2025) and Xu et al. (2026) combine a spherical array with linear ones, first by residual refinement and then by one modal basis for the pair. Both lift the ceiling by adding microphones. The obstruction proved here turns on the microphone count itself, so motion lifts it at fixed hardware, and rooms that the array's translates cover keep the resulting constant bounded uniformly in the mode count.

Estimating better inside that limit. The natural response is a better estimator. Classical regularization supplies parameter-choice rules (Hansen, 1992; Golub et al., 1979; Engl et al., 1996), and the Bayesian formulation replaces the choice by a prior with rates attached (Kaipio & Somersalo, 2005; Stuart, 2010; Knapik et al., 2011). Learning occupies the same slot, through learned variational parameters and Tikhonov regularizers (Kunisch & Pock, 2013; Alberti et al., 2021; Burger & Kabri, 2023) and through iterations unrolled into networks (Gregor & LeCun, 2010; Adler & Öktem, 2018; Aggarwal et al., 2018). Room-acoustic priors go beyond these generic ones, from the physics-informed network of Karakonstantis et al. (2024) to diffusion models, wave-equation-constrained kernels and models conditioned on a reconstructed room geometry (Miotello et al., 2024; Ribeiro et al., 2026; Ahn et al., 2023). Arridge et al. (2019) survey the line, and others report what it cannot buy (Antun et al., 2020; Gottschling et al., 2025). Han & Kim (2026) takes the modal case to its

endpoint, where a trained regularizer converges to the closed-form Bayes filter. Every result here sharpens the estimate on the subspace the array can see, and none changes which subspace that is.

Motion, first as measurement practice. Moving-microphone measurement predates the theory for it. Ajdler et al. (2007) measure impulse responses with a continuously moving microphone, and Katzberg et al. (2017; 2021) develop sampling and reconstruction for moving receivers over a decade. Damiano et al. (2026) release a moving 16-microphone array on a rail with dense static references, and Brunnström et al. (2026) estimate the field from moving microphones under a kernel prior. Deppisch et al. (2024) enlarge the aperture by rotating a static array rather than translating it, the synthetic-aperture idea of array processing. The standardised moving-microphone method (Arachchige et al., 2018) is energy-only, with no phase and no modal identification. Motion is also not free, and ego-noise and position error can leave a static method ahead of a moving one on real hardware (Deppisch et al., 2025). This literature establishes that moving measurement works, without stating how the error falls with observation time or which room property sets the constant.

Motion, as theory. Unnikrishnan & Vetterli (2013) prove sampling advantages for moving sensors observing bandlimited fields, and Salgia & Kumar (2018) recover such a field from sensors that know neither where nor when they sampled. Aldroubi et al. (2026) survey dynamical sampling, where the field evolves under a known operator and the sampler is fixed, and Huang et al. (2024) extend it to convolutional operators. That setting is the inverse of this one, in which the field is time-invariant and the sampler moves. A deterministic path is formally an instance of it, with composition along the path as the known operator. The correspondence stops there, because a random trajectory has no such operator and composition does not preserve the span of the first K modes. All four results fix the signal class by its bandwidth, whereas the object here is parametric, the identification of K modal coefficients with rate constants set by the room and the array.

Choosing where to move. Given that motion restores identifiability, design asks where to move. Optimal experimental design answers by optimization, placing sensors under a Gaussian process model, treating infinite-dimensional and structural inverse problems, and making A-optimal criteria robust to misspecification (Krause et al., 2008; Alexanderian et al., 2014; Liu et al., 2021; Alexanderian, 2021; Attia et al., 2023). Attia (2026) promotes the design variable from placements to trajectories, and Deng et al. (2026) learn an active-sensing policy that beats static and offline-optimized layouts. These methods evaluate an objective numerically, and that objective is the trace of the posterior covariance. This paper works on the objective itself. Theorem 1 evaluates it in closed form, and a covering condition on the array's translates bounds γ_B uniformly in K (Theorem 2).

Numerics with error bars. Because those constants are geometric, they are computed in the spirit of verified numerics for the Laplacian on polygonal domains (Liu & Oishi, 2013; Carstensen & Gedicke, 2014) but weaker in kind. Here γ_B carries an error bar from a three-level mesh ladder and a declared branch rule (Appendix C) rather than a proved bound. Polyhedral clipping decides coverage exactly, without quadrature.

3 Setup

The room is a convex polyhedral domain $\Omega \subset \mathbb{R}^3$ with Neumann Laplace eigenpairs (λ_k, φ_k) and $L^2(\Omega)$-orthonormal eigenfunctions. The field is the K-mode truncation $u(x) = \sum_{k=1}^{K} c_k \, \varphi_k(x)$ with time-invariant coefficients $c \in \mathbb{R}^K$ drawn from the Gaussian prior $\mathcal{N}(0, \tau^2 I_K)$. The array is a rigid planar $n \times n$ microphone grid with spacing s, aperture $(n-1)s$, $m = n^2$ microphones and microphone offsets $p_1, \dots, p_m$, which form the set L. The placement domain $D \subset \Omega$ is the set of array positions whose footprint lies inside Ω, and $|D|$ is its volume. A trajectory is a path $(x_t)_{t=1}^{T}$ in D. At each snapshot, within which the array is treated as at rest, the array reports $y_t = \Phi(x_t)\, c + \varepsilon_t$, where $\Phi(x) \in \mathbb{R}^{m \times K}$ is the mode-sample matrix at the m microphone positions and $\varepsilon_t \sim \mathcal{N}(0, \sigma^2 I_m)$ is independent across snapshots. The trajectory enters the risk only through $\sum_t \Phi(x_t)^\top \Phi(x_t)$, which does not depend on the order of the snapshots, so what lifts the obstruction is the set of positions the array occupies and not the continuity of the path between them. The reference trajectory class is reflected Brownian motion on D. Its stationary law is uniform on D, its generator has spectral gap $\rho_\star$, and it is sampled at spacing Δ. The visibility matrix is $B =$

$I_K - \frac{1}{m}\int_D \Phi(x)^\top \Phi(x)\,dx$, the information the array misses over its placement domain, and $\gamma_B = \lambda_{\max}(B)$. Section 5 turns γ_B into a geometric quantity. The static comparison holds hardware and cost fixed, with the array observing the same model from one $x \in D$ for all T snapshots at equal microphone-seconds. Proposition 1 states this obstruction as an exact null-space fact. That fact is independent of the approximate aperture-to-wavelength bound of Han & Kim (2026), which explains the further collapse of effective rank below m for compact arrays. They identify two static remedies, enlarging the aperture and sampling spatially distinct positions over time. Rank m caps the first, so for $K > m$ motion is the only remedy at fixed hardware, and while their aperture is the physical extent relative to wavelength, aperture here means the $(n-1)s$ translation window of the anchor square.

4 How motion restores the rate

Proposition 1 (Rank obstruction)**.** *In the time-invariant-coefficient model of Section 3, with c free in $\mathbb{R}^K$, any static array of $m < K$ microphones yields a spatial Gram matrix of rank at most m. A $(K-m)$-dimensional subspace of coefficients is therefore unidentifiable from static measurements, at any noise level and any observation time.*

Proof. The Gram matrix is a sum of m rank-one terms, and time-averaging leaves its column space unchanged. Any v in the kernel stays invisible, because shifting c by v leaves the likelihood unchanged at every T and every σ. Section A adds the risk floor this forces on every estimator. □

Assumption 1 (Trajectory contract)**.** *The trajectory is a stationary Markov chain on D, independent of the noise, whose stationary density is bounded above and below. Its time-averaged normalized information concentrates around its mean at the parametric rate. Assumption 2 states the full form. For the reference class this is a consequence, not an assumption, since Corollary 3 proves that reflected Brownian motion at the snapshot rate satisfies it.*

Theorem 1 (Moving-array rate, trace form)**.** *For stationary reflected Brownian motion on D sampled at spacing Δ, after a concentration horizon $T_0 = \widetilde{O}(K/(\rho_\star \Delta))$ the moving array attains*

$$T\,\mathbb{E}\|\hat{c} - c\|^2 \;\longrightarrow\; \sigma^2\,\mathrm{tr}\big((I_K - B)^{-1}\big)\,\frac{|D|}{m}, \tag{1}$$

with matched constants in the upper and lower bounds. More generally, for any trajectory satisfying Assumption 1 with stationary density between $\rho_{\min}$ and $\rho_{\max}$, the scaled risk lies asymptotically between the same constant with $|D|$ replaced by $1/\rho_{\max}$ and by $1/\rho_{\min}$ (Section A).

The limit needs only $\gamma_B < 1$, which Lemma 2 gives at every finite K, while the horizon carries a factor $(1-\gamma_B)^{-2}$ that coverage bounds uniformly in K (Corollary 3). Explicitly, Equation (20) gives $T_0 \lesssim \tau\, M_\star(K)\,|D|/(\theta(1-\gamma_B))^2$ up to logarithms, with $\tau = \widetilde{O}((\rho_\star\Delta)^{-1})$ the minorization time of the sampled chain, $M_\star(K) = \sup_z \sum_{k\le K}\varphi_k(z)^2 = O(K)$ and θ the accepted relative error, so in wall-clock time $T_0\Delta = \widetilde{O}(K/\rho_\star)$ and the spacing drops out.

The time-averaged normalized information converges to its stationary mean $(I_K - B)/|D|$, so Σ_T inverts a matrix that grows linearly in T about that mean, and $T\,\mathrm{tr}\,\Sigma_T$ tends to the right side of Equation (1). Two things need control, the average must concentrate, which Assumption 1 guarantees and which fixes T_0, and the limit must be invertible, which is $\gamma_B < 1$ and therefore a statement about geometry rather than about time. The lower bound uses only the occupation measure, so it covers adaptive designs under the same density bound. The constants match because the stationary density bounds the limiting information from both sides.

Corollary 1 (Interpretable form)**.** $\mathrm{tr}((I_K - B)^{-1}) \le K/(1-\gamma_B)$*, so the constant is at most $\sigma^2\,K\,|D|/\big((1-\gamma_B)\,m\big)$. Section 5 bounds γ_B uniformly in K.*

5 Uniform conditioning from geometry

Let $\kappa(x)$ count the microphone offsets $p \in L$ for which the array fits inside Ω with the offset-p microphone at x, and let $\kappa_{\min} = \operatorname{ess\,inf}_{x\in\Omega}\kappa(x)$. The contraction factor is

$$\gamma_B(K) \;=\; \sup_{\|v\|=1} v^\top B v \;=\; 1 - \lambda_{\min}\Big(\frac{1}{m}\int_D \Phi(x)^\top\Phi(x)\,dx\Big). \tag{2}$$

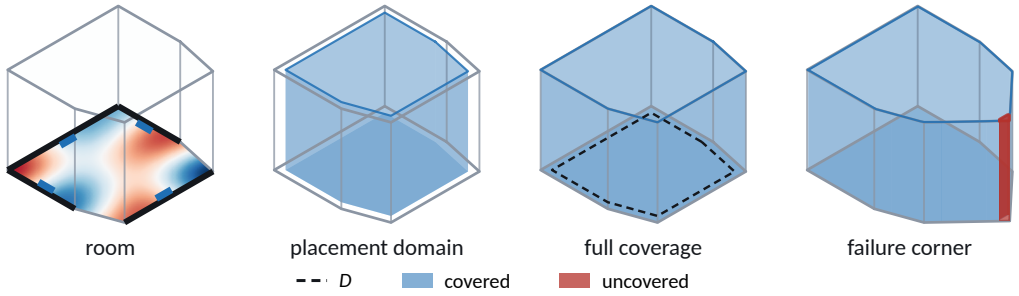

Figure 2: The shape of the room decides whether a moving array can observe all of it (Theorem 2). The field is a sum of room modes, one shown on the floor, and the four axis-aligned walls are drawn in black (room). The array can sit wherever its footprint fits (placement domain). A point is covered when some microphone can reach it, and walls at least one grid spacing longer than the aperture cover every point (full coverage). An oblique corner leaves a prism no microphone reaches, so γ_B approaches one (failure corner).

Its supremum over the modal dimension is a counting statement about apertures.

Proposition 2 (Coverage identity)**.** $\sup_K \gamma_B(K) = 1 - \kappa_{\min}/m$. *In particular* $\sup_K \gamma_B < 1$ *if and only if the translated apertures cover* Ω.

Only $\gamma_B = \lambda_{\max}(B) < 1$ is needed for $(I_K - B)^{-1}$ to exist, so $\lambda_{\min}(B) \to 0$ is harmless.

Lemma 1 (Slicing and the corner dichotomy)**.** *For a horizontal planar array in a prism* $\Omega = P \times [0, H]$*, coverage decomposes over horizontal cross-sections, so the problem is two-dimensional. In the plane, a neighborhood of a polygon vertex is locally coverable if and only if both outward edge normals at that vertex lie in one closed coordinate quadrant.*

Theorem 2 (Coverage from four axis facets)**.** *Let* P *be a convex polygon whose edges are maximal facets, with facets in all four axis directions, each of length at least* ns*. Then* $P = R + [0, ns]^2$ *for a compact convex* R*, the placement domain equals* $R + [0, s]^2$*, and the* n^2 *microphone translates of the placement domain cover every point of* P*. Hence* $\kappa_{\min} \geq 1$ *and, with Lemma 1 and Proposition 2,* $\sup_K \gamma_B \leq 1 - 1/m$ *for the prism* $P \times [0, H]$*, uniformly in* K.

Remark 1 (Equivalent hypothesis). For maximal-facet polygons the four-facet hypothesis is equivalent to corner quadrant-compatibility at every vertex together with axis-facet length at least ns. The corner language carries the failure taxonomy of Section D.

Figure 2 draws the geometry, and three structural consequences follow from the proof (Section B), not from computation. Only axis-facet lengths matter, because Minkowski addition by a square lengthens only axis facets and absorbs oblique geometry unchanged into R, which is why the sampled rooms of Section D show no oblique-edge condition. The constant is exactly ns, the aperture plus one grid spacing, because the placement domain must retain one further $s \times s$ cell whose n^2 translates tile the square summand. A room the array merely fits is not enough, because it must also step (Figure 7).

Remark 2 (Orientation). The square summand exists only in the array's frame. Rotating the array by $45°$ in a box gives $\kappa_{\min} = 0$, hence $\sup_K \gamma_B = 1$ by Proposition 2. Table 10 finds the same at every nonzero angle it tests, while at finite K the constant inflates but stays finite.

Corollary 2 (Aligned grid in a box)**.** *For an axis-aligned box with all sides at least* ns*,* $\kappa_{\min} = 1$ *and* $\sup_K \gamma_B = 1 - 1/m$*. The length hypothesis is sharp by the interval argument, since a side in* $[3s, 4s)$ *fits the array but cannot step and leaves an uncovered band. The polygon family of Section D (Figure 10) shows the same threshold off-axis.*

6 Validated numerics

The constants of Theorems 1 and 2 are geometric quantities of the room and the array. A gated finite-element pipeline computes each γ_B on three nested meshes with refinement ratio two, and

a declared branch rule chooses between Richardson extrapolation and the finest-mesh value and assigns an error bar. An exact polyhedral clipper, not quadrature, decides coverage of the placement domain. On the three rooms with closed-form spectra, every reported interval covers the exact value, and quadratic elements are more accurate than linear ones at matched cost (Tables 4 and 8). Section C gives the protocol, its runtime canaries and validations, and Section D lists every room the sampler rejected.

Ensemble uniformity. Figure 8 reports γ_B for twelve sampled convex prism rooms at three modal targets, snapped per room to the nearest spectral gap, and Table 5 lists the values. The room-to-room spread falls as K grows, and the means rise toward the ceiling $1 - 1/m$ of Section 5.

The certificate detects what finite K cannot. One sampled room, seed 4002, called room 13 below, fails the exact coverage certificate on a small part of its volume (Table 7). Every finite-K diagnostic passes, and its γ_B values sit inside the accepted-ensemble ranges at all three targets. By Proposition 2 its uncovered set forces $\sup_K \gamma_B = 1$. The certificate therefore decides a property no finite-K computation observes.

Effective dimension. The rate uses the whole spectrum of B, not its largest eigenvalue alone. That spectrum defines the effective dimension $d_{\text{eff}} := \operatorname{tr}((I_K - B)^{-1})$, the trace constant of Theorem 1, which is at least K because $B \succeq 0$. That trace is the A-optimality criterion of the parameter risk, not the capacity-style effective dimension of kernel methods, and Section 8 draws the contrast. The measured per-mode value d_{eff}/K sits well below the worst-case factor $1/(1 - \gamma_B)$ at both array scales (Table 6), and across the sampled ensemble it varies little (Table 5). That gap is why the rate carries d_{eff} and $K/(1-\gamma_B)$ survives only as Corollary 1. Mesh resolution sets a computable ceiling $K^*(h)$ on the resolvable K, associated with near-degenerate mode pairs and extrapolated one level ahead as a diagnostic (Section C.4 and Figure 9).

7 Experiments

Every experiment here is synthetic except Section 7.3. The synthetic fields come from validated FEM eigenfunctions of a sampled room (Figure 11), the noise is Gaussian with declared parameters, and no quantity is fit to data. Section E gives the protocol, including the gated eigenfunction evaluator.

Setup. The room is ensemble room 4001 with certified coverage, at $K = 50$ modes snapped to a spectral gap. The array is the 4×4 grid of $m = 16$ microphones at fixed leveled orientation, so $m < K$. The trajectory is reflected Brownian motion on the placement domain D, whose stationary density is uniform, so $\rho_{\min} = 1/|D|$. The simulation uses its Euler scheme, a reflected Gaussian random walk at the declared step, whose invariant law matches that density only up to discretization error (Section E). The estimator is the exact posterior mean under the declared prior.

7.1 The rank gap

The static baseline is optimized, not convenient. It minimizes exact Bayes risk over seeded positions of the same array in D, then refines around the best candidate. Even so, Proposition 1 leaves most modes with most of their content in the unidentified subspace at the declared observation time, and its Bayes risk stays orders of magnitude above the moving array's. The rank floor bounds the static curve from below at every T (Figures 3 and 4), so no observation time closes the gap.

7.2 The parameter-free asymptote

Figure 5 plots $T\,\mathbb{E}\|\hat{c}_T - c\|^2$ against T. The predicted asymptote $\sigma^2 d_{\text{eff}}/(\rho_{\min} m)$ contains no fitted parameter, since σ^2 is declared, $\rho_{\min} = 1/|D|$, and d_{eff} comes from the computed spectrum of B. The measured curve approaches the prediction from above, and the remaining excess is consistent with a $1/T$ correction. Theorem 1 fixes the horizon only up to its constant, and evaluating $\widetilde{O}(K/(\rho_\star \Delta))$ with constant one and the Payne–Weinberger bound places the observed knee inside it, by a criterion fixed after the run and disclosed (Figure 12).

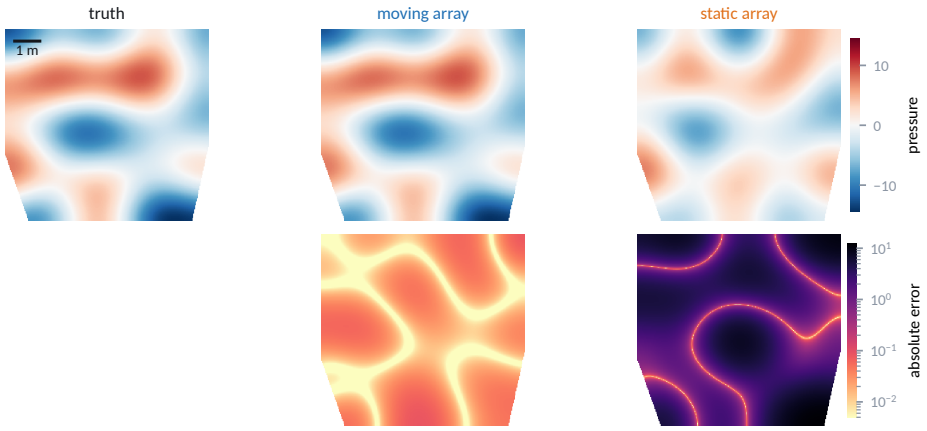

Figure 3: A static array misses whole regions of the field, however long it records. Top, the true field and its two reconstructions at equal microphone-seconds. Bottom, the error of each on one scale. The static error is structured, not random, because it is the unobservable subspace of Proposition 1, the part of the field the static position never observes, while the moving error is far smaller.

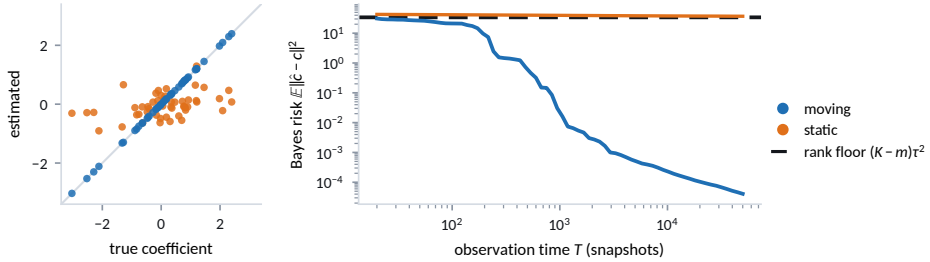

Figure 4: Moving the array removes an error that recording longer cannot. Left, perfect estimates lie on the diagonal, and the static array returns near zero for the modes it cannot see. Right, the static risk never drops below the rank floor however long it records, as Proposition 1 requires, while the moving risk falls through the floor and keeps falling, as Theorem 1 predicts.

Section E ranks raster, Lissajous and reflected-Brownian paths by exact Bayes risk under a declared speed-dependent noise model, and reports the raster's optimal speed. A measured $\sigma^2(v)$ would replace the declared model with no other change. At equal displacement per snapshot that ordering reverses, since a raster comes within a factor of one and a half of the constant of Theorem 1 while a random walk at walking speed explores too little of the room (Section E.4), and mis-known positions cost the moving array far less than the rank floor costs a static one (Section E.5).

7.3 A REAL ROOM

Proposition 1 constrains the rank of a measurement, so a hardware test of it needs no estimator, no prior and no modal basis. We recorded a 4×4 planar 16-microphone array at surveyed positions in a furnished room in everyday use, and froze the schedule, the gates and the analysis code before the first recorded take (Section G, Table 18). The array is at rest while each take is recorded, so the ladder compares one placement with many placements at equal microphone-seconds rather than a static array with a continuously moving one. At one position the array resolves only a few of its 16 possible directions at the declared screening level, the aperture collapse that Han & Kim (2026) report. Across 16 surveyed positions at equal microphone-seconds, the 17th singular value of the stacked measurement matrix exceeds the noise norm measured from the session's own silence recording. By Weyl's inequality the noiseless component then has rank at least 17 (Table 1). A static array of 16 microphones yields 16 rows, so its rank is at most 16, which is Proposition 1 outside the modal model. The measurement replicates across two sessions recorded two days apart, with the array disconnected in between, every position surveyed again and the order reversed. On the same recordings the held-out reconstruction error falls from one position to 16 in both sessions,

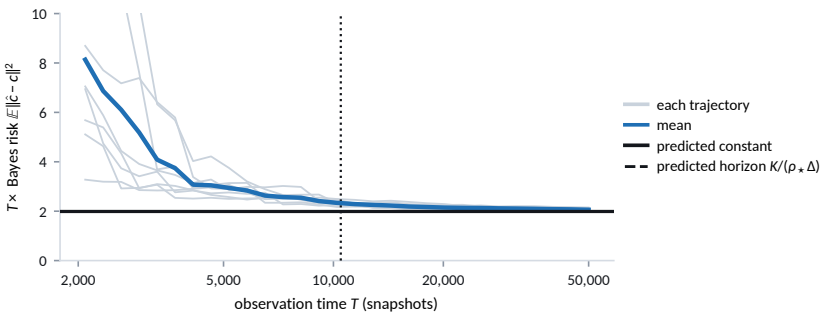

Figure 5: The measured risk settles on the predicted constant, which has no fitted parameter. A flat $T\times$risk means the risk falls exactly as $1/T$, so each doubling of observation time halves the error. The trajectories reach that constant near the predicted horizon, by a criterion fixed after the run. This checks both parts of Theorem 1, the rate and its closed-form constant, inside the assumed model.

	pre-registered		measured	
positions	prediction	noise free	session 1	session 2
1	3	3	3	3
2	6	6	6	6
4	10	10	10	10
8	14	14	14	14
rank of any static array of 16 microphones is at most 16				
16	16	17	17	17

Table 1: Motion raises the number of independent directions the recordings resolve. The count grows with the number of positions and at the largest rung exceeds the rank any static array of these microphones can reach, where it also exceeds the prediction despite higher noise than assumed, the direction of a furnished room carrying more structure than the rigid-box model. Both sessions give identical counts. This is Proposition 1 tested on real data, without a modal model.

as pre-registered, and ends below that of the static position chosen before recording as the best predicted under the rigid-box model (Figure 6). A second pre-registered test, which orders five sets of positions by predicted error, matches the measured order on nearly every resolvable pair. This is one room, pre-registered, and it claims no constant.

7.4 Comparison with third-party estimators

The rank gap is a property of acquisition, so it must not depend on our estimator. On identical synthetic data, the comparison runs our estimator against two independent moving-microphone estimators with public code, over a tenfold range of observation times (Table 2). With the same regularizer rule on each side, the moving array beats the best static position at equal microphone-seconds for both estimators that have one. Tuning each side separately leaves that ordering intact at the longer observation times (Section F). Tenfold observation time barely changes our static arm, while the gap to our moving arm widens. Our moving arm is the least accurate of the three moving arms at every observation time. Its prior is a single global variance over modal coefficients, and the true powers vary by orders of magnitude, whereas the published kernels are physically matched priors for a single-source field. Swapping in the oracle per-coefficient prior, a diagnostic rather than an estimator, reduces our error severalfold, which locates the difference in the prior, not in the information the samples carry. Section F gives the design and the per-bin curves.

8 Discussion

The result of this paper is a ceiling and the geometry that lifts it. Absent sparsity, a compact static array cannot identify more modes than it has microphones, and no estimator repairs that. Motion at equal microphone-seconds removes the obstruction, and the risk it leaves has a limit whose constant

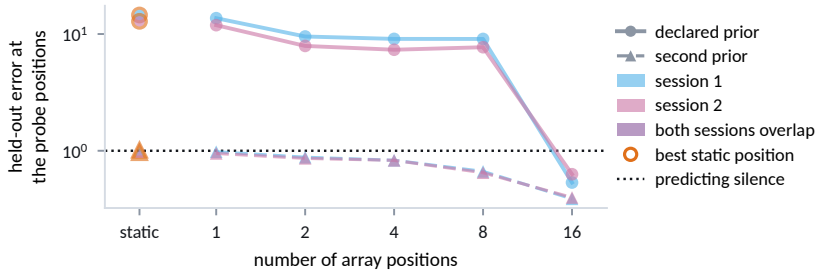

Figure 6: On real recordings, reconstruction error is lower at 16 positions than at one, the trend Theorem 1 predicts in an idealized room. Error is scored at microphones never used in the fit, where 1 means no better than predicting silence. Under the declared prior the error stays near its single-position level until the largest rung, then falls below silence and far below the best-predicted static position, while the second prior arm, declared with the analysis code and using a much smaller prior variance, falls steadily. Both sessions agree.

	ours		Brunnström et al. (2026)		Katzberg et al. (2021)	
obs. time	moving	static	moving	static	moving	static
1 s	0.912	0.971	0.821	1.894	0.855	†
3 s	0.859	0.965	0.685	1.351	0.739	†
10 s	0.492	0.952	0.359	1.295	0.434	†

Table 2: With the same regularizer rule on each side, moving beats the best static position at every observation time under both estimators that have a static counterpart. Here every static arm sits near or above the error of predicting silence, while every moving arm improves with observation time. The gain belongs to the acquisition, not to our estimator, as Proposition 1 implies. † marks an estimator released with no static counterpart.

is computable from the room and the array. The contraction factor γ_B decides the size of the limit and is bounded away from one exactly when the array's translates cover the room. The spectral gap $\rho_\star$ decides only how long the trajectory needs to reach that limit.

The floor is closed-form under the assumed model, so an estimator is graded against what the measurement made available rather than against another estimator. That grading separates an acquisition deficit, the geometry never sampled, from an estimator excess, what the algorithm gives away. The effective dimension d_{eff} carries the rate here, and it is the A-optimality trace of the parameter risk, not the effective dimension of kernel ridge regression. The two coincide in name only, and the contrast matters because d_{eff} inflates the risk above K while a capacity-style one is bounded by it.

What is open is sharpness rather than existence. The horizon of Theorem 1 holds up to constants and polylogarithmic factors, and pinning down the $O(K)$ variance proxy of Equation (20) would sharpen it. The second-order excess beyond that horizon is consistent with a $1/T$ decay that remains to be derived. The coverage identity fixes $\sup_K \gamma_B$ exactly, and the law by which the finite-K values approach it is open (Section E). Optimality over trajectories is untouched, and estimating B and d_{eff} from data would remove the need for a modal basis. The room class is convex and polyhedral, nothing is claimed beyond it, and extending the certificate past it is the next case. The hardware evidence is one room, one array and one mounting orientation, so replication is the next measurement to make. Each take is recorded with the array at rest, so the noise of a moving array is still to be measured, and that measurement would replace the declared speed-dependent model of the synthetic runs. Position error is tested only in the model, where independent jitter and a common offset leave floors of fourth and second order, far below the static one (Section E.5). The noise is independent and Gaussian throughout, except in the rank crossing, which uses the measured noise norm. A proved bound would replace the mesh-ladder error bar the computed constants carry.

The formulation generalizes past microphones, since the argument uses only the eigenbasis and the fixed offsets a rigid array samples from its own frame, so any such array in a convex polyhedral domain with known modes, under the same prior and contract, has the same visibility matrix, coverage count and trace limit.

AI USE STATEMENT

(This section is **required** and does not count toward the page limit.)

We used large language model assistants throughout this work, under author direction. Their use spans the required and recommended category this venue lists, proof development included, apart from translation, qualitative and thematic data analysis, survey design and transcription, none of which arise here.

We have reviewed AI-assisted work and take responsibility for it.

ETHICS STATEMENT

(This section is **recommended** and does not count toward the page limit.)

No ethical issues.

REPRODUCIBILITY STATEMENT

(This section is **recommended** and does not count toward the page limit.)

Assumptions are stated where each claim is made, and every proof is in Sections A and B. Each experimental appendix gives its rundown, its declared parameters, and what was decided after the data. Sections C and D for the numerics, Section E for the synthetic runs, Section F for the baselines, Section G for the hardware and its pre-registration chain.

Code and other data will be made public upon acceptance.

REFERENCES

Jonas Adler and Ozan Öktem. Learned primal-dual reconstruction. *IEEE transactions on medical imaging*, 37(6):1322–1332, 2018.

Hemant K Aggarwal, Merry P Mani, and Mathews Jacob. Modl: Model-based deep learning architecture for inverse problems. *IEEE transactions on medical imaging*, 38(2):394–405, 2018.

Byeongjoo Ahn, Karren Yang, Brian Hamilton, Jonathan Sheaffer, Anurag Ranjan, Miguel Sarabia, Oncel Tuzel, and Jen-Hao Rick Chang. Novel-view acoustic synthesis from 3d reconstructed rooms. *arXiv preprint arXiv:2310.15130*, 2023.

Thibaut Ajdler, Luciano Sbaiz, and Martin Vetterli. Dynamic measurement of room impulse responses using a moving microphone. *The Journal of the Acoustical Society of America*, 122(3): 1636–1645, 2007.

Giovanni S Alberti, Ernesto De Vito, Matti Lassas, Luca Ratti, and Matteo Santacesaria. Learning the optimal tikhonov regularizer for inverse problems. *Advances in Neural Information Processing Systems*, 34:25205–25216, 2021.

Akram Aldroubi, Carlos Cabrelli, Ilya Krishtal, and Ursula Molter. Dynamical sampling: A survey: A. aldroubi et al. *La Matematica*, 5(2):37, 2026.

Alen Alexanderian. Optimal experimental design for infinite-dimensional bayesian inverse problems governed by pdes: A review. *Inverse Problems*, 37(4):043001, 2021.

Alen Alexanderian, Noemi Petra, Georg Stadler, and Omar Ghattas. A-optimal design of experiments for infinite-dimensional bayesian linear inverse problems with regularized \ell_0-sparsification. *SIAM Journal on Scientific Computing*, 36(5):A2122–A2148, 2014.

Niccolo Antonello, Enzo De Sena, Marc Moonen, Patrick A Naylor, and Toon Van Waterschoot. Room impulse response interpolation using a sparse spatio-temporal representation of the sound field. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(10):1929–1941, 2017.

Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C Hansen. On instabilities of deep learning in image reconstruction and the potential costs of ai. *Proceedings of the National Academy of Sciences*, 117(48):30088–30095, 2020.

TMB Arachchige, S Jayawardhana, and CM Kalansuriya. Field measurement of sound insulation in buildings and building elements. 2018.

Simon Arridge, Peter Maass, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta numerica*, 28:1–174, 2019.

Ahmed Attia. A probabilistic approach to trajectory-based optimal experimental design. *arXiv preprint arXiv:2601.11473*, 2026.

Ahmed Attia, Sven Leyffer, and Todd Munson. Robust a-optimal experimental design for bayesian inverse problems. *arXiv preprint arXiv:2305.03855*, 2023.

Richard F Bass and Pei Hsu. Some potential theory for reflecting brownian motion in hölder and lipschitz domains. *The Annals of Probability*, pp. 486–508, 1991.

Terence Betlehem and Thushara D Abhayapala. Theory and design of sound field reproduction in reverberant rooms. *The Journal of the Acoustical Society of America*, 117(4):2100–2111, 2005.

Jesper Brunnström, Martin Bo Møller, Jan Østergaard, Shoichi Koyama, Toon van Waterschoot, and Marc Moonen. Sound field estimation with moving microphones using kernel ridge regression. *IEEE Transactions on Audio, Speech and Language Processing*, 2026.

Martin Burger and Samira Kabri. Learned regularization for inverse problems: Insights from a spectral model. *arXiv preprint arXiv:2312.09845*, 2023.

Carsten Carstensen and Joscha Gedicke. Guaranteed lower bounds for eigenvalues. *Mathematics of Computation*, 83(290):2605–2629, 2014.

Stefano Damiano, Kathleen MacWilliam, Valerio Lorenzoni, Thomas Dietzen, and Toon van Waterschoot. The trajectorir database: room acoustic recordings along a trajectory of moving microphones. *Journal on Audio, Speech, and Music Processing*, 2026(1):19, 2026.

Huayu Deng, Jinghui Zhong, Xiangming Zhu, Yunbo Wang, and Xiaokang Yang. Laser: Learning active sensing for continuum field reconstruction. *arXiv preprint arXiv:2604.19355*, 2026.

Thomas Deppisch, Jens Ahrens, Sebastià V Amengual Garí, and Paul Calamia. Spatial room impulse response identification from rotating equatorial microphone arrays. In *2024 32nd European Signal Processing Conference (EUSIPCO)*, pp. 116–120. IEEE, 2024.

Thomas Deppisch, Sebastià V Amengual Garí, Paul Calamia, and Jens Ahrens. Spatial room impulse response estimation from a moving microphone array. In *2025 33rd European Signal Processing Conference (EUSIPCO)*, pp. 91–95. IEEE, 2025.

Heinz Werner Engl, Martin Hanke, and Andreas Neubauer. *Regularization of inverse problems*, volume 375. Springer Science & Business Media, 1996.

Efren Fernandez-Grande. Sound field reconstruction using a spherical microphone array. *The Journal of the Acoustical Society of America*, 139(3):1168–1178, 2016.

Gene H Golub, Michael Heath, and Grace Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223, 1979.

Nina M Gottschling, Vegard Antun, Anders C Hansen, and Ben Adcock. The troublesome kernel: On hallucinations, no free lunches, and the accuracy-stability tradeoff in inverse problems. *SIAM Review*, 67(1):73–104, 2025.

Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *Proceedings of the 27th international conference on international conference on machine learning*, pp. 399–406, 2010.

Jeahn Han and Pyojin Kim. Why learning rediscovers the closed-form diagonal regularizer. *arXiv preprint arXiv:2609.09656*, 2026.

Per Christian Hansen. Analysis of discrete ill-posed problems by means of the l-curve. *SIAM review*, 34(4):561–580, 1992.

Longxiu Huang, A Martina Neuman, Sui Tang, and Yuying Xie. Convolutional dynamical sampling and some new results. *arXiv preprint arXiv:2406.15122*, 2024.

Jari P Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*. Springer, 2005.

Xenofon Karakonstantis, Diego Caviedes-Nozal, Antoine Richard, and Efren Fernandez-Grande. Room impulse response reconstruction with physics-informed deep learning. *The Journal of the Acoustical Society of America*, 155(2):1048–1059, 2024.

Fabrice Katzberg, Radoslaw Mazur, Marco Maass, Philipp Koch, and Alfred Mertins. Sound-field measurement with moving microphones. *The Journal of the Acoustical society of America*, 141 (5):3220–3235, 2017.

Fabrice Katzberg, Marco Maass, and Alfred Mertins. Spherical harmonic representation for dynamic sound-field measurements. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 426–430. IEEE, 2021.

Bartek T Knapik, Aad W Van Der Vaart, and J Harry van Zanten. Bayesian inverse problems with gaussian priors. *The Annals of Statistics*, pp. 2626–2657, 2011.

Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(2), 2008.

Karl Kunisch and Thomas Pock. A bilevel optimization approach for parameter learning in variational models. *SIAM Journal on Imaging Sciences*, 6(2):938–983, 2013.

Jie Liu, Heng Ouyang, Xu Han, and Guirong Liu. Optimal sensor placement for uncertain inverse problem of structural parameter estimation. *Mechanical Systems and Signal Processing*, 160: 107914, 2021.

Xuefeng Liu and Shin'ichi Oishi. Verified eigenvalue evaluation for the laplacian over polygonal domains of arbitrary shape. *SIAM Journal on Numerical Analysis*, 51(3):1634–1654, 2013.

Rémi Mignot, Gilles Chardon, and Laurent Daudet. Low frequency interpolation of room impulse responses using compressed sensing. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(1):205–216, 2013.

Federico Miotello, Luca Comanducci, Mirco Pezzoli, Alberto Bernardini, Fabio Antonacci, and Augusto Sarti. Reconstruction of sound field through diffusion models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1476–1480. IEEE, 2024.

Charles B Morrey Jr and Louis Nirenberg. On the analyticity of the solutions of linear elliptic systems of partial differential equations. *Communications on Pure and Applied Mathematics*, 10 (2):271–290, 1957.

Esa Nummelin. A splitting technique for harris recurrent markov chains. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 43(4):309–318, 1978.

Lawrence E Payne and Hans F Weinberger. An optimal poincaré inequality for convex domains. *Archive for Rational Mechanics and Analysis*, 5(1):286–292, 1960.

Boaz Rafaely. Spatial sampling and beamforming for spherical microphone arrays. *2008 Hands-Free Speech Communication and Microphone Arrays*, pp. 5–8, 2008.

Juliano GC Ribeiro, Ryo Matsuda, and Jorge Trevino. Exterior sound field estimation based on physics-constrained kernel. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 15167–15171. IEEE, 2026.

Sudeep Salgia and Animesh Kumar. Bandlimited spatiotemporal field sampling with location and time unaware mobile sensors. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4574–4578. IEEE, 2018.

Jonas M Schmid, Efren Fernandez-Grande, Manuel Hahmann, Caglar Gurbuz, Martin Eser, and Steffen Marburg. Spatial reconstruction of the sound field in a room in the modal frequency range using bayesian inference. *The Journal of the Acoustical Society of America*, 150(6):4385–4394, 2021.

Andrew M Stuart. Inverse problems: a bayesian perspective. *Acta numerica*, 19:451–559, 2010.

Sakari Tervo, Jukka Pätynen, Antti Kuusinen, and Tapio Lokki. Spatial decomposition method for room impulse responses. *J. Audio Eng. Soc*, 61(1/2):17–28, 2013.

Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

Jayakrishnan Unnikrishnan and Martin Vetterli. Sampling and reconstruction of spatial fields using mobile sensors. *IEEE Transactions on Signal Processing*, 61(9):2328–2340, 2013.

Thach Pham Vu and Hervé Lissek. Low frequency sound field reconstruction in a non-rectangular room using a small number of microphones. *Acta Acustica*, 4(2):5, 2020.

Shunxi Xu and Craig T Jin. Hierarchical sparse sound field reconstruction with spherical and linear microphone arrays. In *2025 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 83–88. IEEE, 2025.

Shunxi Xu, Thushara Abhayapala, and Craig T Jin. A unified svd-modal solution for sparse sound field reconstruction with hybrid spherical-linear microphone arrays. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 14722–14726. IEEE, 2026.

symbol	meaning
Ω, P, H	room, and its floor polygon and height when $\Omega = P \times [0, H]$
(λ_k, φ_k)	Neumann Laplace eigenpairs, with φ_k orthonormal in $L^2(\Omega)$
K	number of modes in the field
c, τ^2	modal coefficient vector and its prior variance per mode, $c \sim \mathcal{N}(0, \tau^2 I_K)$
Σ_0	prior covariance in the proofs, equal to $\tau^2 I_K$ in the main text
n, s, m	grid side, grid spacing and number of microphones, $m = n^2$
$p_1, \dots, p_m$, L	microphone offsets and their set
D, $\lvert D\rvert$	placement domain and its volume
x_t, T	array position at snapshot t, and observation time in snapshots
$\Phi(x)$, $G(x)$	mode-sample matrix at the microphones of an array at x, and $G(x) = \Phi(x)^\top \Phi(x)$
y_t, ε_t, σ^2	snapshot measurement, its noise, and the noise variance
$\rho_\star$, Δ	spectral gap of reflected Brownian motion on D, and its sampling interval
∇^2	Laplacian, so the reflected Brownian generator is $\frac{1}{2}\nabla^2$
$\rho_{\min}$, $\rho_{\max}$	lower and upper bounds on the stationary density of the trajectory
B, γ_B	visibility matrix and its largest eigenvalue
d_{eff}	effective dimension $\operatorname{tr}((I_K - B)^{-1})$
$\hat{c}$, Σ_T	posterior mean and posterior covariance after T snapshots
T_0	concentration horizon of Theorem 1
$\kappa(x)$, $\kappa_{\min}$	coverage count at x and its essential infimum over Ω

Table 3: Notation of the main text and the proofs.

A Proofs for Section 4

Table 3 collects the notation of Sections 3–5, and each statement keeps only the assumptions it needs.

A.1 Information geometry

Let $\Omega \subset \mathbb{R}^3$ be the unit-volume room, with $L^2(\Omega)$-orthonormal Neumann eigenfunctions $(\varphi_k)_{k \geq 1}$. The array offsets are $p_1, \dots, p_m$, and

$$D = \bigcap_{i=1}^{m} (\Omega^\circ - p_i)$$

is the nonempty placement domain. For $x \in D$, define

$$\Phi_{ik}(x) = \varphi_k(x + p_i), \qquad G(x) = \Phi(x)^\top \Phi(x), \qquad A(x) = \frac{1}{m} G(x).$$

For a density ρ on D, write $\Sigma_\rho = \int_D A(x)\rho(x)\,dx$. The uniform law on D is $\mu(dx) = dx/|D|$.

Let

$$\kappa(z) = \#\{i : z - p_i \in D\}.$$

This is the coverage count κ of Section 5 written with the open placement domain. The two forms differ only on translated boundaries, so

$$\operatorname*{ess\,inf}_{z \in \Omega} \kappa(z) = \kappa_{\min}.$$

The visibility matrix and effective dimension are

$$B = I_K - \frac{1}{m} \int_D G(x)\,dx, \qquad \gamma_B = \lambda_{\max}(B), \qquad d_{\text{eff}} = \operatorname{tr}\big((I_K - B)^{-1}\big). \tag{3}$$

Three standard facts enter. First, Neumann eigenfunctions on a bounded Lipschitz domain are bounded, and finite linear combinations are real analytic in the interior. Bass & Hsu (1991) give the heat-kernel bounds and Morrey Jr & Nirenberg (1957) the interior analytic regularity. Second, reflected Brownian motion on a bounded Lipschitz domain exists and has the Neumann heat kernel as its transition density (Bass & Hsu, 1991). Third, if ν_1 is the first nonzero Neumann eigenvalue

of D, then the reflected Brownian semigroup with generator $\frac{1}{2}\nabla^2$ has spectral gap $\rho_\star = \nu_1/2$, and convexity gives

$$\rho_\star \geq \frac{\pi^2}{2\,\mathrm{diam}(D)^2} \tag{4}$$

by the Payne–Weinberger inequality (Payne & Weinberger, 1960). Finally, with

$$M_\star(K) := \sup_{z\in\Omega} \sum_{k=1}^{K} \varphi_k(z)^2,$$

heat-kernel bounds and Weyl's law give $M_\star(K) = O(K)$ in three dimensions. For an axis-aligned rectangular box, $M_\star(K) \leq 8K$.

Lemma 2 (Structure of B)**.** *For every finite K,*

$$I_K - B = \int_\Omega \frac{\kappa(z)}{m}\,\varphi(z)\varphi(z)^\top dz, \qquad \Sigma_\mu = \frac{I_K - B}{|D|}, \tag{5}$$
$$0 \preceq B \preceq I_K, \qquad \gamma_B < 1.$$

Moreover,

$$\gamma_B \leq 1 - \frac{\kappa_{\min}}{m}, \qquad K \leq d_{\text{eff}} \leq \frac{K}{1-\gamma_B}. \tag{6}$$

Proof. Translate each microphone term:

$$\int_D G(x)\,dx = \sum_{i=1}^{m} \int_{D+p_i} \varphi(z)\varphi(z)^\top dz = \int_\Omega \kappa(z)\varphi(z)\varphi(z)^\top dz.$$

Orthonormality gives Equation (5), and $0 \leq \kappa \leq m$ gives $0 \preceq B \preceq I_K$. For a unit vector v, let $f_v = \sum_k v_k\varphi_k$. Then

$$v^\top(I_K - B)v = \frac{1}{m}\sum_i \int_{D+p_i} f_v(z)^2\,dz.$$

If this were zero, f_v would vanish on the nonempty open set $D + p_1$. Interior analyticity would then force $f_v \equiv 0$, contradicting $\|f_v\|_{L^2} = 1$. Thus $I_K - B \succ 0$ and $\gamma_B < 1$. The first inequality in Equation (6) follows from $v^\top Bv = \int(1-\kappa/m)f_v^2 \leq 1 - \kappa_{\min}/m$. The eigenvalues of $I_K - B$ lie in $[1-\gamma_B, 1]$, which gives the bounds on d_{eff}. □

A.2 Posterior risk and the static floor

A design is a sequence $x_1, \ldots, x_T \in D$. It is deterministic, randomized independently of the data, or adaptive, with x_t measurable with respect to past observations and independent auxiliary randomness. The observations are

$$y_t = \Phi(x_t)c + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma^2 I_m), \qquad c \sim \mathcal{N}(0, \Sigma_0),$$

with $\Sigma_0 \succ 0$. Define

$$I_T = \sigma^{-2}\sum_{t=1}^{T} G(x_t), \qquad \Sigma_T = (\Sigma_0^{-1} + I_T)^{-1}. \tag{7}$$

Lemma 3 (Posterior under adaptive designs)**.** *Conditioned on $(x_{1:T}, y_{1:T})$, the coefficient vector is Gaussian with covariance Σ_T and mean*

$$\hat{c}_T = \Sigma_T\sigma^{-2}\sum_{t=1}^{T} \Phi(x_t)^\top y_t.$$

Hence every estimator $\tilde{c}_T$ satisfies

$$\mathbb{E}\|\tilde{c}_T - c\|^2 \geq \mathbb{E}\,\mathrm{tr}\,\Sigma_T, \tag{8}$$

with equality for the posterior mean.

Proof. Condition on the history before snapshot t. The adaptive rule makes x_t fixed under this conditioning, so the usual linear-Gaussian update applies. Induction over t gives Equation (7) and the stated mean. The policy probabilities depend only on already observed data and auxiliary randomness, so they add no c-dependent factor to the posterior. The conditional mean minimizes squared error, giving Equation (8). □

Proposition 3 (Static risk floor, Proposition 1)**.** *For a static design $x_t \equiv x$ with $m < K$, let $N = \ker G(x)$. Then $\dim N \geq K - m$, and for every estimator*

$$\mathbb{E}\|\tilde{c}_T - c\|^2 \geq \sum_{j=1}^{\dim N} \frac{1}{v_j^\top \Sigma_0^{-1} v_j} \geq (K-m)\lambda_{\min}(\Sigma_0), \tag{9}$$

for any orthonormal basis (v_j) of N. If $\Sigma_0 = \tau^2 I_K$, the floor is $(K-m)\tau^2$.

Proof. Since $G(x) = \Phi(x)^\top \Phi(x)$ has rank at most m, vectors in N remain unobserved for every T. By Lemma 3, it is enough to lower-bound $\operatorname{tr} \Sigma_T$. For unit $v \in N$ and $M = \Sigma_0^{-1} + I_T$,

$$v^\top M^{-1} v \geq \frac{1}{v^\top M v} = \frac{1}{v^\top \Sigma_0^{-1} v},$$

where the inequality is Cauchy–Schwarz and $I_T v = 0$. Summing over an orthonormal basis of N gives Equation (9). □

The floor holds for every estimator, linear or not, because the prior has full support and the likelihood is constant along N. A sparsity assumption escapes it by shrinking the set of admissible coefficients, not by measuring more, and this paper makes none.

A.3 Lower and upper trace bounds

The map $M \mapsto \operatorname{tr}(M^{-1})$ is convex and order-reversing on positive-definite matrices. One convenient variational form is

$$\operatorname{tr}(M^{-1}) = \max_{X = X^\top} \left\{ 2 \operatorname{tr} X - \operatorname{tr}(XMX) \right\}. \tag{10}$$

Theorem 3 (Occupation-measure lower bound)**.** *Suppose a design, adaptive or not, satisfies*

$$\mathbb{E}\left[\sum_{t=1}^{T} G(x_t) \right] \preceq \rho_{\max} T \int_D G(x)\, dx. \tag{11}$$

Then every estimator satisfies

$$\mathbb{E}\|\tilde{c}_T - c\|^2 \geq \operatorname{tr}\left(\Sigma_0^{-1} + \frac{\rho_{\max} m T}{\sigma^2} (I_K - B) \right)^{-1}, \tag{12}$$

and therefore

$$\liminf_{T \to \infty} T\, \mathbb{E}\|\tilde{c}_T - c\|^2 \geq \frac{\sigma^2 d_{\text{eff}}}{\rho_{\max} m}. \tag{13}$$

A stationary design with one-time density $\rho \leq \rho_{\max}$ satisfies Equation (11)*.*

Proof. Apply Lemma 3, Jensen's inequality, Equation (10), and Equation (11). This gives Equation (12). Multiplying by T and letting $T \to \infty$ gives Equation (13). □

Assumption 2 (Trajectory contract, full form of Assumption 1)**.** *The trajectory $(X_t)_{t \geq 1}$ is a stationary Markov chain on D, independent of (c, ε), with stationary density ρ satisfying $0 < \rho_{\min} \leq \rho \leq \rho_{\max} < \infty$. Let*

$$\Sigma_\infty = \int_D A(x)\rho(x)\, dx, \qquad E_T = \frac{1}{T} \sum_{t=1}^{T} \big(A(X_t) - \Sigma_\infty \big).$$

For all $T \geq 1$ and $\delta \in (0,1)$,

$$\mathbb{P}(\|E_T\| \leq \epsilon_K(T,\delta)) \geq 1 - \delta, \qquad \epsilon_K(T,\delta) = \sqrt{\frac{v_K(T,\delta)}{T}} + \frac{b_K(T,\delta)}{T}, \tag{14}$$

where v_K and b_K grow at most polylogarithmically in T and $1/\delta$ for fixed K.

The density bounds and Lemma 2 give

$$\rho_{\min}(I_K - B) \preceq \Sigma_\infty \preceq \rho_{\max}(I_K - B), \qquad \operatorname{tr}(\Sigma_\infty^{-1}) \le \frac{d_{\text{eff}}}{\rho_{\min}}. \tag{15}$$

Theorem 4 (Upper bound). *Under Assumption 2, let $\hat{c}_T$ be the posterior mean. If $\theta, \delta \in (0,1)$ and $\epsilon_K(T,\delta) \le \theta\rho_{\min}(1-\gamma_B)$, then*

$$\mathbb{E}\|\hat{c}_T - c\|^2 \le \frac{\sigma^2 \operatorname{tr}(\Sigma_\infty^{-1})}{(1-\theta)mT} + \operatorname{tr}(\Sigma_0)\,\delta \le \frac{\sigma^2 d_{\text{eff}}}{(1-\theta)\rho_{\min} mT} + \operatorname{tr}(\Sigma_0)\,\delta. \tag{16}$$

Consequently,

$$\limsup_{T\to\infty} T\,\mathbb{E}\|\hat{c}_T - c\|^2 \le \frac{\sigma^2}{m}\operatorname{tr}(\Sigma_\infty^{-1}) \le \frac{\sigma^2 d_{\text{eff}}}{\rho_{\min} m}. \tag{17}$$

Proof. On the event $\|E_T\| \le \epsilon_K(T,\delta)$,

$$I_T = \frac{Tm}{\sigma^2}(\Sigma_\infty + E_T) \succeq \frac{Tm}{\sigma^2}(1-\theta)\Sigma_\infty,$$

so $\operatorname{tr}\Sigma_T \le \sigma^2 \operatorname{tr}(\Sigma_\infty^{-1})/((1-\theta)mT)$. On the complement, $\Sigma_T \preceq \Sigma_0$. Taking expectations gives Equation (16). For Equation (17), take $\delta_T = T^{-2}$ and $\theta_T = \epsilon_K(T,\delta_T)/(\rho_{\min}(1-\gamma_B))$; the contract gives $\theta_T \to 0$. □

A.4 Reflected Brownian motion

Let $X(s)$ be stationary reflected Brownian motion on D with generator $\frac{1}{2}\nabla^2$, and sample it at spacing $\Delta > 0$: $X_t = X(t\Delta)$. Then $\rho = 1/|D|$ and $\Sigma_\infty = (I_K - B)/|D|$.

Fix $t_0 > 0$ and let

$$\tau_D = 2t_0 + \rho_\star^{-1}\log\bigl(2C_{\text{uc}}(t_0)^2\bigr), \qquad \tau = \lceil \tau_D/\Delta \rceil, \tag{18}$$

where $C_{\text{uc}}(t_0) = \sup_x \|K_{t_0}(x,\cdot)\|_{L^2(\mu)} < \infty$.

Proposition 4 (Concentration of the snapshot chain). *Let $F(x) = A(x) - \Sigma_\infty$ and $V = \|\int_D F(x)^2\,d\mu(x)\|$. Then $\sup_x \|F(x)\| \le M_\star(K)$ and $V \le M_\star(K)/|D|$. For every $T \ge \tau$ and $\delta \in (0,1)$,*

$$\|E_T\| \le 2\sqrt{\frac{\tau L V \ell}{T}} + \frac{5}{3}\frac{\tau L M_\star(K)\ell}{T} \tag{19}$$

with probability at least $1-\delta$, where

$$\ell = \log\frac{4K\tau}{\delta}, \qquad L = \left\lceil \log_2 \frac{6T}{\delta} \right\rceil.$$

Proof. The heat-kernel semigroup and the spectral gap imply a uniform minorization. Indeed, for $s = 2t_0 + u \ge \tau_D$,

$$|K_s(x,y) - 1| = |\langle K_{t_0}(x,\cdot), (P_u - \Pi)K_{t_0}(y,\cdot)\rangle| \le C_{\text{uc}}(t_0)^2 e^{-\rho_\star u} \le \tfrac{1}{2}.$$

Hence the τ-step kernel satisfies $P^\tau(x,\cdot) \ge \frac{1}{2}\mu(\cdot)$.

Split the sampled chain into its τ interleaved phases. Nummelin splitting (Nummelin, 1978) represents each phase as regeneration blocks with geometric lengths, generated from disjoint randomness and hence independent conditional on those lengths. Two further facts are not part of the standard statement. First, invariance $\mu P^\tau = \mu$ reads $\mu = \beta\mu + (1-\beta)\mu R$ for the residual kernel R, so $\mu R = \mu$ and every state in every block has marginal law μ. Each block sum Z_j is therefore centered, with $\mathbb{E}[Z_j \mid N] = N_j \int F\,d\mu = 0$. Second, for symmetric $V_1,\dots,V_n$ one has $(\sum_i V_i)^2 \preceq n\sum_i V_i^2$, which with the marginals gives $\mathbb{E}[Z_j^2 \mid N] \preceq N_j^2 \int F^2\,d\mu$. With probability at least $1 - \delta/2$, all block lengths needed up to time T are at most L. Conditional on this event, each block sum has norm at most $LM_\star(K)$ and total matrix variance at most Ln_rV in a phase of length n_r. Matrix Bernstein (Tropp, 2012) followed by a union bound over the τ phases gives Equation (19). The variance bound uses $A(x)^2 \preceq M_\star(K)A(x)$ and $\Sigma_\infty = (I_K - B)/|D| \preceq I_K/|D|$. □

Corollary 3 (Contract and concentration horizon for reflected Brownian motion)**.** *The snapshot chain satisfies Assumption 2 with $\rho_{\min} = \rho_{\max} = 1/|D|$. A sufficient horizon is the least $T \geq \tau$ such that*

$$T \geq \max\left\{ \frac{16\tau LV\ell|D|^2}{\theta^2(1-\gamma_B)^2}, \frac{10\tau LM_\star(K)\ell|D|}{3\theta(1-\gamma_B)} \right\}. \tag{20}$$

For a fixed room, array, snapshot spacing, and confidence parameters, $M_\star(K) = O(K)$ gives

$$T_0 = \widetilde{O}\left(\frac{K}{\rho_\star \Delta}\right) \tag{21}$$

whenever the coverage condition keeps $(1-\gamma_B)^{-1}$ uniformly bounded. Because L grows logarithmically in T, Equation (20) *is implicit in T, and solving it changes T_0 only by the logarithmic factor that Equation* (21) *already absorbs.*

Proof. Equation (19) has the form required by Assumption 2. Requiring each term in Equation (19) to be at most $\theta(1-\gamma_B)/(2|D|)$ gives Equation (20). Equation (21) then follows from $V \leq M_\star(K)/|D|$, $M_\star(K) = O(K)$, and $\tau = \widetilde{O}((\rho_\star\Delta)^{-1})$ on a fixed domain. □

A.5 Moving-array rate

Theorem 5 (Moving-array trace rate, Theorem 1)**.** *For stationary reflected Brownian motion sampled at spacing $\Delta > 0$,*

$$\lim_{T\to\infty} T\,\mathbb{E}\|\hat{c}_T - c\|^2 = \frac{\sigma^2|D|}{m}\, d_{\text{eff}} = \sigma^2\,\mathrm{tr}\big((I_K - B)^{-1}\big)\frac{|D|}{m}. \tag{22}$$

The same constant lower-bounds every estimator for every design whose expected occupation measure is uniform on D. More generally, for a stationary trajectory satisfying Assumption 2,

$$\frac{\sigma^2 d_{\text{eff}}}{\rho_{\max} m} \leq \liminf_T T\,\mathbb{E}\|\hat{c}_T - c\|^2 \leq \limsup_T T\,\mathbb{E}\|\hat{c}_T - c\|^2 \leq \frac{\sigma^2 d_{\text{eff}}}{\rho_{\min} m}. \tag{23}$$

Proof. For the uniform reflected-Brownian law, Theorem 3 has $\rho_{\max} = 1/|D|$, while Theorem 4 and Corollary 3 have $\Sigma_\infty = (I_K - B)/|D|$. Their asymptotic constants coincide, giving Equation (22). The density-bounded statement is the same pair of bounds without setting $\rho_{\min} = \rho_{\max}$. □

Corollary 4 (Interpretable form, Corollary 1)**.** *Since $d_{\text{eff}} \leq K/(1-\gamma_B)$,*

$$\lim_{T\to\infty} T\,\mathbb{E}\|\hat{c}_T - c\|^2 \leq \frac{\sigma^2 K|D|}{(1-\gamma_B)m}.$$

On covered rooms, Section 5 bounds $(1-\gamma_B)^{-1}$ independently of K.

Remark 3 (Simulation versus theorem). Section E uses the reflected Euler walk $X_{t+1} = \mathrm{refl}(X_t + s_v\xi_t)$ rather than exact snapshots of reflected Brownian motion. Theorem 5 applies to the exact process. The Euler walk is used as a numerical approximation, and exact stationarity of its invariant law is not assumed.

B Coverage proofs

For the geometric arguments, let

$$S = [0,(n-1)s]^2, \qquad L = \{0,s,\dots,(n-1)s\}^2, \qquad D_2 = P \ominus S.$$

Here $P \ominus S = \{y : y + S \subseteq P\}$ is the closed erosion. Define

$$\kappa(x) = \#\{p \in L : x \in D_2 + p\}, \qquad \kappa_{\min} = \operatorname*{ess\,inf}_{x\in P} \kappa(x),$$

as in Section 5. The closed and open placement domains differ only on boundaries, so this count has the same essential infimum as the one entering the information operator.

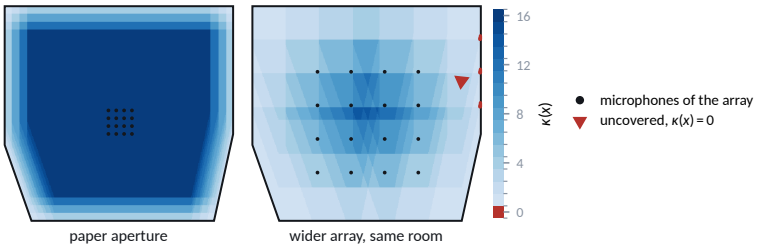

Figure 7: Coverage decides whether the contraction factor stays bounded away from one as the model grows, which is Proposition 2. A point is covered when at least one microphone can reach it. With the array of the paper every point is covered, so $\sup_K \gamma_B$ stays below one (paper aperture). A wider array in the same room cannot reach small patches along one wall, and γ_B then tends to one as K grows (wider array, same room).

B.1 Coverage identity

Proof of Proposition 2. Let $V_K = \operatorname{span}\{\varphi_1, \dots, \varphi_K\}$ and set $q(x) = 1 - \kappa(x)/m$. Lemma 2 gives

$$\gamma_B(K) = \sup_{\substack{f \in V_K \\ \|f\|_{L^2}=1}} \int_\Omega q(x)|f(x)|^2 \, dx.$$

The spaces V_K are nested and their union is dense in $L^2(\Omega)$. Since $q \in L^\infty(\Omega)$, the quadratic form above is continuous in L^2. Therefore

$$\sup_K \gamma_B(K) = \sup_{\|f\|_{L^2}=1} \int_\Omega q|f|^2 = \|q\|_{L^\infty} = 1 - \frac{\kappa_{\min}}{m}.$$

Thus $\sup_K \gamma_B < 1$ exactly when $\kappa(x) \geq 1$ almost everywhere. For polygonal rooms, the closed-domain formulation above upgrades this to geometric coverage of the closure. □

B.2 Slicing and corner compatibility

Proof of Lemma 1. For a horizontal planar array in $P \times [0, H]$, feasibility depends only on the horizontal coordinates. The three-dimensional coverage problem is therefore the same two-dimensional problem on every horizontal slice.

Now consider a vertex v of P with incident outward normals u_1, u_2. Suppose the microphone at offset $p \in L \subset S$ is placed at v. The anchor $v - p$ is feasible near this corner if and only if

$$h_S(u_j) = u_j \cdot p, \qquad j = 1, 2,$$

because $u_j \cdot v$ is the supporting value of the incident edge and $p \in S$ already implies $u_j \cdot p \leq h_S(u_j)$. Thus one offset must maximize both linear functionals over the square S.

A square has a common maximizing vertex for two directions exactly when those directions lie in one closed coordinate quadrant. Hence the corner is locally coverable if and only if u_1 and u_2 lie in a common closed quadrant. If the condition holds, the corresponding corner microphone covers a neighborhood of v inside P, provided the array placed at v clears the nonincident facets. The statement is local in this sense, and the length hypothesis of Theorem 2 supplies that clearance at every vertex. If it fails, no offset is feasible at v, so $v \notin D_2 + p$ for every $p \in L$. Each translate is closed and there are finitely many, so their union is closed and misses an open ball around v. That ball meets P in a set of positive area on which $\kappa = 0$, so $\operatorname{ess\,inf} \kappa = 0$ and Proposition 2 gives $\sup_K \gamma_B = 1$. □

B.3 Four-axis-facet condition

Proof of Theorem 2. Let the four axis-aligned edge lengths of P be $\ell_b, \ell_r, \ell_t, \ell_l$, each at least ns.

Square summand. Traverse the maximal edge vectors of P counterclockwise. Shorten each of the four axis-aligned vectors by ns and leave each oblique edge unchanged. The removed vectors sum to zero, so the residual vectors still form a cyclically ordered closed polygonal chain. Hence they are the edge vectors of a compact convex set R, possibly degenerate. Restoring the removed square gives

$$P = R + [0, ns]^2. \tag{24}$$

Erosion. For compact convex sets Q, S, support functions give

$$(Q + S) \ominus S = Q.$$

Indeed, $y + S \subseteq Q + S$ implies $u \cdot y + h_S(u) \le h_Q(u) + h_S(u)$ for every u, hence $y \in Q$. Since $[0, ns]^2 = [0, s]^2 + [0, (n-1)s]^2$, Equation (24) yields

$$D_2 = P \ominus [0, (n-1)s]^2 = R + [0, s]^2. \tag{25}$$

Coverage. The n^2 cells $[0, s]^2 + (is, js)$, $i, j \in \{0, \dots, n-1\}$, cover $[0, ns]^2$. Therefore

$$\bigcup_{p \in L} (D_2 + p) = R + \bigcup_{p \in L} ([0, s]^2 + p) = R + [0, ns]^2 = P.$$

Thus $\kappa_{\min} \ge 1$. Proposition 2 gives $\sup_K \gamma_B \le 1 - 1/m$, and the slicing part of Lemma 1 lifts the result to $P \times [0, H]$. □

The threshold ns is sharp for this family. The constructed examples in Appendix D have positive uncovered area for an axis facet shorter than ns and zero uncovered area at and above ns.

C Numerical validation and reproducibility

The numerical pipeline approximates Neumann eigenfunctions with quadratic tetrahedral finite elements and evaluates the restricted information matrix on the array placement domain. Rooms are normalized to $|\Omega| = 1$. The numerical uncertainty reported here is an empirical discretization estimate, not a rigorous a posteriori FEM bound. The placement domain cuts through mesh elements, so each element is classified as inside, outside or cut. A cut element is clipped into sub-tetrahedra, and a degree-four rule on each integrates the product of two quadratic shape functions exactly, so the restriction adds no quadrature error to the discretization error. The coverage test is separate and uses polyhedral clipping rather than quadrature.

C.1 Mesh ladder and uncertainty rule

Each room is solved on three uniformly refined meshes with refinement ratio two. Let $\gamma_0, \gamma_1, \gamma_2$ denote the coarse, middle, and fine values of γ_B. When

$$\frac{\gamma_0 - \gamma_1}{\gamma_1 - \gamma_2} > 0, \qquad p_3 = \log_2\left(\frac{\gamma_0 - \gamma_1}{\gamma_1 - \gamma_2}\right)$$

is defined. If $p_3 \in [3, 5]$ and the fourth-order Richardson correction exceeds the numerical floor, the rule reports

$$\gamma_R = \gamma_2 + \frac{\gamma_2 - \gamma_1}{15}$$

with uncertainty $\max(|\gamma_R - \gamma_2|, \text{floor})$, the RICH branch of Table 4. Otherwise it reports γ_2 with uncertainty $\max(|\gamma_2 - \gamma_1|, \text{floor})$, the DRIFT branch. The floor is 2×10^{-6}, the level at which assembly and the eigensolver stop resolving γ_B. The order of a restricted integral of finite-element eigenfunctions need not be four, because the orthogonality that doubles the order of an eigenvalue holds on the whole room and not on the placement domain. The rule therefore extrapolates only where the observed order supports it, and most rows of Table 4 fall to the DRIFT branch. Of the 36 sampled-room values of Table 5, 17 use the RICH branch and 19 the DRIFT branch. Modal cutoffs are snapped to spectral gaps, so the realized K can differ slightly from the target.

Table 4 compares the rule with closed-form values on the cube and two rectangular-box scalings. The box is a room of 6.4 by 6.9 by 4.7 m scaled to unit volume. Its faithful scaling keeps the 42 mm microphone spacing at that scale, and its fixed-relative scaling keeps the spacing of the cube, 0.042 in normalized units. Every reported bar is larger than the actual error, and the table lists the margin for each row.

room	K	branch	reported γ_B	bar	$\lvert\hat{\gamma} - \gamma^{\text{exact}}\rvert$	margin
cube	28	DRIFT	0.572172845	8.7×10^{-6}	6.13×10^{-7}	14.2×
cube	50	RICH	0.627317093	3.3×10^{-6}	1.53×10^{-7}	21.7×
cube	110	DRIFT	0.727127618	6.2×10^{-6}	9.68×10^{-7}	6.4×
cube	204	DRIFT	0.787693771	2.5×10^{-6}	1.27×10^{-6}	2.0×
cube	403	RICH	0.834934168	2.0×10^{-5}	3.87×10^{-6}	5.1×
box (faithful)	25	DRIFT	0.110382839	3.9×10^{-6}	4.66×10^{-8}	84.8×
box (faithful)	50	DRIFT	0.145486973	2.0×10^{-6}	2.17×10^{-7}	9.2×
box (faithful)	110	DRIFT	0.187035568	1.3×10^{-5}	1.06×10^{-6}	12.2×
box (faithful)	203	DRIFT	0.229223233	1.9×10^{-5}	2.89×10^{-6}	6.7×
box (fixed-rel.)	25	DRIFT	0.529212237	2.8×10^{-6}	9.32×10^{-8}	29.6×
box (fixed-rel.)	50	DRIFT	0.634224209	2.0×10^{-6}	1.37×10^{-8}	145.6×
box (fixed-rel.)	110	DRIFT	0.720564111	4.7×10^{-6}	1.80×10^{-7}	25.8×
box (fixed-rel.)	203	DRIFT	0.781414859	1.3×10^{-5}	6.36×10^{-7}	19.8×

Table 4: Wherever the exact value is known, the reported interval contains it. On all three closed-form rooms every reported interval contains the exact γ_B, with room to spare. The same mesh rule produces every other reported γ_B, so these checks back the sampled-room values of Section 6. The intervals validate the rule on these rooms and are not rigorous error bounds for other rooms.

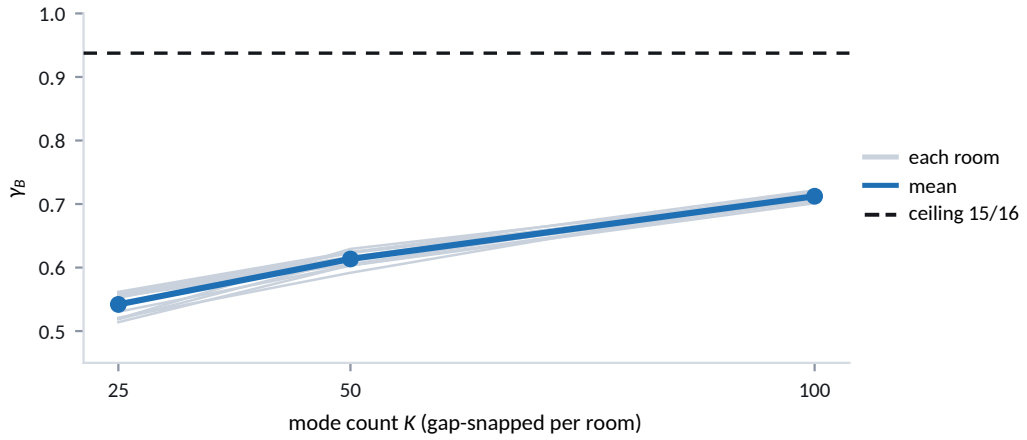

Figure 8: Covered rooms behave alike, which a uniform guarantee needs. Across the sampled rooms γ_B stays well below the ceiling of Theorem 2, and the spread between rooms shrinks as K grows. A value closer to the ceiling would mean a larger worst-case rate constant, so rooms of this class share nearly the same worst case.

C.2 Numerical checks and coverage test

The pipeline applies four checks before it reports a numerical result. They are mass-matrix orthonormality of the computed eigenvectors, agreement of the modal count across adjacent refinement levels, an independent eigenvalue recount below each spectral cutoff, and invariance under a rigid rotation of the coordinate frame. The reported solves satisfy the orthonormality and rotation checks to numerical precision and agree on the modal counts.

Coverage is evaluated by clipping the translated array footprint against the polyhedral room. No quadrature enters this test. Floating-point geometry is the remaining numerical limitation. On the analytic sharpness family of Appendix D, the computed uncovered area agrees with the closed form to floating-point precision on the nonzero cases.

seed	K	γ_B	d_{eff}/K	K	γ_B	d_{eff}/K	K	γ_B	d_{eff}/K
		target 25			target 50			target 100	
4001	25	0.5617	1.54	50	0.6247	1.61	101	0.7060	1.68
4007	27	0.5523	1.51	51	0.6235	1.58	104	0.7216	1.64
4008	26	0.5186	1.49	52	0.6130	1.53	100	0.7044	1.59
4015	25	0.5552	1.52	50	0.6131	1.58	102	0.7189	1.65
4016	26	0.5577	1.53	50	0.6149	1.60	101	0.7190	1.67
4023	25	0.5585	1.54	51	0.6228	1.58	103	0.7075	1.65
4025	25	0.5140	1.47	50	0.6036	1.52	103	0.7058	1.59
4026	25	0.5606	1.52	50	0.6167	1.59	103	0.7157	1.66
4029	25	0.5209	1.48	50	0.5918	1.52	108	0.7163	1.59
4033	25	0.5304	1.50	50	0.6032	1.53	104	0.7013	1.61
4035	25	0.5195	1.55	50	0.6293	1.63	100	0.7134	1.70
4036	28	0.5544	1.52	50	0.6052	1.57	104	0.7171	1.64
mean		0.5420	1.51		0.6135	1.57		0.7123	1.64
spread (range/mean)		8.8%	5.3%		6.1%	7.5%		2.9%	7.1%

Table 5: Every covered room behaves alike at every tested mode count. Both γ_B and the effective dimension per mode vary little from room to room, and the effective dimension per mode stays below the worst-case factor $1/(1-\gamma_B)$, increasingly so as K grows. The per-mode factor d_{eff}/K in the constant of Theorem 1 is therefore nearly the same across this room class.

room	K	γ_B	actual d_{eff}/K	worst case $1/(1-\gamma_B)$
cube	403	0.835	1.89	6.06
box (faithful)	268	0.248	1.08	1.33
box (fixed-rel.)	268	0.807	1.76	5.17

Table 6: The worst-case bound overstates the rate constant. On the exact-spectrum rooms the actual effective dimension per mode lies below the worst-case factor $1/(1-\gamma_B)$, by a wide margin where γ_B is largest. This is why Theorem 1 carries d_{eff}, and why $K/(1-\gamma_B)$ appears only as the interpretable bound of Corollary 1.

C.3 Ensemble values and effective dimension

Figure 8 and Table 5 report the covered-room ensemble. The finite-K values remain below the geometric ceiling, and the per-mode effective dimension is substantially smaller than the worst-case factor $1/(1-\gamma_B)$.

Table 7 shows why a statement uniform in K needs the geometric test. The uncovered region of room 13 is about 0.06% of its volume. An uncovered room looks ordinary at every finite cutoff tested.

C.4 Element order and resolution limit

Uniform refinement makes the P_1 degree-of-freedom count one level finer equal to the P_2 count one level coarser. Table 8 therefore compares the two element orders at matched algebraic size.

Define $K^*(h)$ as the first modal cutoff at which the eigenvalue error exceeds half the adjacent spectral gap. Beyond this point the ordering of near-degenerate modes is not numerically reliable. A swap of two modes inside the first K only rotates the basis, which changes B by a similarity and leaves γ_B and d_{eff} unchanged, so only a swap across the cutoff matters. The limit is measured on the rectangular box alone, where the exact spectrum gives the eigenvalue error. On sampled rooms the snap to a spectral gap and the agreement of the modal count across refinement levels guard the cutoff instead. On the rectangular-box ladder, the observed limits are associated with near-degenerate mode pairs. A simple extrapolation from the two finest observed levels predicts the next limit, and Figure 9 shows that prediction as a diagnostic, not a guarantee.

The tables report the actual gap-snapped modal dimension used in each solve.

target	accepted-ensemble range	K	room-13 γ_B	inside range
25	[0.5140, 0.5617]	25	0.5160	yes
50	[0.5918, 0.6293]	50	0.6084	yes
100	[0.7013, 0.7216]	103	0.7056	yes

Table 7: Finite-K values cannot tell an uncovered room from a covered one. Room 13 has a small uncovered region, yet its γ_B lies inside the covered-room range at every tested mode count. By Proposition 2 its γ_B still approaches one as K grows, so only the geometric coverage test separates it, which is why Theorem 2 states a condition on the room and not a threshold on γ_B.

K	$\gamma_B^{P_1}$	e^{P_1}	$\gamma_B^{P_2}$	e^{P_2}	$\lvert e^{P_1}/e^{P_2}\rvert$
28	0.572208376	3.61×10^{-5}	0.572172845	6.13×10^{-7}	58.9×
50	0.627485287	1.68×10^{-4}	0.627320422	3.48×10^{-6}	48.3×
110	0.727024005	-1.05×10^{-4}	0.727127618	-9.68×10^{-7}	108.0×
204	0.787628877	-6.62×10^{-5}	0.787693771	-1.27×10^{-6}	52.3×
403	0.835393147	4.63×10^{-4}	0.834953959	2.37×10^{-5}	19.6×

Table 8: Quadratic elements are more accurate than linear elements at equal cost. At the same number of unknowns on the cube, their error in γ_B is at least an order of magnitude smaller at every tested cutoff. This is why the pipeline uses quadratic elements for every reported value.

D COVERAGE STRESS TESTS

The sampled-room study screens rooms for geometry and mesh quality before the FEM solve. Rooms that fail the polyhedral coverage test stay in the analysis as negative examples for the geometric condition, and a constructed family tests the sharp length threshold.

D.1 CONSTRUCTED THRESHOLD FAMILY

Consider a polygon with one axis-aligned facet of length ℓ, flanked by edges whose outward normals are $(\pm\cos\theta, -\sin\theta)$. For an $n \times n$ array of spacing s, the uncovered area below the threshold is

$$A_{\mathrm{unc}}(\ell, \theta) = \frac{(n-1)(ns-\ell)^2}{4\tan\theta}, \qquad (n-1)s \le \ell < ns, \tag{26}$$

and it is zero for $\ell \ge ns$. Below $(n-1)s$ the array cannot reach the facet at all, a whole strip stays uncovered, and the formula no longer applies. In a prism the uncovered volume fraction is this area divided by the area of the polygon, by Lemma 1. Thus the transition occurs at the aperture plus one grid spacing, not at the aperture itself.

Equation (26) also shows that the uncovered area vanishes quadratically as $\ell \uparrow ns$. Changing the flank angle changes the prefactor but not the threshold. This isolates the mechanism used in Theorem 2.

D.2 SAMPLED-ROOM AUDIT

The sampled rooms exhibit two failure modes predicted by the geometry. Most uncovered rooms contain a corner whose incident outward normals do not lie in one closed coordinate quadrant. The remaining uncovered cases fail the axis-facet length condition. Conversely, the rooms that satisfy both conditions are covered by the clipping test.

Table 9 accounts for every room on record. The exact clipping certificate classifies every room that passes the geometric screens as covered or uncovered. The remaining rooms fail a geometric screen on inradius, aspect ratio, placement-domain volume or shortest polygon edge before any FEM solve, so they are never certified. Sampling under quadrant compatibility at every vertex leaves no uncovered room.

Across the rooms for which a coverage verdict is available, the theorem's two conditions agree with the clipping verdict. The sampled covered rooms do not approach the ns length threshold closely, so

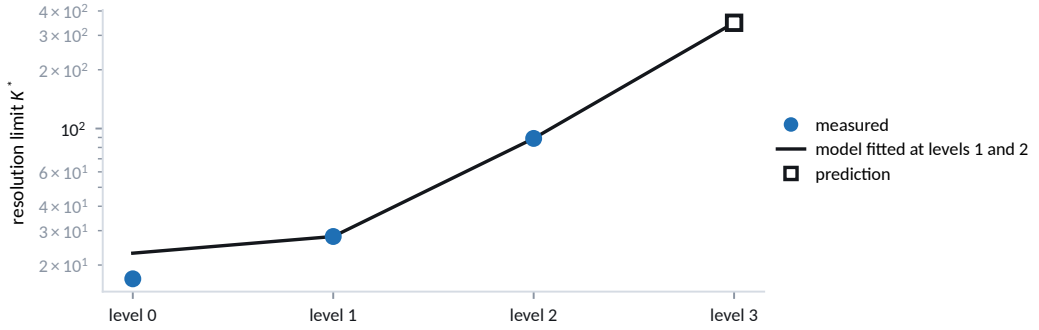

Figure 9: Each mesh resolves modes only up to a limit K^*, beyond which close modes swap order, and each refinement raises that limit. The model is fitted at the two finest measured levels, so it matches them by construction, and the coarsest level, the only independent check, falls below it. The square extrapolates one level further and is a prediction, not a measurement. The limit sets which mode counts the reported values can use.

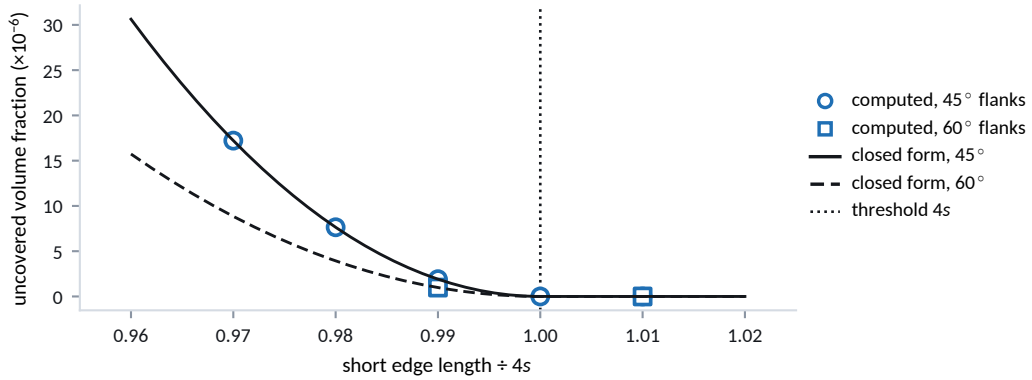

Figure 10: Coverage switches on exactly when the short wall reaches ns, the aperture plus one grid spacing. Below that length a sliver of the room is uncovered, and it shrinks quadratically to zero at the threshold. The exact computation lands on the closed-form curve at every tested length and both flank angles. The length condition of Theorem 2 therefore cannot be shortened.

they do not test its sharpness. That role is played by the constructed family in Figure 10. Likewise, the absence of a length condition on oblique facets follows from the proof, not from the finite room sample.

D.3 Orientation

The certificate is stated in the array's frame, so a rotated array is a different geometry. Table 10 sweeps the rigid array through orientation in the rooms of this paper, with the exact oracle and exact arrangement-cell coverage and no finite elements. The coverage count drops to zero at the smallest nonzero angle tested in every configuration, so $\sup_K \gamma_B$ reaches one at every nonzero angle tested and not only at $45°$. The uncovered set grows from measure zero with the angle, and the finite-K values move far less, so the quantity that sets the rate at a fixed mode count survives a misalignment that destroys the ceiling.

E Synthetic experiment protocol

The synthetic experiments use the FEM eigenfunctions and the Gaussian observation model of Section 3. All design parameters are fixed before the reported runs, except the two knee statistics of

quadrant compatibility	spacings certified	rooms	covered	uncovered	not certified
no	one	36	12	19	5
no	two	20	7	13	0
yes	one	40	40	0	0
all rooms on record		96	59	32	5

Table 9: Every sampled room is accounted for, including every rejected one, so no failure is hidden. Each room either carries an exact coverage verdict or fails a geometric screen before the test runs. The uncovered rooms fail the conditions of Theorem 2, and sampling under quadrant compatibility at every vertex, the corner condition of Lemma 1, leaves none uncovered.

room and array	$\kappa_{\min}$ aligned	$\kappa_{\min}$ at 0.1°	uncovered at 45°	γ_B at 45°	d_{eff} ratio at 45°
cube s=0.042	1	0	2.117%	0.954	1.779
box fixed-relative	1	0	1.680%	0.955	1.713
box faithful 42mm	1	0	0.048%	0.386	1.040
measured room, board in plane	1	0	0.246%	0.550	1.099
measured room, board yaw	1	0	0.043%	0.433	1.003

Table 10: Misalignment removes the uniform-in-K ceiling at once and the finite-K constant only slowly. The coverage count falls to zero at the smallest nonzero angle tested in every configuration, so $\sup_K \gamma_B$ reaches one, which is Proposition 2. At 45°, far beyond any plausible mounting error, the uncovered fraction is still small, while d_{eff} at the largest tested mode count changes little for the faithful-scale box and the measured room and inflates substantially for the two relative-scaled rooms. Exact oracle and exact clipping throughout, with no finite elements.

Figure 12, which were defined after the run and are declared as post-hoc. Static and moving arrays use the same microphones, total observation time, noise model, and posterior estimator, and only the array path changes. Table 11 lists the declared values.

E.1 Numerical setup

FEM eigenfunctions are evaluated at interior points with the local quadratic basis. The evaluator is checked on a rectangular room with a closed-form spectrum and by cross-level consistency on the production mesh.

The static baseline is selected by minimizing exact Bayes risk over uniformly sampled positions in the placement domain, followed by local refinement around the best candidate. This makes the static baseline a best-found position rather than an arbitrary one.

The stochastic trajectory follows

$$X_{t+1} = \mathrm{refl}(X_t + s_v \xi_t), \qquad \xi_t \sim \mathcal{N}(0, I_3), \tag{27}$$

starting from a uniform draw on D. With Brownian time step $\Delta = s_v^2$, Equation (27) is an Euler approximation to the reflected Brownian motion used in Appendix A. Nothing here assumes that the discrete reflected walk has an exactly uniform invariant law. The theoretical constant is therefore a reference prediction for the limiting process, not an exact finite-step theorem for the simulator.

E.2 Rate experiment

The asymptotic constant

$$\frac{\sigma^2 |D|}{m} d_{\text{eff}}$$

is computed from the noise variance, the placement-domain volume, and the spectrum of B, and no parameter is fit to the risk curve. Figure 12 shows the normalized risk over the full simulated range. The observed mean approaches the predicted constant from above. The two knee statistics were defined after the run. The ratio-two knee is the first checkpoint at which $T\times$Bayes risk is at

declared input	value
room	ensemble room 4001, volume normalized to one, with one unit taken as 5.92 m wherever a speed or a position error is physical
modes	$K = 50$, snapped to a spectral gap, shortest modal wavelength 3.1 m
array	4×4 grid, $m = 16$, spacing $s = 0.042$ in normalized units, leveled, fixed orientation
placement domain	$\lvert D \rvert = 0.791$ in normalized units
prior	$c \sim \mathcal{N}(0, \tau^2 I_K)$ with $\tau^2 = 1$
noise	per-sample SNR 20 dB at rest, so $\sigma_0^2 = \tau^2 K / 10^{\mathrm{SNR}/10} = 0.5$
reflected walk	$s_v = 0.05$ per coordinate in normalized units, $\Delta = s_v^2 = 0.0025$, 8 trajectories from seed 4001
observation times	$T = 2000$ snapshots for the field and coefficient figures, up to $T = 50000$ for the rate
static baseline	512 seeded uniform positions, then local refinement at step 0.01
speed model	$\sigma^2(v) = \sigma_0^2[1 + (v/v_0)^p]$ with $v_0 = 0.5$ m/s and $p = 2$, 50 snapshots per second

Table 11: Every synthetic experiment runs on one declared setup, fixed before the runs, so no value was tuned to a result. The room is solved in normalized units, and a physical length enters only where a speed or a position error is stated. These are inputs, not measurements, and the constant of Theorem 1 follows from them and the spectrum of B.

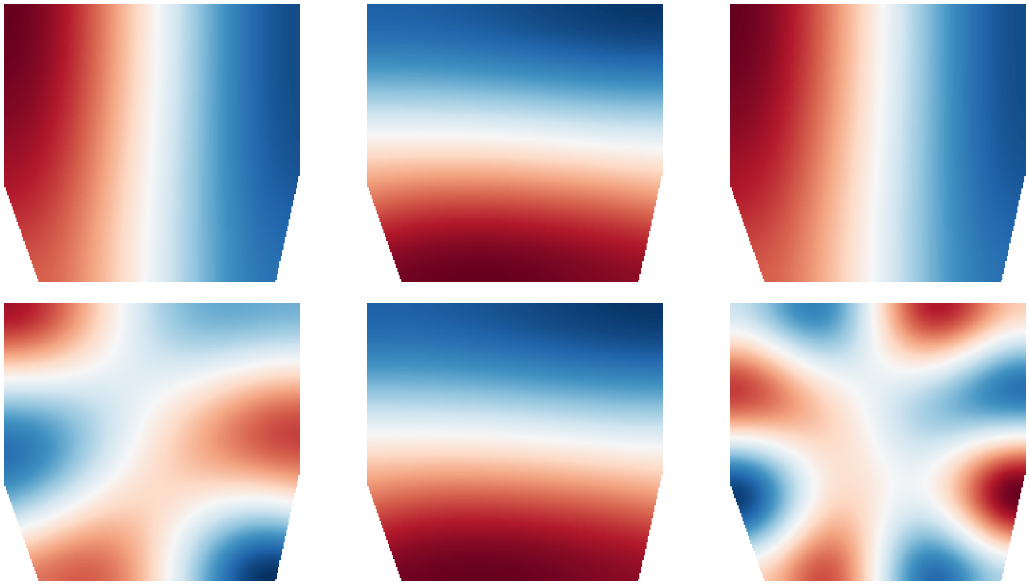

Figure 11: The sound field in the model is a sum of room modes like these, K of them in total. Higher-index modes vary faster in space. A static array must separate all of them from a fixed set of microphones, which Proposition 1 shows is impossible once K exceeds the number of microphones.

most twice the predicted constant. The $1/T$ scaling holds once $T\times$excess stays within a declared band of its final-decade mean. Over the final range, the residual is approximately consistent with a $1/T$ correction, but no second-order rate is claimed.

The moving design falls below the static rank floor at the same microphone-seconds. The static curve uses the lower envelope over candidate positions at each observation time. The comparison therefore isolates the acquisition difference within the assumed Gaussian model.

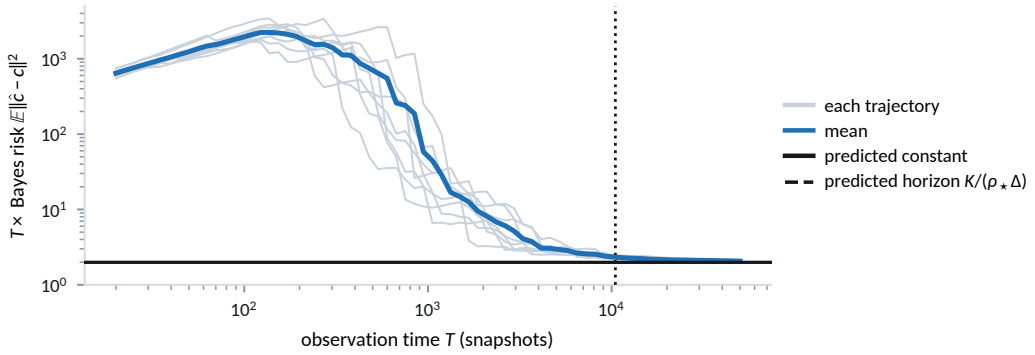

Figure 12: The risk approaches its predicted constant from above after a concentration horizon, consistent with Theorem 1. The mean over trajectories flattens onto the constant near the horizon $K/(\rho_\star \Delta)$ drawn with constant one. The theorem fixes that horizon only up to constants and logarithmic factors, and the knee statistics were defined after the run, so this is a scale check and not a test.

trajectory	risk at v_0	path covered at v_0	risk at own optimal speed
reflected Brownian motion (idealized reference, declared step)	2.18×10^{-4}	–	– (no speed)
Lissajous 3:2:1	6.91×10^{-3}	100%	3.60×10^{-3} at 0.1 m/s
3D raster	7.11×10^{-3}	68%	6.19×10^{-4} at 1 m/s

Table 12: Speed is a design variable. At the reference speed the raster covers only part of its path in the observation time, and at its best speed it closes most of the gap to the idealized reflected-Brownian reference. The Lissajous path covers its whole length at every speed, so it does best at the slowest speed. These paths lie outside Assumption 1, so exact Bayes risk ranks them instead of the constants of the theorem.

E.3 Trajectory and speed sensitivity

The comparison covers three-dimensional raster, Lissajous and reflected-Brownian paths. The deterministic paths are clipped to the placement domain. Speed affects the deterministic paths through the per-snapshot displacement and the noise model

$$\sigma^2(v) = \sigma_0^2 \left[1 + (v/v_0)^p\right]. \tag{28}$$

The reflected-Brownian trajectory is an idealized mixing reference and is not assigned a physical traversal speed.

The raster has an interior optimum on the tested speed grid (Table 12). The Lissajous path is fully covered at all tested speeds, so its best value occurs at the slowest tested speed. These results are specific to the stated noise model and observation time. The deterministic paths do not satisfy Assumption 1, so exact Bayes risk ranks them, rather than mixing constants.

E.4 Trajectory classes at equal speed

The reflected-Brownian reference of Table 12 is not speed-matched, since its declared step moves the array much farther per snapshot than any deterministic path on the declared speed grid. Table 13 therefore repeats the comparison at equal displacement per snapshot, matching the reflected-Brownian path to each deterministic path by root-mean-square step, and reports the plateau of $T\times$risk against the constant of Theorem 1 at that speed's declared noise level. The declared-step reference attains the constant to within a few percent at every speed, which is the check of Figure 5 repeated across the noise model. The raster, a path anyone can drive, comes within a factor of one and a half of the constant once it covers its length, although it lies outside the trajectory contract. The Lissajous curve stays more than an order of magnitude above it at every speed, because a single closed curve covers a three-dimensional placement domain poorly. The speed-matched reflected-Brownian path is far worse than either, because a diffusive walk at walking speed explores only a

path	at $v_0 = 0.5$ m/s	at 1 m/s	path covered at v_0, $T = 50000$
reflected Brownian, declared step	1.03	1.03	–
3D raster	1.32	1.26	100%
Lissajous 3:2:1	34.97	34.71	100%
reflected Brownian, speed-matched	2.8×10^5	6.7×10^4	–

Table 13: At equal speed a practical raster path comes close to the constant of Theorem 1, and random walking does not. Each entry is $T\times$risk at the largest simulated T divided by the predicted constant at that speed, so one means the constant is attained. The declared-step reference attains it, the raster comes within a factor of one and a half once its path is covered, and a reflected-Brownian path slowed to the same displacement per snapshot explores too little of the room to approach it.

	$T = 2000$			$T = 20000$		
error (mm)	moving, jitter	moving, offset	static	moving, jitter	moving, offset	static
0	0.00339	0.00339	39.1	0.000105	0.000105	37.5
1	0.0034	0.00351	39.1	0.000106	0.000146	37.5
2	0.00343	0.00435	39.1	0.000107	0.000501	37.5
5	0.00364	0.00936	39.1	0.000113	0.00318	37.5
10	0.00424	0.017	39.1	0.000139	0.0166	37.6
20	0.00754	0.0943	39.1	0.000276	0.0611	37.6
50	0.025	0.649	39.3	0.00165	0.341	37.7
100	0.0888	1.5	39.3	0.0139	1.4	38.2

Table 14: Mis-known positions cost the moving array far less than the rank obstruction costs a static one. Entries are the expected squared error when the estimator's assumed positions are wrong by the stated Gaussian error, at two observation times, against the strongest static placement with one survey error. Independent per-snapshot jitter leaves only a fourth-order floor, visible at the largest errors, while a survey offset leaves a floor quadratic in the error, yet even the largest offset tested stays more than an order of magnitude below the static rank floor of Proposition 1.

small part of the domain within the observation time. The ranking of Table 12 therefore reflects the displacement per snapshot of its reference, not an advantage of random paths over planned ones.

E.5 Position error

The theory assumes the estimator knows where every snapshot was taken. Table 14 breaks only that knowledge, keeping the declared model of this appendix otherwise unchanged, and reports the exact expected squared error when the posterior mean uses the assumed positions while the measurements come from positions in error by a Gaussian offset of the stated size, mirrored back at the wall when they would leave the room. Two regimes bracket a real survey. Jitter gives each snapshot its own independent offset, as when every surveyed position carries its own error, and a survey offset shifts the whole path by one draw. The comparator is the strongest static placement of Figure 4, placed once and therefore carrying one survey offset. Three predictions were written before the run. The first two hold: the jittered moving array stays below the static placement at every error in the grid, by orders of magnitude, and its excess risk grows quadratically in the error at small error. The third predicted a visible floor under the jittered risk at the declared error, and it failed as declared. To first order independent jitter averages out, so at small error the risk keeps falling with observation time. Jitter still leaves a floor, because averaging a mode over a Gaussian offset of size ϵ per axis attenuates it by $e^{-\epsilon^2\lambda_k/2}$, a bias in the assumed model that no observation time removes. That floor is of fourth order in the error, which a post-hoc computation of the limit confirms, so it is negligible at the smaller errors of Table 14 and most of the risk at the largest error and the longer observation time. The survey offset leaves a larger floor, quadratic in the error and independent of observation time. Even the largest offset tested, a few percent of the shortest modal wavelength of Table 11, leaves that floor far below the static placement, because the static floor is set by rank and not by noise. The run does not test the noise of a continuously moving array, which Section 8 lists as untested.

n	99%-energy rank	top-3 energy fraction	rank at 10^{-2}	rank at 10^{-3}
1	3 / 3	0.999 / 0.998	5 / 5	9 / 10
2	5 / 5	0.957 / 0.971	10 / 8	16 / 15
4	7 / 7	0.872 / 0.902	13 / 12	24 / 22
8	8 / 7	0.879 / 0.897	15 / 13	31 / 28
16	7 / 6	0.905 / 0.929	14 / 12	36 / 30

Table 15: Real impulse responses show the rank obstruction. Each entry gives the two loudspeakers in turn. Stacking rail positions raises the number of directions above the tightest threshold at every step, as acquiring new spatial directions predicts, while at the looser threshold and for the energy rank the count rises through eight positions and then falls slightly at sixteen. This is consistent with the static ceiling of Proposition 1 and with added positions lifting it, as a measurement diagnostic.

E.6 Finite-K approach to the coverage ceiling

Proposition 2 fixes $\sup_K \gamma_B$, but it does not give a rate at which finite-K values approach that supremum. Table 4 shows that approach on the exact-spectrum rooms. On the cube and the fixed-relative box, γ_B is still well below the ceiling $1 - 1/m$ at the largest validated K, and the paper claims no convergence law for that approach.

All moving trajectories vary in three dimensions. A trajectory restricted to one horizontal plane samples only horizontal translations of the planar array and need not identify the vertical modal variation targeted by the experiment.

F Public-data and estimator comparisons

These comparisons test whether the acquisition effect survives changes in hardware and estimator. They do not extend the theorem beyond its stated model.

F.1 Public moving-array data

The trajectoRIR dataset (Damiano et al., 2026) contains impulse responses from a 16-microphone circular array at multiple rail positions. The analysis forms the microphone-by-frequency data matrix over the stated band and reports three diagnostics from its singular values, the 99%-energy rank, the share of energy in the first three directions, and the numerical rank at fixed relative thresholds (Table 15). The former measures energy concentration. The latter depends on the threshold and is not a theorem-level rank statement.

At one rail position, the 99%-energy rank is small for both loudspeakers (Table 15). This is consistent with strong aperture-induced concentration, but Proposition 1 itself gives only the hard upper bound m for a static array. Stacking distinct positions increases the numerical rank through eight positions, as expected when acquisition samples additional spatial directions, with a slight fall at sixteen. Because the dataset contains no modal ground truth for this analysis, these quantities serve only as measurement diagnostics.

F.2 Third-party moving and static estimators

The comparison runs our posterior mean against two published moving-microphone estimators from a public library, using identical synthetic fields, microphone positions, observation times, and noise, so that only the estimator varies. The first is the kernel estimator of Brunnström et al. (2026), which has a static counterpart, and both run at their published settings. The second is the spherical-harmonic estimator of Katzberg et al. (2021), which has no static counterpart in the library. A moving-versus-static contrast is therefore not defined for it, and it is not evidence either way on that point. The common score is energy-weighted relative RMS field error over the scored band, pooled over bins and grid points, with the per-bin curves in Figure 13.

The design was declared in full before the code (Table 16). Two interpretation clauses were fixed in advance. A nonzero moving-side floor of a baseline reflects mismatch between its kernel class and

declared input	value
observation times per array	1, 2, 3, 5, 7 and 10 s at 256 Hz, of which Table 2 reports three
excitation	0.5 s periodic
damping	uniform modal, $T_{60} = 0.5$ s
trajectory	the reflected-Brownian path of Figure 3
static position	the best static position of Figure 3
noise	the same level for both arrays
scored band	2–110 Hz, with 112–128 Hz reported separately
baseline estimators	published settings, with the regularizer set to the known noise variance on both sides
truth per bin	$c_k(\omega) = \varphi_k(x_s)/(\omega_k^2 - \omega^2 + 2i\delta\omega)$, $k \le K$

Table 16: The comparison design, declared before the code, so no parameter was tuned to the results. These are inputs, not measurements.

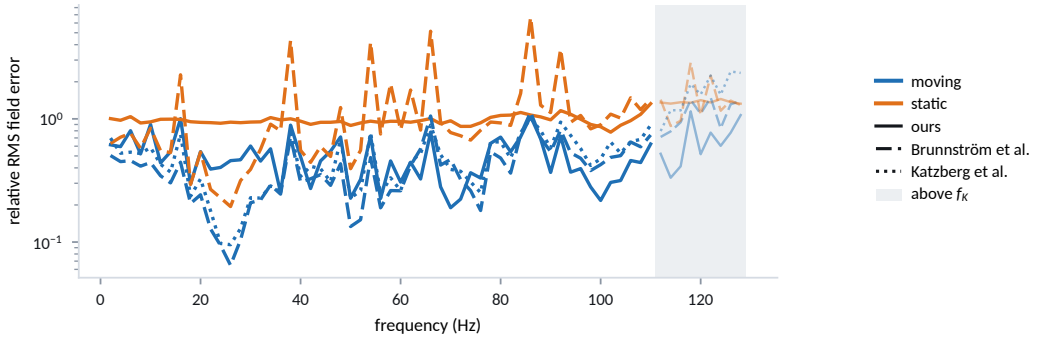

Figure 13: The moving advantage holds across frequency, not only in the pooled score. At the longest observation time the moving array has lower error than the static one in nearly every in-model bin, under our estimator and under the estimator of Brunnström et al. (2026). The gain is therefore a property of the acquisition across nearly all in-model frequencies. Above the K-mode validity limit, shaded, the model no longer applies and no claim is made.

the modal truth, and is reported as such. No claim is made that a baseline's static line sits on the floor of Proposition 1. The only claim the comparison was declared to support is that moving beats static at equal microphone-seconds for every estimator that has a static counterpart.

Katzberg et al. (2021)'s estimator does not execute on the pinned software stack, because it calls a special function removed from the current scientific library. It runs here under a shim that maps the removed function to its replacement, checked against three closed-form values, and every cell of that method raised a library deprecation warning. Its expansion order depends on the choice of origin, so its output is origin-dependent by construction.

The synthetic truth is generated from the same K-mode field used elsewhere in the paper. For each frequency bin, the modal coefficients follow the damped modal Green's function. Bins beyond the K-mode validity range are reported separately rather than used to support the in-model comparison.

The regularizer is the known noise variance on both sides, and under that rule the moving array wins at every observation time for both estimators that have a static counterpart. The rule is not neutral between the two sides. Swept over twelve decades, it is the exact optimum for both arms of our estimator and for the moving arm of Brunnström et al. (2026), while their static arm needs about a hundred times more before it becomes usable at all, where it reaches 0.77, 0.72 and 0.67 at the three observation times. With each side at its own optimum the moving array leads at three and ten seconds and trails at one, where it has visited a tenth of the plane it is scored on. Theorem 1 fixes the rate only after a concentration horizon, so no advantage is claimed before it. A rank-deficient static Gram matrix therefore has to be damped heavily before it produces anything, which is a property of what the acquisition made available rather than of the estimator. Both third-party moving estimators are more accurate than ours at every observation time (Table 2). Replacing our

method	array aperture	room plane
learned physics prior, static	0.039	0.999
posterior mean, static	0.015	0.799
posterior mean, moving	0.174	0.426

Table 17: In this comparison a learned prior does not recover what a static array never measured. Inside the array footprint both static methods are accurate, but across the room plane their error rises toward that of predicting silence. The moving posterior mean gives up accuracy inside the footprint for the lowest error across the room. The learned prior does not beat the static posterior mean in either region, and in this experiment only the moving acquisition brings room-plane error well below the level of predicting silence.

single global prior variance by the true per-coefficient variances reduces our error severalfold, so prior mismatch accounts for much of the difference. This oracle-prior calculation is a post-hoc diagnostic, not a deployable estimator.

F.3 Published physics-informed method on the static arm

The last comparison evaluates the public physics-informed reconstruction method of Karakonstantis et al. (2024) on static impulse responses. Inputs are band-limited and resampled to the method's required rate. It scores that method and our static and moving posterior means on two regions, the array aperture and the full room plane at the same height (Table 17).

The learned static method performs well within the microphone hull. On the full-room plane its error is much larger, while the moving posterior is lowest among the three methods in this experiment. The two regions answer different questions, so the aperture result does not contradict the full-room acquisition claim. We do not run the published network directly on the three-dimensional moving samples because that would require changing its input architecture.

G Real-room measurement

The hardware study uses one room and a 4×4 planar array. It tests the rank obstruction and the reconstruction trend on measured data, and it fits none of the asymptotic constants in Theorem 1.

G.1 Apparatus and protocol

Table 18 lists the room, the array, the source and the excitation. The temperature and the scored lines were measured in each session, and every other value was declared before recording. A room in everyday use cannot be fully specified, because its furniture and fittings change. The rank crossing needs no model of the room. The predicted counts and the reconstruction use a rigid-box model, which leaves the furnishings out.

A declared master ordering of positions fixes the experiment. A rung is a prefix of that ordering, so each rung contains the one before it, and every rung has the same total observation time (Table 18). The second session recorded the same ordering in reverse, so no rung inherits the first session's traversal order. Between the sessions the array was disconnected, and every position was surveyed again before the second session.

The probe positions lie in distinct regions of the room, are held out of every fit and never belong to any rung. The reconstruction error is scored at those probes as the per-line relative RMS error on magnitudes, which does not depend on the source's per-line phase and scores a prediction of zero at one. The posterior mean uses the rigid-box modal basis of the surveyed room (Table 18). Its declared prior takes the modal prior variance from an anchor position recorded at the start of every recording, and the noise variance from the silence recording of each session.

item	value
room	$3.487 \times 7.026 \times 2.469$ m, or 60.5 m^3, furnished, in everyday use and not modified for the experiment, with the door closed and the air conditioning off
reverberation time	1.31 s from 60 to 90 Hz, a resolution-limited value, and 1.52 s from 90 to 140 Hz
array	miniDSP UMA16v2, 4×4 planar grid, $m = 16$ microphones facing the source, in a vertical plane with its normal along the long axis, centre 0.885 m above the floor at every position
microphone spacing	42 mm, so the aperture is 126 mm
source	Genelec 8010A, woofer centre 0.17 m above the floor near a corner at one end of the long axis, firing along it. It was knocked once during an abandoned recording in session 1 and put back at its original position without a new measurement
excitation	flat multisine on a 2 Hz grid from 60 to 140 Hz, Schroeder phases, sampled at 48 kHz
scored lines	35 of the 41 grid lines. The four detection lines, at 74, 102, 110 and 134 Hz, and the 120 Hz line carry no excitation, and the 60 Hz line fails the mains check on the silence recording of each session
take	20 s per position, 40 periods of 0.5 s, of which the last 38 are analyzed
temperature	29.2 to 30.0 °C, read before every recording
held-out probes	4 positions, every probe microphone at least 1 m from every microphone of every rung, ranked set and static position
rungs	prefixes of one declared ordering with 1, 2, 4, 8 and 16 positions, each with 32 periods in total, taken as the last $32/k$ analyzed periods at each of its k positions
modal basis	rigid box, the 63 modes below 180 Hz at 26 °C, which is 40 Hz above the excitation band, with each mode frequency scaled to the temperature read before its recording

Table 18: The hardware setup, with every value declared before recording except the measured temperature and scored lines. At the top of the band the array spans a small fraction of a wavelength, so from one position its microphones see nearly the same pressure and resolve few independent directions. That small aperture is where a static array resolves the fewest directions and motion has the most to add.

G.2 Rank crossing

A static array with m microphones produces a data matrix with at most m rows, so its rank cannot exceed m. The multi-position measurement crosses this bound without invoking a modal basis, a channel map or a reconstruction algorithm. The noise matrix has the shape and the averaging of the data matrix. The silence recording of the session is cut into as many consecutive blocks as the rung has positions, each of the per-position length, and each block is averaged coherently and stacked like a position. Its norm stands in for the noise carried by the data, which no recording observes directly. At 16 surveyed positions the 17th singular value exceeds the session noise norm in every run, by a factor of about six on the raw data matrix and of at least 1.7 under the column-normalised variant. The column-normalised variant is the declared instrument, the raw matrix is the direct measurement, and the crossing holds under both. Weyl's inequality therefore implies that the noiseless stacked measurement has rank at least 17, above the maximum rank of any static 16-microphone measurement.

The count in Table 1 of the main paper repeats in both sessions. The count itself uses the declared screening threshold. The rank-crossing claim above does not, because it uses the measured noise norm directly.

G.3 Reconstruction error falls with the number of positions

The pre-registered test uses the declared prior and predicts lower pooled held-out error at 16 positions than at one. The error falls from one position to 16 in both passes of session 1 and in session 2, well outside the repeat bands, and its rank correlation with the number of positions has the declared

relative threshold	0.005	0.0075	0.01	0.015	0.02
count at 16 positions	19	18	17	16	15

Table 19: The screened count moves with its threshold. The rank-crossing claim therefore rests on the singular-value test, which uses no threshold, and Weyl's inequality turns that test into a rank bound.

negative sign in every run (Figure 6). Most of the drop occurs at the largest rung. Between four and eight positions the error rises slightly in two runs, a single step that the trend test does not assess.

Under the declared prior the error at one position is worse than predicting silence. The second prior arm, declared with the analysis code, scales the declared prior variance by 10^{-3}. That factor was fixed on a pilot recording before either session and was never re-selected on session data. Its pooled error stays near that of silence at one position and falls with every added rung. In every session and under both priors, 16 positions end below the static comparator, the candidate position with the lowest predicted error under the rigid-box model, chosen before recording, although it records for longer than any rung.

G.4 Protocol and sensitivity

The design, the schedule, the gates and the analysis code were fixed before the first recording. No recording was seen before the analysis code was fixed. The second session repeated the same excitation set and used the same analysis. Timing and surveyed positions passed the predefined validity checks.

The screened direction count admits a singular direction when its singular value exceeds both a relative threshold and a noise-based threshold. At the rungs reported in the main paper, the relative threshold is the active one. Table 19 therefore reports the sensitivity of the 16-position count to that threshold.

The measured noise level lies above the range the count prediction assumed, so that prediction is approximate. The measured count at 16 positions exceeds it nonetheless, the direction of a room carrying more spatial structure than the rigid-box model. The noise level does not affect the algebraic static rank bound, and the direct singular-value comparison still places the noiseless stacked rank above 16.

A second pre-registered hypothesis ranks five sets of positions by predicted error, and the second session recorded all five sets. The test counts the resolvable pairs of sets whose measured order matches the predicted order. Under the declared prior every resolvable pair but one matches, and under the second prior arm every pair matches. Under the second prior arm the measured set errors lie close to their predictions, whereas under the declared prior they are at least several times larger.