Research Square

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

Robust Fisher-Information Echo Scheduling for Two-Component CPMG NMR Relaxometry Under Uncertain T2 Values: A Computational Study

Connor Nitchals

nobleconnor257@gmail.com

Independent Scholar

Research Article

Keywords: NMR relaxometry, CPMG, T2, biexponential decay, Fisher information, D-optimal design, optimal sampling, inverse problem, low-field NMR

Posted Date: September 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10898895/v1

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

Additional Declarations: No competing interests reported.

Robust Fisher-Information Echo Scheduling for Two-Component CPMG NMR Relaxometry Under Uncertain T2 Values: A Computational Study

Connor Nitchals

Independent Scholar, Tonganoxie, Kansas, United States

Corresponding author: Connor Nitchals

Email: nobleconnor257@gmail.com

Article type: Original Paper

Target journal: Applied Magnetic Resonance

Acknowledgments: None stated in this draft.

Abstract

Multiexponential transverse-relaxation measurements are sensitive not only to signal-to-noise ratio but also to where the decay is sampled. This computational study develops a robust Fisher-information design for selecting 32 echo/sample times for two-component Carr-Purcell-Meiboom-Gill (CPMG) NMR relaxometry when the short and long T2 values and component fraction are not known in advance. A biexponential signal model was parameterized by amplitude, T2,short, T2,long, and short-component fraction. A 120-member prior ensemble spanning T2,short = 8-35 ms, T2,long approximately 55-300 ms, and fractions 0.15-0.85 was used to construct a greedy robust D-optimal schedule on a 0.4-400 ms candidate raster. The resulting schedule was compared with 32-point linear and logarithmic schedules using Fisher-information metrics and 150-repetition nonlinear Monte Carlo fits in three benchmark regimes at signal-to-noise ratios of 40 and 80. The robust schedule increased the prior-averaged log determinant of the information matrix by 2.41 relative to linear sampling, equivalent to an approximately 11-fold determinant ratio under the common parameter scaling. At SNR 80, robust sampling reduced T2,short relative RMSE from 10.1% to 7.1% in the well-separated case and from 31.9% to 19.4% in the low-short-fraction case. Performance remained difficult when the two relaxation times were close, revealing a fundamental identifiability limitation rather than an optimizer failure. The results provide a reproducible route to acquisition design for sparse or multi-train CPMG relaxometry.

Keywords: NMR relaxometry; CPMG; T2; biexponential decay; Fisher information; D-optimal design; optimal sampling; inverse problem; low-field NMR

1 Introduction

Transverse-relaxation measurements are among the most widely used time-domain nuclear magnetic resonance (NMR) experiments. The Carr-Purcell and Meiboom-Gill developments established echo-train methods that refocus static field inhomogeneity while allowing the irreversible transverse decay to be sampled over time [1,2]. In many materials, biological samples, porous systems, and complex mixtures, the measured decay is not adequately described by a single exponential. Multiple microscopic environments or exchange-limited populations can produce sums or distributions of exponential components, making the inverse problem substantially more difficult than monoexponential T2 estimation [4-11].

The difficulty is partly mathematical. Two exponentials with similar time constants can mimic one another over a finite time window, and a weak fast component can disappear into noise before enough information has been collected to identify it. Inverse Laplace approaches convert time-domain decays into distributions of relaxation times but require regularization because the inversion is ill posed [4,5]. Discrete multiexponential models avoid a predefined relaxation-time grid but introduce nonlinear parameter coupling and model-order questions [8-11]. Regardless of the estimator, the information available to the analysis is determined by the acquisition times at which the decay was sampled.

Sampling design is therefore a magnetic-resonance methodology question, not merely a post-processing choice. Jones, Hodgkinson, Barker, and Hore used Cramer-Rao theory to derive optimal sampling patterns for monoexponential T2

measurement and showed that optimal designs can be very different from uniformly spaced conventional sampling [3]. More recently, Parziale and co-workers studied the dependence of matrix-pencil relaxometry resolution on the uniform sampling interval and developed a procedure for choosing a time-domain spacing that improves separation of relaxation coefficients under noise [7]. Applied Magnetic Resonance has also published new approaches for summarizing relaxation and diffusion distributions and for deriving multiexponential relaxation behavior in model geometries [8,9].

A practical complication is that the relevant T2 values are usually not known before the measurement. A schedule optimized for one exact pair of relaxation times may be poor for another pair, which limits the usefulness of purely local design. In addition, a two-component decay contains more than two time constants: amplitude and mixture fraction also enter the fit and can correlate strongly with the relaxation parameters. A robust acquisition should therefore remain informative over a plausible ensemble of T2 values and fractions rather than at a single nominal point.

The present study formulates sparse two-component relaxometry as a robust D-optimal design problem. A prior ensemble of biexponential decays is generated over broad short- and long-T2 ranges. A greedy algorithm then selects 32 unique times from a candidate raster to maximize the average log determinant of the Fisher information matrix. The determinant criterion rewards designs that enlarge the joint information volume in the four-parameter space, while averaging over the prior ensemble reduces sensitivity to an uncertain starting guess.

Two conventional comparators are included. A linear schedule spreads 32 samples uniformly between 0.4 and 400 ms, which resembles the dense late-time coverage produced by a constant-spacing echo train when only a subset of echoes is retained. A logarithmic schedule allocates more points early in the decay and spans the full time range multiplicatively. The robust design is then assessed using both local information metrics across the prior ensemble and nonlinear Monte Carlo estimation in three deliberately different regimes.

The simulation is intended to represent the timing problem rather than the full pulse-sequence physics. Echo amplitudes are assumed to follow an ideal biexponential decay with independent Gaussian noise. B1 errors, stimulated echoes, diffusion in background gradients, imperfect refocusing, exchange during the train, and Rician magnitude statistics are not included. Those effects are important experimentally, but beginning with the ideal model allows the information benefit of timing alone to be quantified cleanly.

The main questions are: (i) where does a robust design place its samples when it must learn both fast and slow relaxation components; (ii) how much information is lost by uniform-linear or simple logarithmic sampling; (iii) do Fisher-information rankings predict actual nonlinear-fit accuracy; and (iv) how does a larger minimum accessible echo time reduce the achievable information? These questions are directly relevant to low-field NMR, benchtop relaxometry, sparse echo selection, and multi-train acquisition strategies where acquisition time or stored data volume is constrained.

2 Theory and Computational Methods

2.1 Biexponential relaxation model

The normalized signal model is a two-component transverse decay,

$$S(t) = A\,[\, f \exp(-t/T2,s) + (1 - f) \exp(-t/T2,l)\,] + \mathrm{epsilon}(t)$$

where A is the initial amplitude, T2,s is the short relaxation time, T2,l is the long relaxation time with T2,l > T2,s, f is the short-component fraction, and epsilon is independent zero-mean Gaussian noise. This discrete model is the simplest representation that still contains the central multiexponential identifiability problem. Similar sums underlie discrete relaxometry analyses and provide local approximations to broader relaxation-time distributions [4,7-11].

For design calculations, parameters were expressed in scaled coordinates q = (log A, log T2,s, log T2,l, logit f). This transformation makes the Fisher derivatives dimensionless, preserves positivity of amplitudes and relaxation times, and prevents the determinant criterion from being dominated by arbitrary units. The model sensitivities were evaluated analytically. The derivatives with respect to the two log-relaxation times are proportional to t/T2 multiplied by the corresponding component amplitude, which means that each component contributes most strongly over a characteristic time window rather than uniformly over the decay.

2.2 Robust prior ensemble

A robust design requires a set of plausible parameter values. One hundred twenty prior samples were generated with T2,s log-uniformly distributed from 8 to 35 ms. The ratio T2,l/T2,s was sampled log-uniformly from 2.5 to 12, and T2,l was clipped to approximately 55-300 ms while enforcing separation from T2,s. The short fraction f was uniformly distributed from 0.15 to 0.85. The amplitude was fixed at one for design because an overall amplitude scale does not alter the relative timing structure after noise normalization.

This ensemble is not a population prior for a particular material. It is a design prior that forces the schedule to remain informative across short-lived and long-lived components of varying relative abundance. The range was chosen to create both easy and challenging biexponential cases while remaining compatible with time-domain relaxometry windows commonly encountered in low-field and materials NMR [7-11].

2.3 Fisher information and robust D-optimal selection

For a schedule containing times t_k and Gaussian noise variance sigma^2, the scaled Fisher information matrix is F = sum_k j(t_k)^T j(t_k) / sigma^2, where j(t_k) is the row vector of model derivatives at time t_k. Because all schedules in the design comparison contain the same number of samples and are compared under the same noise variance, sigma only adds a common determinant offset and was omitted during schedule selection.

D-optimality maximizes log det F and therefore minimizes the volume of the local joint covariance ellipsoid under standard regularity assumptions. A small diagonal ridge of 10^-6 was added at the beginning of the greedy search so the determinant was defined before four independent directions had been selected. Candidate times were placed on a 500-

point linear raster from 0.4 to 400 ms. At each greedy step, the unselected candidate that produced the largest average log-determinant increment across the 120 prior parameter sets was added. The process continued until 32 unique times had been selected.

The matrix determinant lemma was used to evaluate each rank-one information update efficiently. This allows the design to account for all four parameters and all prior models without repeatedly recomputing a full determinant from scratch. The resulting robust design is an exact 32-point discrete schedule on the stated candidate raster, not a continuously optimized density.

2.4 Comparator schedules

The uniform-linear schedule used 32 equally spaced times from 0.4 to 400 ms. Because most of its points lie at late times on a linear scale, it is expected to estimate the slow component well but may undersample the rapid initial curvature generated by T2,s. The logarithmic schedule used 32 geometrically spaced times over the same interval. Logarithmic timing is a common intuitive response to multiscale decay because it allocates more observations at early times while retaining late-time coverage.

The robust schedule was not forced to resemble either comparator. In particular, it was allowed to cluster samples in time. Such clustering should be interpreted as repeated information allocation around highly informative regions. In an experiment, closely spaced selected times could be implemented by retaining selected echoes from one or more CPMG trains, by varying the echo spacing between acquisitions, or by allocating repeated averages near the same characteristic time. The present model does not prescribe one hardware implementation.

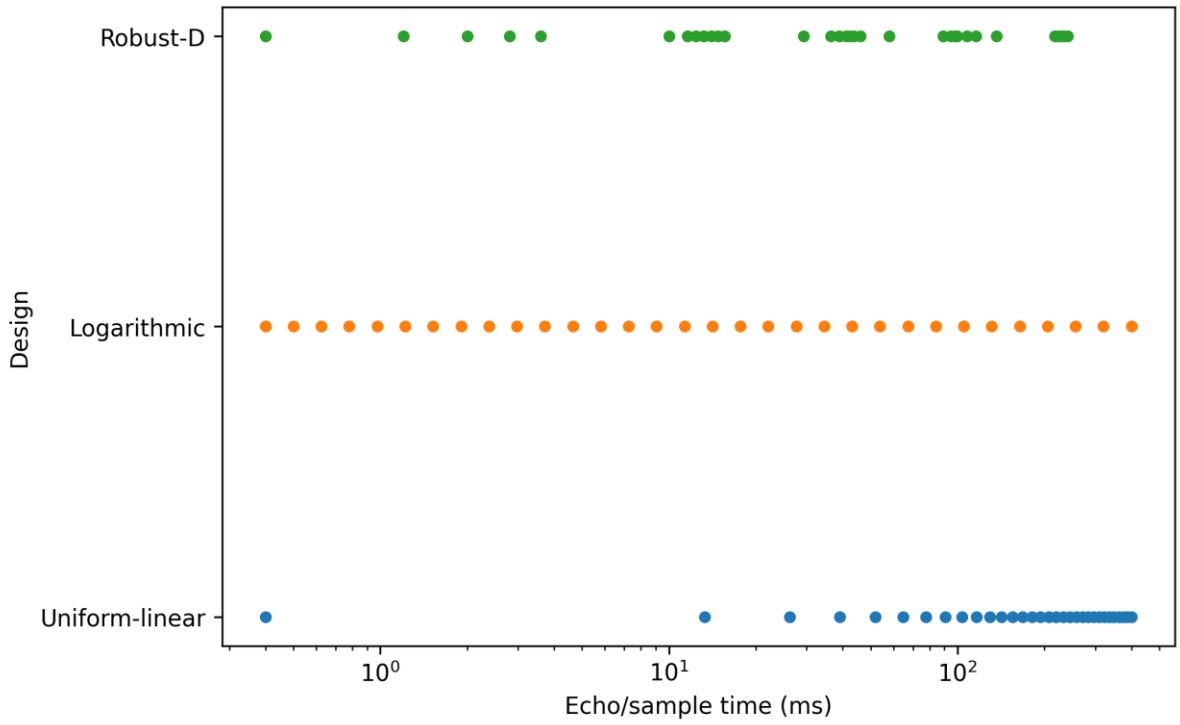


Fig. 1 The three 32-point sampling designs on a logarithmic time axis. The robust D-optimal schedule concentrates observations into several information-rich time regions rather than distributing them uniformly

2.5 Information metrics and Monte Carlo validation

Each schedule was evaluated at all 120 prior models. The recorded metrics were mean and 10th-percentile log det F, the median Fisher-matrix condition number, and the median and 90th-percentile Cramer-Rao standard deviations in the scaled coordinates. These summarize average information, poor-case information, numerical conditioning, and marginal precision.

Monte Carlo validation used three benchmark parameter sets: a well-separated case (A = 1, T2,s = 15 ms, T2,l = 150 ms, f = 0.40), a moderately separated case (28 ms, 90 ms, f = 0.50), and a low-short-fraction case (12 ms, 130 ms, f = 0.15). For each schedule, scenario, and signal-to-noise ratio (SNR = 40 or 80), 150 independent noise realizations were generated. Noise standard deviation was defined as 1/SNR relative to unit initial amplitude.

Nonlinear least squares used transformed parameters that guaranteed positivity and ordering: log A, log T2,s, log(T2,l - T2,s), and logit f. Starting points were randomly perturbed around the true values, and broad bounds were imposed. Relative RMSE was calculated for each fitted parameter. Because all simulations use the correct forward model, the Monte Carlo errors measure statistical difficulty rather than model mismatch.

2.6 Minimum-time/dead-time sensitivity

The minimum accessible echo time can be limited by probe dead time, pulse duration, resonator ring-down, digitizer timing, and sequence constraints. To examine this effect independently of the main schedule comparison, the robust design was recomputed for minimum candidate times of 0.4, 1.0, 2.0, and 5.0 ms while retaining the 400 ms upper limit

and 32 selected points. The mean prior-ensemble log determinant was then compared across these designs. This calculation estimates the information cost of losing access to the earliest portion of the decay.

2.7 Computational implementation

Calculations were performed in Python 3.13.5 with NumPy 2.3.5 and SciPy 1.17.0; figures were generated with Matplotlib 3.10.8 [18,19]. Random number generation used a fixed seed of 20260901. The design and Monte Carlo algorithms use only deterministic model evaluations after the specified random draws. No experimental data were used.

Table 1 Qualitative structure of the three 32-point designs

Design	First selected times (ms)	Late-time coverage
Uniform-linear	0.4, 13.3, 26.2, 39.1, ...	to 400
Logarithmic	0.4, 0.50, 0.62, 0.78, ...	to 400
Robust-D	0.4, 1.20, 2.00, 2.80, 3.60, ...	clusters near ~10-16, ~29-58, ~89-137, and ~217-241

3 Results

3.1 Robust design allocates samples to distinct sensitivity windows

The robust schedule in Fig. 1 is visibly clustered. The first five samples occur at 0.4, 1.20, 2.00, 2.80, and 3.60 ms, capturing the initial amplitude and the earliest curvature of fast components. A dense group appears between approximately 10 and 16 ms, another group between roughly 29 and 58 ms, and broader late-time groups extend through approximately 89-137 ms and 217-241 ms. The exact 32 selected times are reported in Appendix A.

This structure is consistent with the sensitivity functions in Fig. 8. Information about log amplitude is largest early. Sensitivity to T2,s rises and falls on the scale of the short component, whereas sensitivity to T2,l peaks later. The fraction sensitivity depends on the difference between the two component exponentials and becomes largest when the components have separated enough to produce distinct curvature but before both have decayed to very low amplitude. A robust schedule therefore has a reason to revisit several time regions rather than sweep smoothly from early to late times.

The logarithmic schedule captures the early decade efficiently but spreads observations continuously across log time. The uniform-linear schedule places very few points in the first tens of milliseconds: after the first sample at 0.4 ms, the next occurs near 13 ms. This leaves a large information gap for short T2 values at the lower end of the prior range.

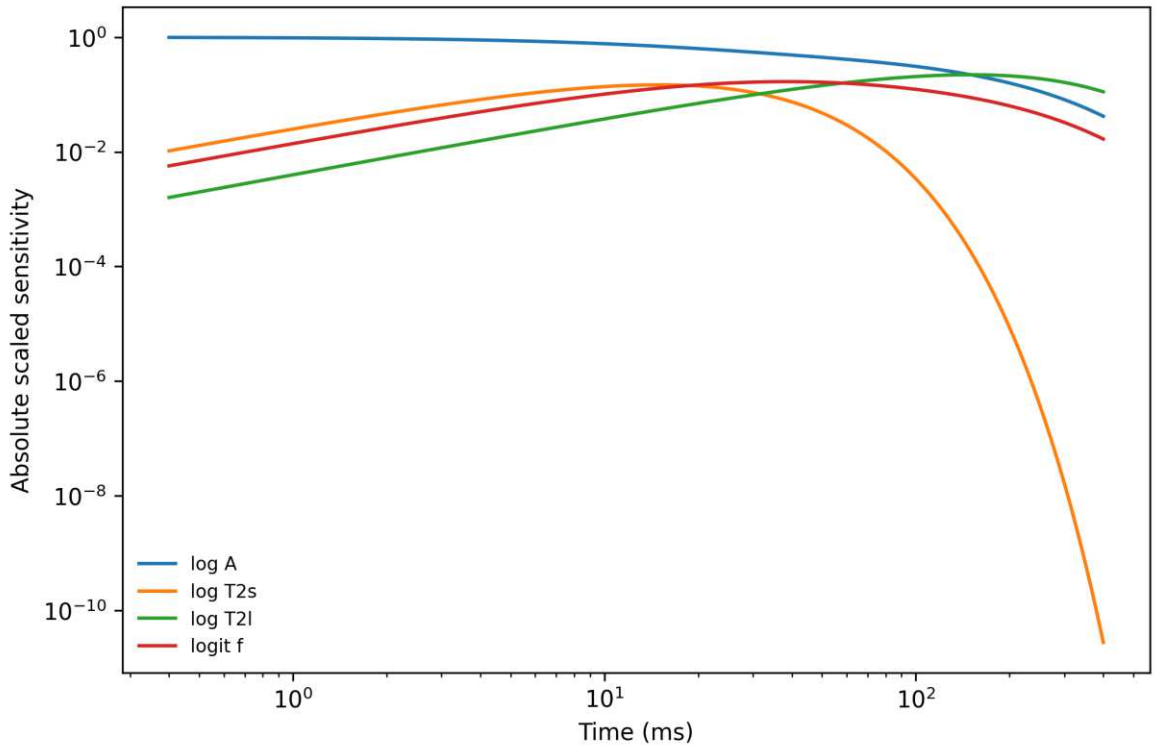


Fig. 2 Absolute scaled parameter sensitivities for the well-separated benchmark decay (T2,s = 15 ms, T2,l = 150 ms, f = 0.40). Different parameters are informed by different portions of the decay

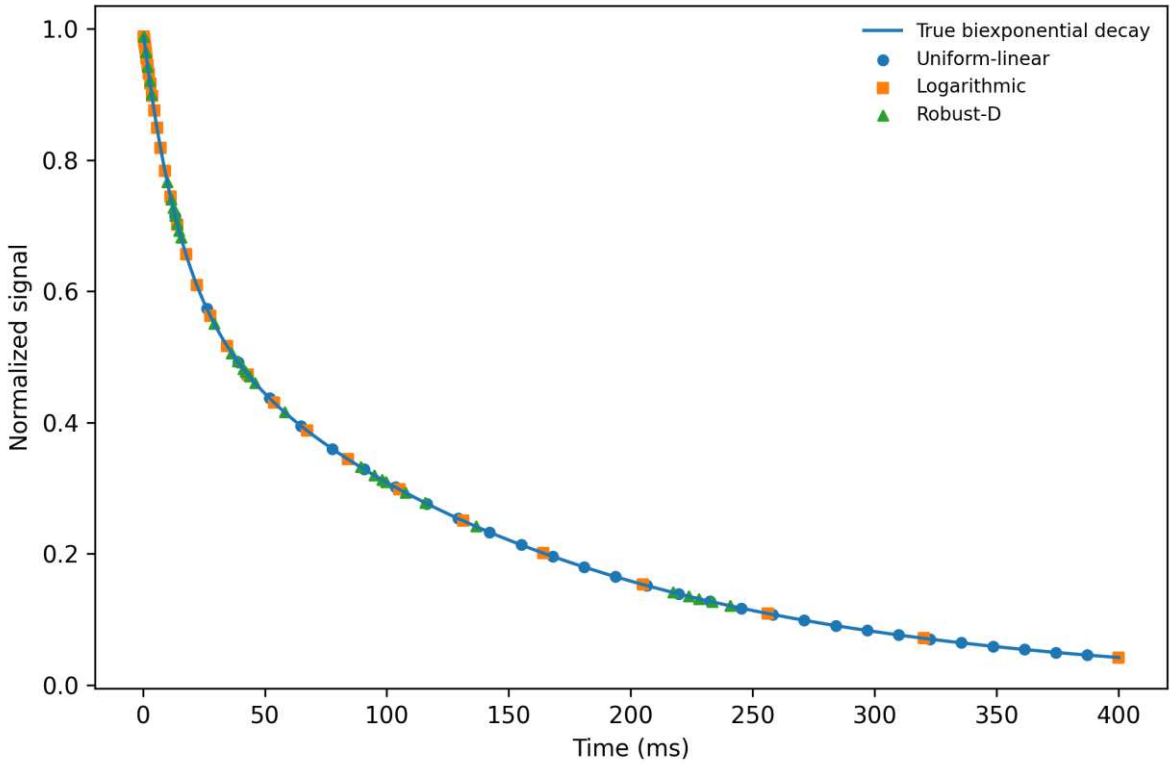


Fig. 3 Well-separated biexponential decay with the sample locations from the three schedules superimposed. Marker density illustrates the different allocation of the common 32-sample budget

3.2 Robust design increases joint information across the prior ensemble

Table 2 Prior-ensemble Fisher-information performance

Design	Mean log det F	10th percentile	Median condition no.	Median scaled SDs
Uniform-linear	-9.276	-11.826	801	1.03, 9.29, 4.69, 14.00
Logarithmic	-6.965	-9.130	6219	0.38, 7.69, 6.85, 15.38
Robust-D	-6.869	-8.867	2885	0.59, 7.50, 5.71, 13.59

The uniform-linear design has a mean log determinant of -9.276 on the common scaled-information convention. The logarithmic design improves this to -6.965, and the robust design achieves -6.869. Relative to linear sampling, the robust improvement is 2.407 log units, corresponding to an approximately 11.1-fold determinant ratio under identical scaling. The logarithmic design also performs strongly, with an approximately 10.1-fold determinant ratio relative to linear.

The 10th-percentile log determinant improves from -11.826 for linear sampling to -8.867 for robust sampling. Thus the robust advantage is not confined to the easiest members of the prior ensemble. The median condition number is 801 for linear sampling, 6,219 for logarithmic sampling, and 2,885 for robust sampling. The condition number does not track determinant monotonically because the three designs distribute information differently among parameter directions. The logarithmic schedule, for example, has high total information but can be anisotropic in parameter space.

The median scaled Cramer-Rao deviations show the tradeoff more directly. Relative to the logarithmic design, the robust schedule sacrifices some amplitude precision but improves the median T2,l and fraction directions. The D-optimal objective rewards the joint covariance volume, so it need not minimize every single marginal variance independently.

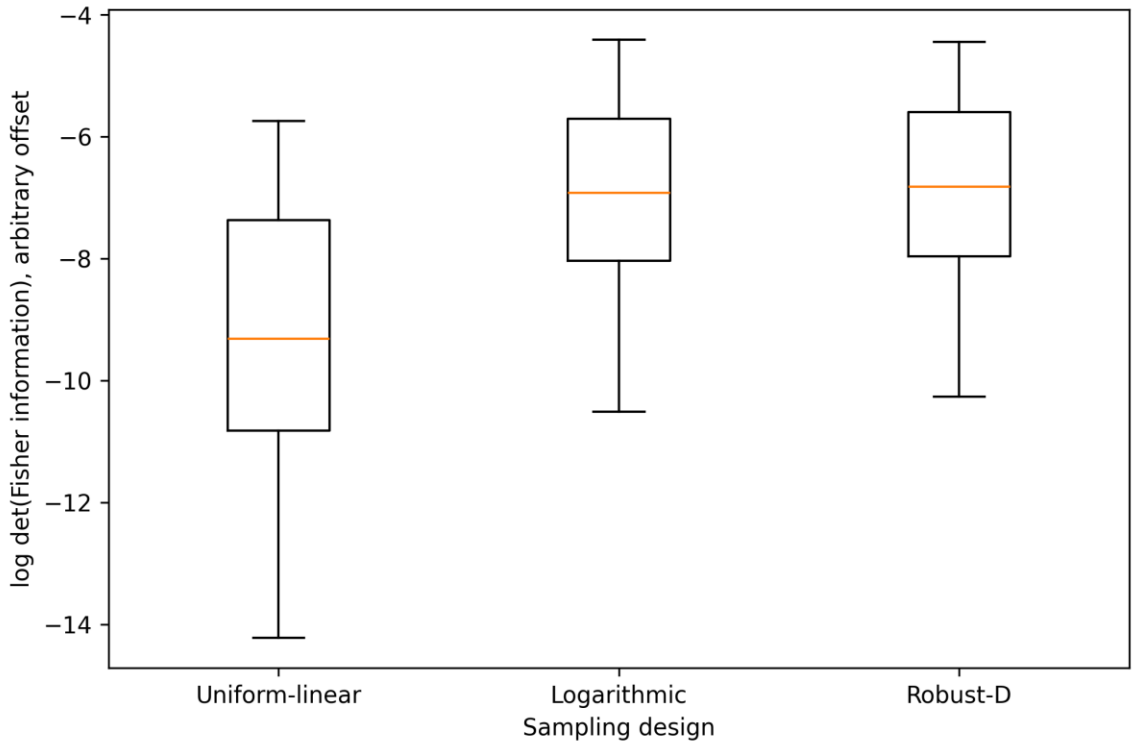


Fig. 4 Distribution of log det(Fisher information) across the 120-member prior ensemble. Both nonuniform designs dominate uniform-linear sampling, with the robust design providing the highest average information

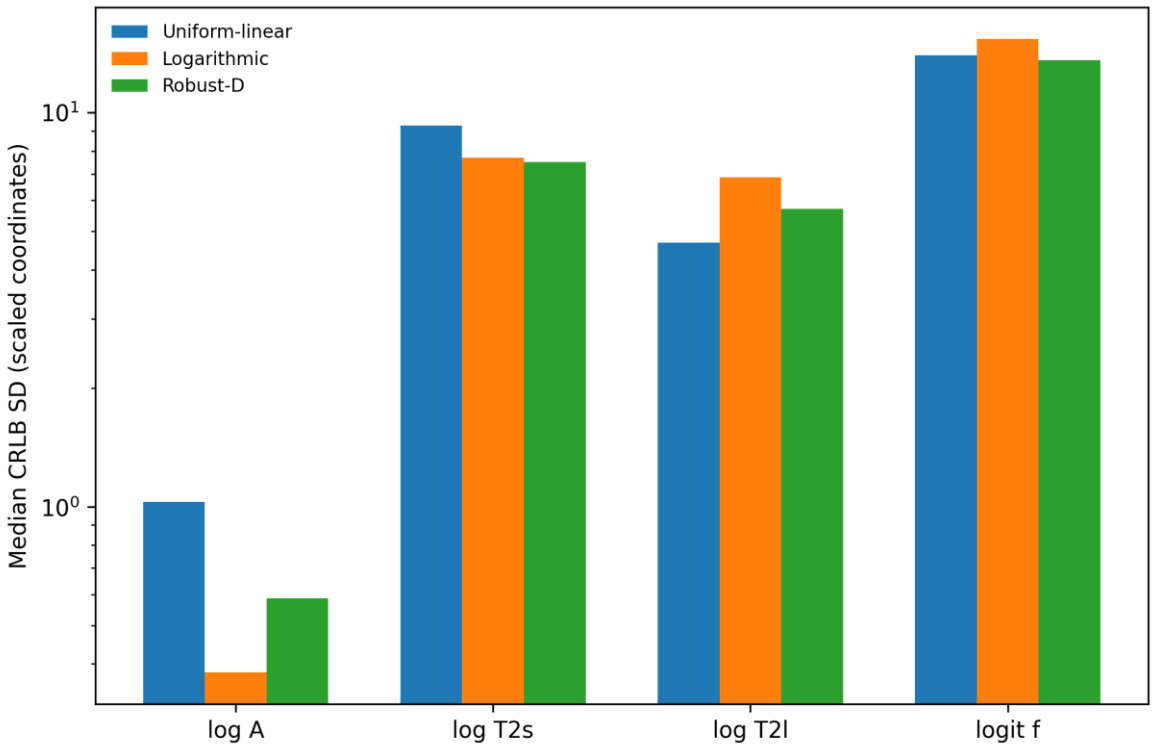


Fig. 5 Median Cramer-Rao standard deviations in the scaled parameter coordinates across the prior ensemble. The logarithmic axis emphasizes the different precision tradeoffs among parameters

3.3 Well-separated components: robust timing improves fast-component recovery

In the well-separated benchmark, all three schedules recover the two time constants at SNR 80, but their precision differs substantially. Uniform-linear sampling gives relative RMSEs of 10.1% for T2,s and 2.65% for T2,l. Logarithmic sampling improves T2,s to 8.30% but worsens T2,l to 4.47%. The robust design gives 7.13% for T2,s and 3.54% for T2,l. The fraction relative RMSE is 4.15%, 5.14%, and 3.67% for the linear, logarithmic, and robust designs, respectively.

At SNR 40 the same pattern persists under greater noise. T2,s RMSE is 21.3% for linear sampling, 16.5% for logarithmic sampling, and 14.8% for robust sampling. The robust schedule therefore reduces the fast-component RMSE by roughly 30% relative to the uniform design while maintaining useful long-component precision.

The Monte Carlo distribution in Fig. 6 shows that the gain is a contraction of the estimator spread rather than a change produced by a small number of outliers. All fits use the same nonlinear optimizer, bounds, and starting-point perturbation rule, so the improvement arises from the timing of the observations.

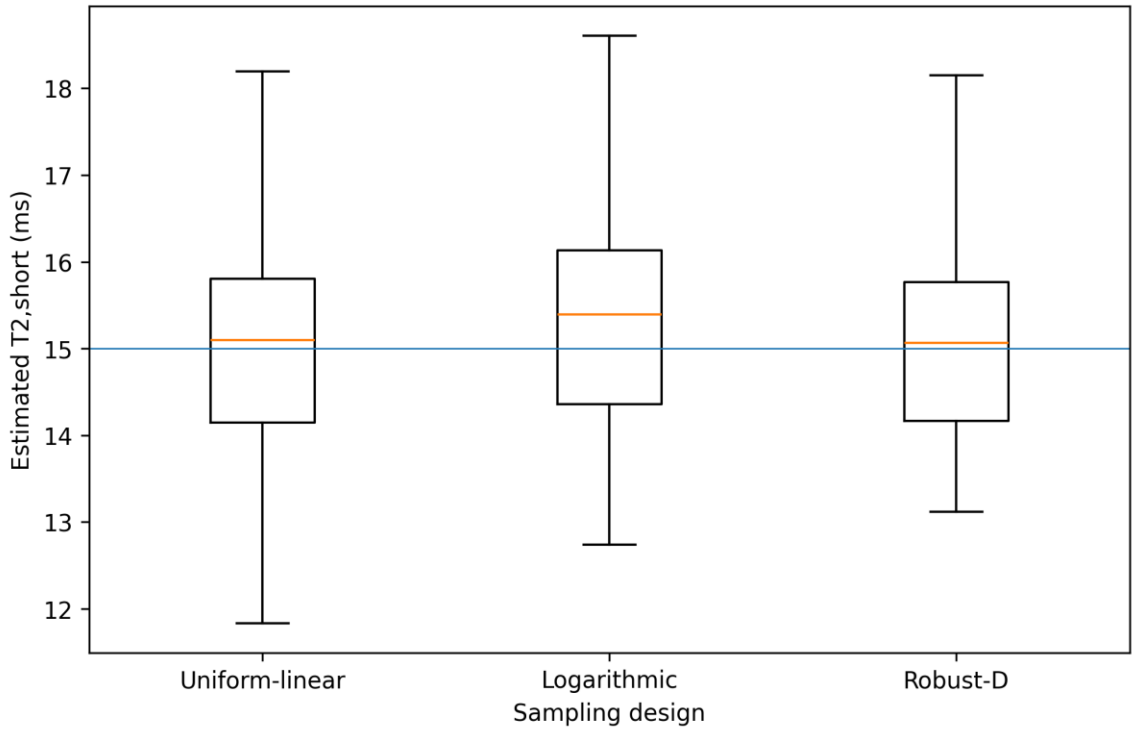


Fig. 6 Monte Carlo estimates of T2,s for the well-separated benchmark at SNR 80. The horizontal reference line marks the true value of 15 ms

3.4 Challenging component separation reveals a fundamental limit

The moderately separated benchmark (T2,s = 28 ms, T2,l = 90 ms) is much harder. At SNR 40, relative RMSE for T2,l ranges from approximately 47.6% for linear sampling to 76.9% for logarithmic sampling, while robust sampling gives approximately 50.0%. At SNR 80 the errors improve but remain substantial: T2,s RMSE is 15.0%, 16.9%, and 14.1% for linear, logarithmic, and robust designs, while T2,l RMSE is 11.2%, 20.4%, and 13.7%.

This case illustrates an important limitation of acquisition optimization. A better schedule can increase Fisher information, but it cannot create strong identifiability when two exponential basis functions are intrinsically similar over the available time window. In such regimes, the fitted fraction and time constants trade off, and the estimator may require higher SNR, more total samples, stronger prior constraints, a different acquisition contrast, or an alternative physical model.

The low-short-fraction benchmark presents a different difficulty: the fast component is well separated in time but contributes only 15% of the initial amplitude. Here early sampling is particularly important. At SNR 80, uniform-linear sampling gives a T2,s relative RMSE of 31.9%. Logarithmic sampling reduces this to 20.1%, and the robust schedule to 19.4%. The robust fraction RMSE is also lowest at 10.6% compared with 12.2% and 12.0% for linear and logarithmic sampling.

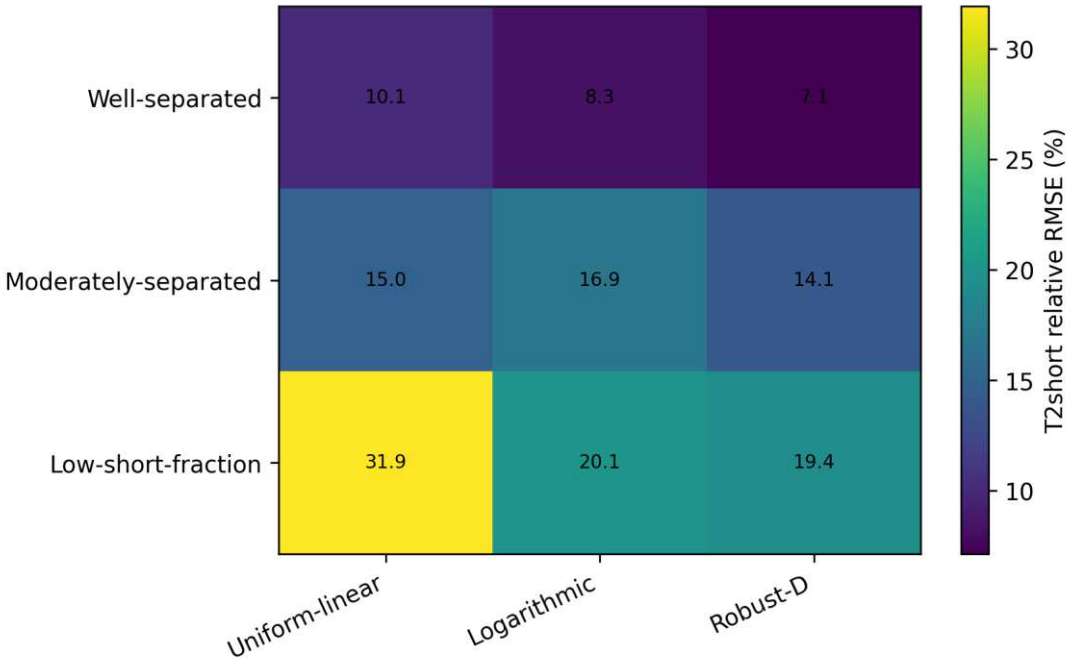


Fig. 7 Relative RMSE of T2,s at SNR 80 for the three benchmark regimes and sampling designs

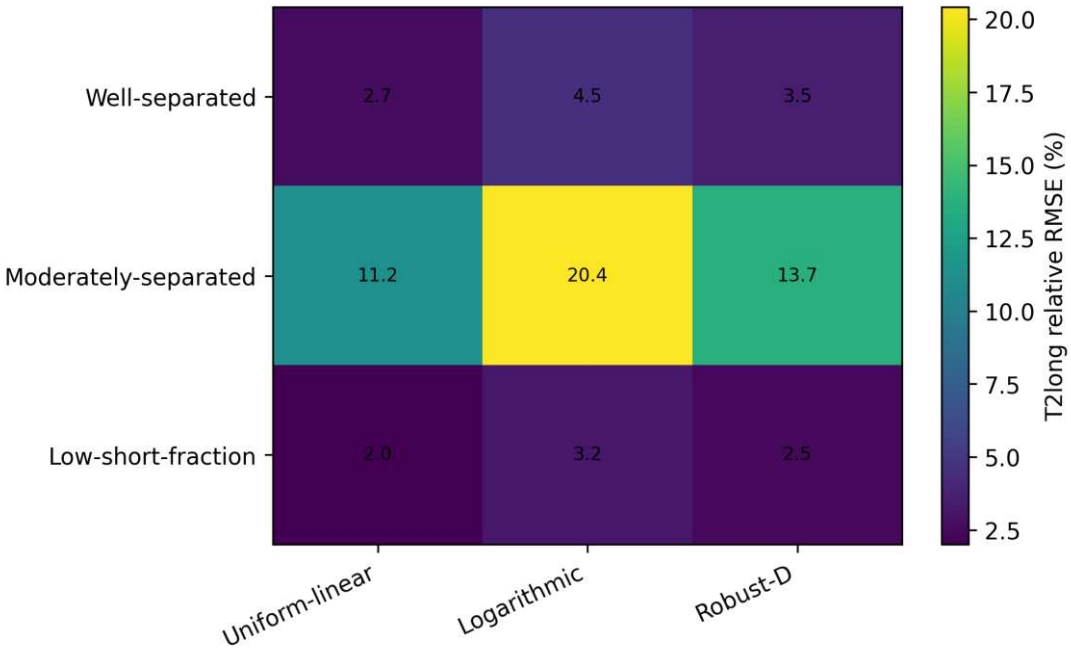


Fig. 8 Relative RMSE of T2,l at SNR 80 for the three benchmark regimes and sampling designs

3.5 Increasing the minimum accessible time reduces robust information

The mean robust-design log determinant decreases monotonically as the minimum candidate time is increased. The values are -6.869 for a 0.4 ms minimum, -7.042 for 1.0 ms, -7.330 for 2.0 ms, and -8.197 for 5.0 ms. Raising the minimum from 0.4 to 5 ms therefore loses approximately 1.33 log-information units, corresponding to a determinant reduction of roughly a factor of 3.8 under the common scaling.

The result is physically intuitive. Very early observations anchor the total amplitude and provide leverage on fast decays before they have substantially attenuated. Once that information is lost, later samples cannot completely replace it because amplitude, fraction, and T2,s become more correlated. The quantitative curve in Fig. 9 can be interpreted as the cost of probe or sequence dead time for the broad prior ensemble used here.

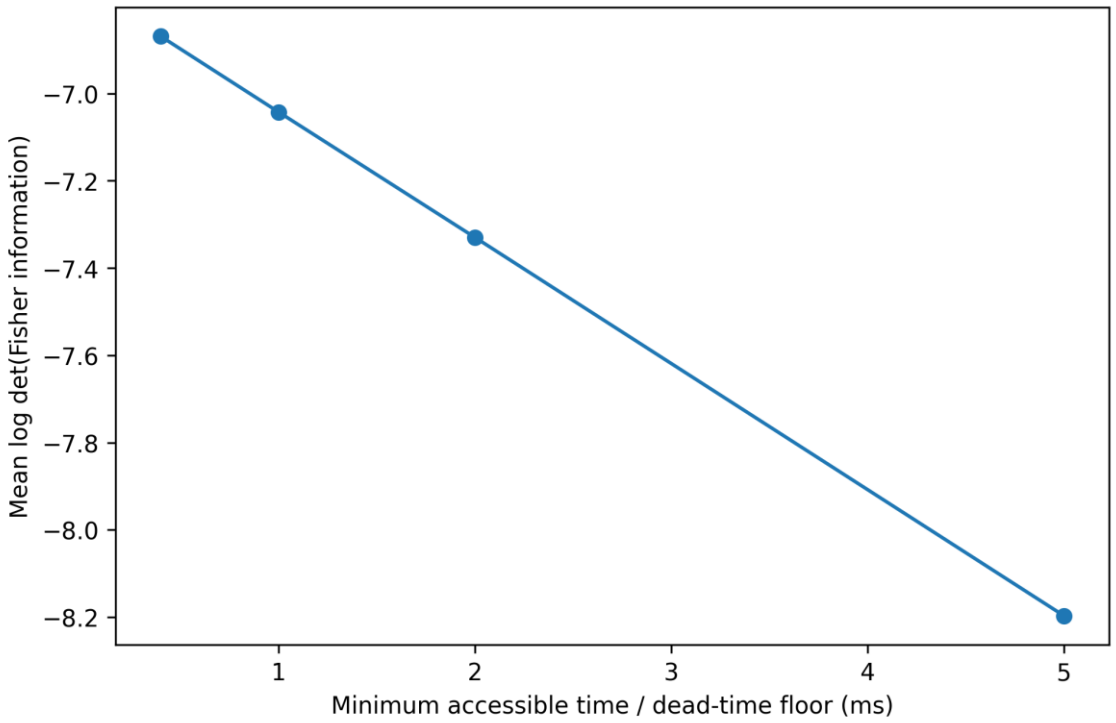


Fig. 9 Prior-averaged robust-design information as a function of the minimum accessible sample time. Later dead-time floors reduce the achievable joint parameter information

3.6 Detailed Monte Carlo summary

Table 3 Monte Carlo relative RMSE at SNR 80

Scenario	Design	T2,s RMSE %	T2,l RMSE %	f RMSE %
Well-separated	Uniform-linear	10.1	2.7	4.1
Well-separated	Logarithmic	8.3	4.5	5.1
Well-separated	Robust-D	7.1	3.5	3.7
Moderately-separated	Uniform-linear	15.0	11.2	17.8
Moderately-separated	Logarithmic	16.9	20.4	24.4
Moderately-separated	Robust-D	14.1	13.7	19.8
Low-short-fraction	Uniform-linear	31.9	2.0	12.2
Low-short-fraction	Logarithmic	20.1	3.2	12.0
Low-short-fraction	Robust-D	19.4	2.5	10.6

Table 3 emphasizes that there is no schedule that dominates every marginal parameter in every regime. Uniform-linear timing can estimate the long component efficiently when it receives many late observations, whereas logarithmic timing strongly favors early-scale coverage. The robust design is a compromise optimized for joint information over uncertainty

in the true model. Its main advantage is consistency: it is never the worst design for T2,s in the three SNR-80 benchmarks and provides the lowest fraction RMSE in the well-separated and low-short-fraction cases.

This distinction is important when the acquisition objective is unknown at design time. If one knows in advance that only a single long T2 is important, a dedicated monoexponential optimal design such as that derived by Jones et al. is preferable [3]. The robust biexponential design is intended for a different situation: two components may exist, their time scales and fraction are uncertain, and all parameters matter.

4 Discussion

4.1 Why the robust schedule is clustered

The clustered schedule is a characteristic feature of optimal-design problems. Repeated or near-repeated observations can be statistically efficient because parameter information is concentrated where the model derivatives are largest or most linearly independent. For a monoexponential decay, Jones and co-workers found an optimal pattern that places a substantial fraction of observations at only a small number of characteristic times rather than uniformly over the decay [3]. The present four-parameter, prior-robust problem produces more support regions because it must separate two relaxation time scales and a fraction simultaneously.

From an NMR implementation perspective, closely spaced support points should not be interpreted rigidly as a requirement for arbitrarily variable interpulse delays within one CPMG train. The information density can instead guide allocation of retained echoes or averages. For example, a multi-train protocol could use several echo spacings chosen to place acquired echoes near the high-value regions. A constant-spacing train could retain a nonuniform subset of echoes for storage and fitting, provided the underlying pulse train itself does not introduce additional sequence-dependent effects.

A future hardware-constrained design can replace the continuous timing candidates used here with the exact echo times achievable by a particular sequence. The same greedy information calculation then returns a schedule that is directly executable on that instrument.

4.2 Relationship to inverse Laplace and discrete fitting methods

Relaxometry analysis is often discussed as a choice between inverse Laplace distributions and discrete multiexponential fitting. That distinction is important for interpretation but does not remove the sampling problem. Both classes of estimators depend on the singular-value structure and noise amplification of the forward exponential kernel [4,5,8,10,11]. If the acquisition contains little sensitivity to a fast component, no regularization method can fully recover information that was never measured.

The present D-optimal criterion is derived for a four-parameter discrete model, so its exact schedule is not guaranteed to be optimal for a broad continuous T2 distribution. Nevertheless, the design principle should transfer. One can define an alternative parameterization of a distribution, a set of basis weights, moments, or matrix-pencil coefficients and then

calculate the corresponding information or resolution criterion. Parziale et al. showed that time-domain spacing materially affects matrix-pencil recovery [7], while Sørland and colleagues demonstrated the use of synthetic multiexponential data for testing relaxation-analysis methodology [8]. The current study adds prior-robust nonuniform time selection to that methodological landscape.

For inverse Laplace analysis, one promising extension is to optimize a resolution matrix or expected posterior covariance under regularization rather than a finite-dimensional determinant. Another is to design for specific distribution summaries, such as mean T2 or component fractions within predefined relaxation windows.

4.3 Implications for low-field and benchtop NMR

Low-field permanent-magnet systems are attractive for relaxometry because they are compact and can provide rapid time-domain measurements, but they often operate with finite dead time and modest SNR. In that setting, collecting more echoes is not always the most efficient response: many late-time points can contain little signal while consuming storage or repetition resources. Timing design offers an additional route to precision improvement that does not require a stronger magnet.

The robust schedule in this study suggests that a sparse measurement budget should deliberately preserve early echoes and several intermediate-time clusters. The dead-time calculation shows why reducing the earliest accessible time can have a measurable information benefit. Instrumental changes that shorten ring-down or receiver recovery may therefore improve multicomponent identifiability even if the nominal SNR per echo is unchanged.

At the same time, real CPMG data may be distorted by imperfect 180-degree pulses, B1 inhomogeneity, stimulated echoes, diffusion in internal gradients, and exchange. If those effects are appreciable, the optimal design should be based on an extended Bloch or coherence-pathway forward model rather than the ideal exponential equation. The information framework remains the same; only the Jacobian changes.

4.4 Limitations

The study assumes additive independent Gaussian noise. Phase-corrected real-valued echo amplitudes can often be treated approximately this way, but magnitude reconstruction produces Rician or noncentral-chi behavior at low SNR. Correlated receiver noise or echo-to-echo drift would also change the Fisher matrix. Generalized least squares can accommodate a known noise covariance by replacing the simple J^T J form with J^T Sigma^-1 J.

The two-component model assumes fixed exponential time constants and no exchange during the acquisition. Chemical exchange can modify CPMG relaxation behavior and may carry additional spectral-density information [12,13]. Similarly, porous media can exhibit multiexponential relaxation arising from diffusion and boundary conditions rather than discrete noninteracting compartments [9]. A robust design for such systems should be built from the appropriate physical forward model.

The prior ensemble is deliberately broad but arbitrary. A different application may require T2 values outside 8-300 ms, much smaller fractions, or more than two components. Robust design performance depends on that prior design space.

The correct interpretation of the selected times is therefore conditional: they are optimal for the specified ensemble and 32-sample budget, not universal CPMG timings.

Finally, the greedy algorithm is not guaranteed to find the global optimum among all 32-point subsets. It is computationally efficient and produced substantial gains over the comparators, but exchange algorithms, mixed-integer optimization, continuous approximate-design methods, or stochastic search may identify still better schedules. The present work establishes the magnitude and structure of the timing effect rather than claiming a mathematically unique global solution.

4.5 Practical workflow for adapting the design

To adapt the method to a real experiment, the user can follow five steps. First, define the forward signal equation including sequence imperfections that materially affect the echoes. Second, construct a prior ensemble spanning the plausible T2 values, fractions, nuisance parameters, and SNR. Third, restrict candidate times to those actually executable by the spectrometer. Fourth, select an information criterion that matches the scientific goal - determinant for joint precision, trace for average variance, or a weighted criterion if one parameter is primary. Fifth, validate the proposed timing with Monte Carlo fits before collecting samples.

This workflow turns echo scheduling into an explicit part of pulse-sequence and acquisition design. It also exposes tradeoffs before instrument time is spent. If no 32-point design achieves acceptable uncertainty for closely separated components, the remedy may be to increase SNR, collect a second contrast, alter field strength, or reformulate the scientific question rather than simply adding more points at the same uninformative times.

5 Conclusions

A 32-point robust D-optimal timing design for two-component CPMG relaxometry was derived over a broad ensemble of unknown T2 values and fractions. The design concentrates observations into several characteristic time regions and increases the prior-averaged Fisher-information determinant by about elevenfold relative to a uniformly linear schedule under common parameter scaling. Monte Carlo nonlinear fits confirm that the information advantage can translate into lower T2,s and fraction errors, particularly when the fast component is weak or when uniform timing undersamples the early decay.

The results also define the limits of timing optimization. When T2,s and T2,l are too close, parameter uncertainty remains high even for the robust schedule because the basis functions are intrinsically difficult to distinguish. Increasing the minimum accessible echo time further reduces information, quantifying the value of early-time access. Together, these findings support treating CPMG sampling times as a design variable that can be optimized for the expected parameter space rather than chosen only by convention.

Statements and Declarations

Funding: [To be completed by the author before submission.]

Competing interests: [To be completed by the author before submission.]

Ethical approval: Not applicable; this is a computational study using synthetic data only.

Data availability: No experimental data were used. All reported results were generated from the equations, parameter ranges, schedules, and random seed specified in the manuscript.

Appendix A: Robust 32-Point Schedule

Table A1 Exact robust-D selected times on the 0.4-400 ms candidate raster

Index	Time (ms)
1	0.400
2	1.201
3	2.002
4	2.802
5	3.603
6	10.010
7	11.611
8	12.412
9	13.213
10	14.014
11	14.814
12	15.615
13	29.229
14	36.436
15	38.838
16	41.241
17	42.042
18	43.643
19	46.046
20	58.058
21	89.289

22	94.895
23	98.098
24	99.699
25	107.707
26	115.715
27	136.536
28	217.417
29	223.824
30	227.828
31	233.433
32	240.640

References

1. H.Y. Carr, E.M. Purcell, Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments, Phys. Rev. 94, 630-638 (1954). https://doi.org/10.1103/PhysRev.94.630
2. S. Meiboom, D. Gill, Modified Spin-Echo Method for Measuring Nuclear Relaxation Times, Rev. Sci. Instrum. 29, 688-691 (1958). https://doi.org/10.1063/1.1716296
3. J.A. Jones, P. Hodgkinson, A.L. Barker, P.J. Hore, Optimal Sampling Strategies for the Measurement of Spin-Spin Relaxation Times, J. Magn. Reson. B 113, 25-34 (1996). https://doi.org/10.1006/jmrb.1996.0151
4. K.P. Whittall, A.L. MacKay, Quantitative interpretation of NMR relaxation data, J. Magn. Reson. 84, 134-152 (1989). https://doi.org/10.1016/0022-2364(89)90011-5
5. S.W. Provencher, CONTIN: A general purpose constrained regularization program for inverting noisy linear algebraic and integral equations, Comput. Phys. Commun. 27, 229-242 (1982). https://doi.org/10.1016/0010-4655(82)90174-6
6. P. Galvosas, P.T. Callaghan, Multi-dimensional inverse Laplace spectroscopy in the NMR of porous media, C. R. Physique 11, 172-180 (2010). https://doi.org/10.1016/j.crhy.2010.06.014
7. M. Parziale, D. Woertge, B. Mohebbi, J. Claussen, M.P. Augustine, Improving the Resolution of MPM Recovered Relaxometry Parameters with Proper Time Domain Sampling, Appl. Magn. Reson. 54, 1391-1409 (2023). https://doi.org/10.1007/s00723-023-01596-x
8. G.H. Sorland, H.W. Anthonsen, S. Simon, A Method for Analyzing Data from 1- and 2-Dimensional Relaxation and Diffusion NMR Experiments by Determination of their Expectation Values and Standard Deviations, Appl. Magn. Reson. 55, 1437-1453 (2024). https://doi.org/10.1007/s00723-024-01718-z
9. A. Afrough, Derivation of Multi-Exponential Magnetic Resonance Relaxation Equations in Simple Pore Geometries, Appl. Magn. Reson. 56, 313-358 (2025). https://doi.org/10.1007/s00723-024-01729-w
10. A. Afrough, R. Mokhtari, K.L. Feilberg, Simple MATLAB and Python scripts for multi-exponential analysis, Magn. Reson. Chem. 62, 698-711 (2024). https://doi.org/10.1002/mrc.5453

11. G.S. Ioannidis et al., Inverse Laplace transform and multiexponential fitting analysis of T2 relaxometry data: a phantom study with aqueous and fat containing samples, Eur. Radiol. Exp. 4, 28 (2020). https://doi.org/10.1186/s41747-020-00154-5
12. R.P.A.G. Teixeira, S.J. Malik, J.V. Hajnal, Joint system relaxometry and Cramer-Rao lower bound optimization of sequence parameters: A framework for enhanced precision of DESPOT T1 and T2 estimation, Magn. Reson. Med. 79, 234-245 (2018). https://doi.org/10.1002/mrm.26670
13. J. Stepisnik, A. Mohoric, Insight into Details of Chemical Exchange Kinetics Studied by NMR CPMG Method, Appl. Magn. Reson. (2024). https://doi.org/10.1007/s00723-024-01679-3
14. A. Mohoric, J. Stepisnik, Chemical Exchange Rate Study by NMR CPMG Method, Appl. Magn. Reson. 54, 1411-1422 (2023). https://doi.org/10.1007/s00723-023-01621-z
15. S. Fleischer, S. Lehmkuhl, L. Lohmann et al., Approaching the Ultimate Limit in Measurement Precision with RASER NMR, Appl. Magn. Reson. 54, 1241-1270 (2023). https://doi.org/10.1007/s00723-023-01597-w
16. A statistical learning perspective on the inversion of NMR relaxation data, AIP Adv. 12, 065230 (2022). https://doi.org/10.1063/5.0092178
17. C.R. Harris et al., Array programming with NumPy, Nature 585, 357-362 (2020). https://doi.org/10.1038/s41586-020-2649-2
18. P. Virtanen et al., SciPy 1.0: fundamental algorithms for scientific computing in Python, Nat. Methods 17, 261-272 (2020). https://doi.org/10.1038/s41592-019-0686-2