Journal of Biomolecular NMR (2006) 35: 117–135
DOI 10.1007/s10858-006-9007-z

Article

Model-free model elimination: A new step in the model-free dynamic analysis of NMR relaxation data

Edward J. d'Auvergne* & Paul R. Gooley
Department of Biochemistry and Molecular Biology, Bio21 Institute of Biotechnology and Molecular Science, University of Melbourne, Parkville, Victoria, 3010, Australia

Received 15 September 2005; Accepted 21 March 2006

Key words: data analysis, error analysis, mathematical modelling, model-free analysis, model elimination, model selection, model validation, NMR relaxation

Abstract

Model-free analysis is a technique commonly used within the field of NMR spectroscopy to extract atomic resolution, interpretable dynamic information on multiple timescales from the R_1, R_2, and steady state NOE. Model-free approaches employ two disparate areas of data analysis, the discipline of mathematical optimisation, specifically the minimisation of a χ^2 function, and the statistical field of model selection. By searching through a large number of model-free minimisations, which were setup using synthetic relaxation data whereby the true underlying dynamics is known, certain model-free models have been identified to, at times, fail. This has been characterised as either the internal correlation times, τ_e, τ_f, or τ_s, or the global correlation time parameter, local τ_m, heading towards infinity, the result being that the final parameter values are far from the true values. In a number of cases the minimised χ^2 value of the failed model is significantly lower than that of all other models and, hence, will be the model which is chosen by model selection techniques. If these models are not removed prior to model selection the final model-free results could be far from the truth. By implementing a series of empirical rules involving inequalities these models can be specifically isolated and removed. Model-free analysis should therefore consist of three distinct steps: model-free minimisation, model-free model elimination, and finally model-free model selection. Failure has also been identified to affect the individual Monte Carlo simulations used within error analysis. Each simulation involves an independent randomised relaxation data set and model-free minimisation, thus simulations suffer from exactly the same types of failure as model-free models. Therefore, to prevent these outliers from causing a significant overestimation of the errors the failed Monte Carlo simulations need to be culled prior to calculating the parameter standard deviations.

Abbreviations: AIC – Akaike's Information Criteria; χ^2 – chi-squared function; c_k – constraint value; CSA – Chemical Shift Anisotropy; DMG – Double Motion Grid; ϵ_i – elimination value; NOE – nuclear Overhauser effect; r – bond length; R_1 – spin-lattice relaxation rate; R_2 – spin-spin relaxation rate; R_{ex} – chemical exchange relaxation rate; RG – R_{ex} Grid; S^2, S_f^2, and S_s^2 – model-free generalised order parameters; τ_e, τ_f, and τ_s – model-free effective internal correlation times; τ_m – global rotational correlation time.

*To whom correspondence should be addressed.
E-mail: edward@nmr-relax.com

Introduction

Model-free analysis of heteronuclear NMR relaxation data, specifically the R_1 and R_2 relaxation rates together with the steady state NOE, is used to extract easily interpretable dynamical information describing both the overall tumbling and the internal motions of a macromolecule (Lipari and Szabo, 1982a, b; Clore et al., 1990). The internal dynamics is quantified by three types of parameter: the generalised order parameter S^2 which characterises the amplitude of the motion; the effective internal correlation time τ_e which links the amplitude to a timescale; and the chemical exchange relaxation parameter R_{ex} which is an indicator of slower microsecond to millisecond timescale dynamics. The order parameters range from one for complete rigidity to zero for high mobility and the timescale of extractable motions can range from the sub-picosecond range to the global tumbling time which is within the nanosecond range. The theory has been extended to include motions on two different timescales (Clore et al., 1990) where the faster of the motions is parameterised by S_f^2 and τ_f, the amplitude and timescale respectively, while the slower is parameterised by S_s^2 and τ_s.

The internal model-free dynamics are decoupled by construction from the global diffusion (Lipari and Szabo, 1982a). The overall rotational diffusion of the entire molecule is described by a three-dimensional Cartesian tensor. Apart from orientation, the tensor can be described by three independent terms dependent on its eigenvalues. If the three eigenvalues of the tensor are equal, the molecule tumbles isotropically and can be described by the single parameter, τ_m, the global correlation time. If only two of the eigenvalues are equal, the diffusion is anisotropic and axially symmetric and is described by the addition of a parameter accounting for the anisotropy. If all three eigenvalues are different, the global correlation time, anisotropy, and asymmetry (or rhombicity) characterise the fully anisotropic diffusion tensor. In certain situations, the assumption that all residues of the molecule under study will experience the same global rotational diffusion may not be the best model of the entire system. To redress this problem, another model whereby all residues are assumed to tumble independently can be optimised. Each residue has its own global correlation time parameter called the local τ_m.

While appearing to be a contradiction in terms, the expression 'model-free model' encapsulates two different concepts. 'Model-free' refers to the absence of a physical model delineating the superposition of all trajectories and microstates of the internal motions of the system. These motions are solely characterised by amplitude and timescale. In contrast, the second use of the word 'model' originates from the various 'model-free' mathematical models used to represent different classes of motion. By using various combinations of model-free parameters, diverse mathematical models can be constructed. These models overlap significantly and have no clear boundaries as, not only are many models parametric restrictions of others but, the simpler models will satisfactorily approximate the more complex ones. The distinctions between them are not physical but purely statistical and depend on the data collected and, most importantly, the errors. Although eight models of varying complexity can be created (Fushman et al., 1997; Orekhov et al., 1999; Korzhnev et al., 2001; Zhuravleva et al., 2004),

$$m1 = \{S^2\}, \tag{1.1}$$

$$m2 = \{S^2, \tau_e\}, \tag{1.2}$$

$$m3 = \{S^2, R_{ex}\}, \tag{1.3}$$

$$m4 = \{S^2, \tau_e, R_{ex}\}, \tag{1.4}$$

$$m5 = \{S^2, S_f^2, \tau_s\}, \tag{1.5}$$

$$m6 = \{S^2, \tau_f, S_f^2, \tau_s\}, \tag{1.6}$$

$$m7 = \{S^2, S_f^2, \tau_s, R_{ex}\}, \tag{1.7}$$

$$m8 = \{S^2, \tau_f, S_f^2, \tau_s, R_{ex}\}, \tag{1.8}$$

often only the five simplest of these are used in the literature. Although models *m6* to *m8* are not utilised, the concepts introduced in this paper apply to these higher, more complex model-free models equally as well. When the local τ_m parameter is assumed rather than a global

rotational diffusion tensor, a new set of models can be created which include an additional dimension to the above models. These are

$$tm1 = \{\tau_m, S^2\}, \tag{2.1}$$

$$tm2 = \{\tau_m, S^2, \tau_e\}, \tag{2.2}$$

$$tm3 = \{\tau_m, S^2, R_{ex}\}, \tag{2.3}$$

$$tm4 = \{\tau_m, S^2, \tau_e, R_{ex}\}, \tag{2.4}$$

$$tm5 = \{\tau_m, S^2, S_f^2, \tau_s\}, \tag{2.5}$$

$$tm6 = \{\tau_m, S^2, \tau_f, S_f^2, \tau_s\}, \tag{2.6}$$

$$tm7 = \{\tau_m, S^2, S_f^2, \tau_s, R_{ex}\}, \tag{2.7}$$

$$tm8 = \{\tau_m, S^2, \tau_f, S_f^2, \tau_s, R_{ex}\}. \tag{2.8}$$

The model-free interpretation consists purely of data analysis methods using techniques from the mathematical fields of modelling and optimisation and the statistical field of model selection. Data analysis usually involves two major steps, model-free minimisation (Palmer et al., 1991; Mandel et al., 1995; Orekhov et al., 1995; Fushman et al., 1997) followed by model-free model selection (Mandel et al., 1995; d'Auvergne and Gooley, 2003; Chen et al., 2004). Standard model-free analysis uses a predefined set of mathematical models representing different types of motion, optimises the parameters of these models to fit the relaxation data, and then uses model selection to determine which of the model-free models best describes the data. However, no attempt is made to identify whether the optimised parameters of any of these models are adequate. Mathematical modelling usually involves a stage whereby a model is validated to check if it is credible (Bellomo and Preziosi, 1994). The user can then decide whether the model is adequate, whether the model should be adjusted, or whether the model should be rejected. Model modification is irrelevant, however, as the model-free mathematical models were constructed prior to model-free analysis (Models 1.1–1.8 and 2.1–2.8), yet certain models should be rejected or eliminated prior to model selection. The effect of not removing failed models is that a model with wildly incorrect parameter values may be selected rather than one with parameters close to the true values.

Apart from its application between model-free minimisation and model-free model selection, model elimination can also be used to improve the accuracy of model-free error analysis. Multiple iterations of model-free minimisation, model elimination, and model selection are required to obtain convergence of the global model, the combination of the global diffusion tensor with the selected model-free model for each residue. After convergence and model selection of the global model, error analysis by Monte Carlo simulations can begin. The simulations consist of taking the model-free parameters which were fitted to the collected data and errors and back-calculating synthetic relaxation data. This data is then randomised n times using the experimental errors, creating n randomised synthetic relaxation data sets. The data sets are assumed to have the same error as the original collected data. Model-free parameters for each simulation are optimised by χ^2 minimisation using exactly the same methods as for the collected data. Each model-free parameter has then been fitted n times creating a probability distribution, the standard deviation of which is a good approximation of the true standard deviation of the model-free parameter fitted to the collected data.

Monte Carlo simulations essentially involve a large number, n, of independent model-free minimisations. Even if the model-free parameters fit the collected relaxation data well and are a reasonable approximation to the truth, the randomisation used to create the synthetic relaxation data sets can create a situation whereby a number of the simulations fail by precisely the same mechanism by which model-free models can fail. These have been identified in the past and have been removed by trimming the upper and lower tails of the χ^2 distribution of simulated results. For example, this has been put into practise as an option in the program Modelfree (Palmer et al., 1991; Mandel et al., 1995). However, the χ^2 distribution is not a good differentiator of failed verses reliable simulations. Using model elimination methods a highly selective and accurate approach can be implemented for the removal of failed simulations.

Methods

Searching the model-free space

To identify failure of model-free models, model-free minimisations were carried out on large numbers of independent data sets. The methodology employed to select and create the different data sets was similar to those presented in d'Auvergne and Gooley (2003). Two three-dimensional grids were created to cover a significant proportion of the model-free parameter space. The first grid, labelled the R_{ex} Grid (RG), consists of the three dimensions $\{S^2, \tau_e, R_{ex}\}$ and covers all single model-free motions with chemical exchange relaxation. This grid corresponds to model *m4* (Model 1.4). When the parameters τ_e or R_{ex} are statistically zero, the models *m1*, *m2*, and *m3* (Models 1.1 to 1.3), which are parametric restrictions of the higher model *m4*, are perfect approximations. The second grid is labelled the Double Motion Grid (DMG) and consists of the dimensions $\{S_f^2, S_s^2, \tau_s\}$. This grid covers part of the model-free space encapsulating the two separate motions of model *m5* (Model 1.5). If one of these two motions is insignificant or indistinguishable from the other, in which case one of the order parameters, S_f^2 or S_s^2, is not statistically different from one, then model *m2* (Model 1.2) is a perfect approximation of the model-free motions. In addition, if the parameter τ_s is statistically zero then model *m1* (Model 1.1) is a perfect approximation. These two grids cover the entire model-free space of the five simplest models (Models 1.1–1.5). The more complex double timescale motions which include the τ_f and R_{ex} parameters, *m6*, *m7*, and *m8* (Models 1.6–1.8) have not been used to identify failed models, however, the conclusions apply to these models equally as well.

The increments of each model-free parameter dimension in the two three-dimensional grids are

$$S^2, S_f^2, S_s^2 = \{0.001, 0.388, 0.582, 0.698, 0.776, 0.831, 0.873, 0.905, 0.931, 0.952, 0.970\}, \quad (3.1)$$

$$\tau_e, \tau_s = \{0.1, 0.5, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192\}, \quad (3.2)$$

$$R_{ex} = \{0, 0.149, 0.223, 0.332, 0.495, 0.739, 1.102, 1.644, 2.453, 3.660, 5.460, 8.145, 12.151, 18.127, 27.043\}, \quad (3.3)$$

where the values of the correlation time parameters τ_e and τ_s are in picoseconds and the values of the chemical exchange parameter R_{ex} are measured in inverse seconds. The R_{ex} values are scaled quadratically with field strength and those given are for 600 MHz data. A total of four separate grids were used, the Perfect R_{ex} Grid with noise-free synthetic relaxation data, the Random R_{ex} Grid with randomised, noisy relaxation data, the Perfect Double Motion Grid with noise-free data, and the Random Double Motion Grid with randomised data. The total number of points in all four grids embodies 9,152 independent data sets.

The data created for these grids consisted of synthetic ^{15}N relaxation data. For each data set or grid point, six data points were generated, the ^{15}N R_1 and R_2, and the $\{^1H\}$-^{15}N steady state NOE, both at 600 and 500 MHz. By specifying the values of the model-free parameters and the specific model-free model used for each grid point, exact relaxation data was created by back-calculation. The data was then randomised using errors of 2% for the R_1 and R_2 and 0.04 for the 600 MHz NOE and 0.05 for the 500 MHz NOE. These errors were chosen to best mimick the white noise associated with experimental NMR data (d'Auvergne and Gooley, 2003).

The assumptions used in the calculation of the relaxation data include that the CSA for all cases is -160 ppm and that the bond length r is 1.02 Å. The major axes of the axially symmetric dipolar and chemical shift tensors were assumed to be co-linear and chemical exchange was scaled quadratically with field strength. All stated R_{ex} values correspond to the chemical exchange contribution to R_2 at 600 MHz data. The diffusion tensor was assumed isotropic and fixed to the value of 10 ns.

Model-free parameter constraints

The simple model-free models (Models 1.1–1.5) and the simple model-free models with a local τ_m parameter (Models 2.1–2.5) were minimised separately. During minimisation, the constraints

$$0 \leq S^2, S_f^2, S_s^2 \leq 1, \tag{4}$$

$$0 \leq \tau_e, \tau_s, \tag{5}$$

$$0 \leq R_{ex}, \tag{6}$$

$$0 \leq \tau_m \tag{7}$$

were applied. For the failed model-free models which possess a local τ_m parameter whereby its value at the minimum is infinite, the curvature of the model-free space can cause the number of iterations required to find the minimum to be unreasonably large. Therefore the constraint

$$\tau_m \leq 200 \text{ ns} \tag{8}$$

was applied to decrease computation time.

Model-free model elimination

For the implementation of model elimination within model-free data analysis, a series of empirical rules can be applied prior to model selection. These rules consist of inequalities containing model-free parameters together with arbitrarily chosen constants, ϵ_i, such that $\epsilon_i \in \mathbb{R}^+$. The ensemble of ϵ_i creates a set of elimination values whereby the elements are specified by the index i. If the inequality has been violated, the model is judged to have failed. By searching over all grid points by eye, only the correlation time parameters can be singled out as suitable for the discrimination of failed models. In all identified cases of model failure, either the internal correlation times, τ_e or τ_s, or the global rotational diffusion correlation time, local τ_m, will head towards infinity rather than minimising to a position in the model-free space close to the true value.

In situations where one of the model-free internal correlation times heads towards infinity, the inequality

$$\tau_e, \tau_f, \tau_s < \epsilon_i \tag{9}$$

could be used for elimination. However, to account for the anisotropy of macromolecules and the variability of the global correlation time in real systems, the inequality

$$\tau_e, \tau_f, \tau_s < \epsilon_i \cdot \tau_m \tag{10}$$

would be more flexible as the internal correlation times are compared directly with the global correlation time. As τ_m is the statistically limiting factor for the extraction of model-free motions from relaxation data, this inequality would be better suited for practical applications. A reasonable and generous value of ϵ_i accounting for significant anisotropy would be 1.5. This inequality is not only valid for models *m1* to *m8* (1.1–1.8) where a single τ_m exists as a global parameter for all residues but is also valid for models *tm1* to *tm8* (2.1–2.8) where an independent local τ_m parameter exists for each residue. In cases where τ_m tends towards infinity, the rule

$$\tau_m < \epsilon_i \tag{11}$$

can be used for elimination. Unlike in Inequality 10, no correlation time parameter can act as a reference for altering the limit in response to the variability of the system under study, therefore, a constant, arbitrary constraint must be used. A reasonable value of ϵ_i in most cases could be 50 ns. If the protein exhibits significant anisotropy and the bond vector is parallel to the long axis of the global rotational diffusion tensor, because the parameter τ_m is a local rather than global parameter in these models, the empirical value of ϵ_i may need to be increased accordingly. Also, if the protein is thought to tumble close to or slower than the fixed value of ϵ_i then its value should be increased as well.

During minimisation, linear constraints have been implemented by Inequalities 4–7 and, in certain cases, Inequality 8. These involve constraint values c_k where k is the constraint index. For example, as both upper and lower limits on τ_m have been set (Inequalities 7 and 8) the parameter is bound by

$$c_{k-1} \leq \tau_m \leq c_k,$$

where $c_{k-1} = 0$ and $c_k = 200$ s. If an upper constraint c_k has been set on τ_m, then the inequality

$$\tau_m < \epsilon_i < c_k \tag{12}$$

must be used for model elimination. The second inequality is necessarily exclusive due to the imperfections associated with machine precision during minimisation.

Data analysis and visualisation

After the constrained minimisation of the model-free model, the errors associated with the model-free parameters were determined by Monte Carlo simulations. This error analysis involved the constrained minimisation of five hundred independent simulations using exactly the same methods as the original model-free model. All model-free calculations were carried out using the in-house program 'relax'. The isosurface representations of the χ^2 values in the model-free parameter space were created using the program Open Visualisation Data Explorer 4.3.2, or OpenDX, from IBM (http://www.opendx.com). The projections of the Monte Carlo simulations in the four-dimensional model-free space of Figure 4 were composed using the program Grace (http://www.plasma-gate.weizmann.ac.il/Grace/), formerly called ACE/gr or Xmgr.

Results and discussion

Model-free model failure

By searching over the tens of thousands of model-free minimisations spread throughout the four separate grids (the Perfect R_{ex} Grid, the Random R_{ex} Grid, the Perfect Double Motion Grid, and the Random Double Motion Grid), numerous examples of failed model-free models are evident. These cases can be grouped into three categories, designated as Categories One, Two, and Three, dependent on how the model failed as well as which model-free parameters are present in the model. The first type of failure occurs in the standard model-free models (Models 1.1–1.8) and results in one of the internal correlation times, τ_e, τ_f, or τ_s, heading towards infinity (Table 2). The second category involves model-free models containing the local τ_m parameter (Models 2.1–2.8) and results in τ_m heading towards infinity (Table 5). The third category involves the same models as the second, where a local τ_m parameter is present, except that the internal correlation time heads towards infinity while the global correlation time, local τ_m, remains close to the true value (Table 7). For simplicity, only the first five model-free models where minimised. Four test cases (Tables 3, 4, 6, and 8) were selected from the many failed models to represent these three categories of failure, the randomised data of which is shown in Table 1. The first of the test cases involves perfect, noise-free data, hence, as the relaxation data and errors can be recalculated from the original model-free parameter values, is not included in the table. To understand the latent issues of model failure, these models will be discussed irrespective of whether they are selected by AIC model selection or not.

Category One failures

The first category of failure, whereby one of the internal correlation time parameters, τ_e, τ_f, or τ_s, heads towards infinity, is associated with model-free analysis using *m1* (Model 1.1) to *m8* (Model 1.8). These cases are distinguished by testing against Inequality 10 where ϵ_i is set to the value of 1.5. As the global correlation time of all grid points was fixed to 10 ns, failure was judged by the internal correlation time value being greater than or equal to 15 ns. A comprehensive summary of the Category One failures for model-free models *m1* to *m5* in each of the four grids is presented in Table 2. Within this table, two distinct subclasses of failure are evident. The first subclass of failure occurs in the models which correspond to those used in the creation of the grids. This is the most critical type of failure as, although less prominent prior to model selection, the majority of these failures appear post model selection. For the R_{ex} Grid the failed model is *m4* (Model 1.4), while for the Double Motion Grid the failed model is *m5* (Model 1.5). This subset of Category One failures occurs solely within the randomised grids, thus experimental noise is the genesis of failure. Potentially, as only synthetic data was tested, noise may not be the sole factor which induces these failures in real systems. The second subset of failures are most evident when the true underlying model experiences chemical exchange relaxation yet the fitted model is lacking the R_{ex} parameter. This R_{ex} compensatory subset of Category One failures occurs in both *m2* (Model 1.2) and *m5* (Model 1.5) in cases involving perfect, noise-free relaxation data and, although not statistically testable, the presence of noise may slightly increase the likelihood of failure. For individual data sets within the R_{ex} Grids, failure may occur solely in *m2*, solely in *m5*, or in both models. This subclass

Table 1. The full precision randomised ^{15}N relaxation data and errors used for the model-free minimisation, model-free model elimination, and model-free model selection of three of the test cases representing the three categories of model failure

Data type	Value	Error
Test Case 1		
600 MHz R_1	1.3607438126077531	0.028212431372334083
600 MHz R_2	15.285185187439794	0.29273912435875549
600 MHz NOE	0.83126690996568853	0.04
500 MHz R_1	1.8453066136464509	0.03603242845760548
500 MHz R_2	14.126806539264148	0.27274019888341894
500 MHz NOE	0.75788616726447011	0.05
Test Case 3		
600 MHz R_1	1.2779288762002963	0.026171608253937394
600 MHz R_2	18.601379356945866	0.38878447784580444
600 MHz NOE	0.82445156329111324	0.04
500 MHz R_1	1.7217246390621328	0.034023907561236126
500 MHz R_2	17.446689867478245	0.33972716512593082
500 MHz NOE	0.77090412852935197	0.05
Test Case 4		
600 MHz R_1	1.3645031921412598	0.027003467055169744
600 MHz R_2	13.502004134930948	0.27434770250574586
600 MHz NOE	0.83775052138736938	0.04
500 MHz R_1	1.7596749434196766	0.034823901208962289
500 MHz R_2	12.749980907526275	0.25906731578629066
500 MHz NOE	0.84114634537543753	0.05

is less concerning, however, as only a tiny fraction of these models are favoured over all others by AIC model selection.

A representative instance of the critical failures of models *m4* and *m5* in the Random R_{ex} and Double Motion Grids is displayed in Table 3. This example, which will be referred to as Test Case 1, is from the Random R_{ex} Grid. The original model-free parameter values used to generate the randomised ^{15}N relaxation data shown in Table 1 are $\{S^2 = 0.952,\ \tau_e = 2048$ ps, $R_{ex} = 0.739\ s^{-1}\}$. In this situation, both subclasses of failure are evident. Not only are the R_{ex} compensatory failures apparent in *m2* and *m5* but, crucially, model *m4* has failed. The χ^2 values and AIC criteria demonstrate that model *m4* is the most parsimonious of all the models and would be selected by AIC model selection. This is despite the internal correlation time value of 279 ns which, when compared with the global correlation time of 10 ns, is clearly a failure. Interestingly, the S^2 value is zero. None of the five models in the table have S^2 or internal correlation time values close to the true values. All models, excluding the failed model *m4*,

Table 2. The number of Category One failures prior to and post model selection[a]

Model[b]	Perfect RG[c]	Random RG[c]	Perfect DMG[c]	Random DMG[c]
m1	0 (0)	0 (0)	0 (0)	0 (0)
m2	190 (0)	210 (5)	0 (0)	1 (1)
m3	0 (0)	0 (0)	0 (0)	0 (0)
m4	0 (0)	91(70)	0 (0)	1 (0)
m5	336 (0)	372 (2)	0 (0)	78 (52)

[a]Category One failures are defined for models *m1* to *m8* and are when the internal correlation time parameter τ_e, τ_f, or τ_s heads towards infinity. Failure was judged as the internal correlation time being 1.5 times greater than the global correlation time, τ_m, of 10 ns. In brackets are the number of failed models chosen by AIC model selection. [b]The parameters of *m1* to *m5* are given in Models 1.1–1.5. [c]The Perfect grid is an ensemble of noise-free relaxation data while the Random grid consists of the same data which has been randomised to add noise. The R_{ex} Grids, composed of the three dimensions $\{S^2, \tau_e, R_{ex}\}$ corresponding to model *m4*, are an ensemble of 2640 relaxation data sets. The Double Motion Grids, composed of the three dimensions $\{S_f^2, S_s^2, \tau_s\}$ corresponding to model *m5*, are an ensemble of 1936 relaxation data sets.

Table 3. Test Case 1[a] model-free parameter values, χ^2 value, and AIC model selection criteria for the five minimised model free models

Model[b]	S^2	S_f^2	τ_e or τ_s (ps)	R_{ex} (s^{-1})	χ^2	AIC
m1	1.000 ± 0.006				20.14	22.14
m2	1.000 ± 0.418		$4.446e^7 \pm 2.246e^7$		20.14	24.14
m3	1.000 ± 0.008			0.831 ± 0.279	7.60	11.60
m4	0.000 ± 0.463		$2.791e^5 \pm 2.401e^5$	1.335 ± 0.320	3.27	9.27
m5	1.000 ± 0.340	1.000 ± 0.006	$3.964e^7 \pm 1.943e^7$		20.14	26.14
Truth	0.952		2048	0.739	13.68	

[a]The randomised 600 and 500 MHz ^{15}N relaxation data used for the fitting of these parameters are shown in Table 1. This case corresponds to the first category of failure whereby the internal correlation time parameter (τ_e, τ_f, or τ_s) shoots towards infinity.[b]The parameters of *m1* to *m5* are given in Models 1.1–1.5. The parameters of the model *Truth* correspond to the true, underlying values used in the creation of the randomised relaxation data.

have S^2 values of one and all internal correlation times are significantly overestimated in spite of the fact that these values are meaningless when taken together with an S^2 value of one. Clearly the best of the five models is *m3* as it is the only model which has a parameter value within errors to the true values with a R_{ex} value of 0.831 ± 0.279 where the true value is 0.739. The implementation of model elimination using Inequality 10 would remove models *m2*, *m4*, and *m5* and subsequent model selection using the AIC criterion would choose *m3*. With the given relaxation data and errors, this result is the most accurate and is best of all the models despite the inability to extract the motion at 2 ns.

To visualise exactly how model *m4* failed, the χ^2 values within three-dimensional model-free parameter space were mapped to a very fine resolution (Figure 1). The dimensions of this space are S^2, τ_e, and R_{ex} while the χ^2 values are represented as a fourth dimension, the curvature of the space, which is illustrated using isosurfaces contouring regions of equal value. The position in this space where the original, true parameter values lie is indicated by a black sphere. Normally noise would induce a slight shift from the true parameter values and cause the area surrounding the minimum to be perturbed but, in this test case, there appears to be no curvature in the space indicating the existence of a minimum proximal to the true values. The χ^2 value at the position of true parameters is 13.68 and minimisation initialised from this point will follow the four-dimensional valley delineated by the funnel shaped isosurfaces shown in Figure 1. This χ^2 minimisation along the non-linear route of the funnel from the initial location to the final minimum can be expressed as

$$\min_{\theta} \chi^2(\theta) \Leftrightarrow \begin{cases} S^2: & 0.952 \rightarrow 0, \\ \tau_e \text{ (ns)}: & 2.048 \rightarrow 279.145, \\ R_{ex}(s^{-1}): & 0.739 \rightarrow 1.335. \end{cases}$$

where $\theta = \{S^2, \tau_e, R_{ex}\}$ is the model-free parameter vector. Termination is due to the constraint $S^2 \geq 0$ preventing minimisation from following the funnel beyond the limit. If the constraints are lifted, a different minimum is found

$$\min_{\theta \in \mathbb{R}^3} \chi^2(\theta) \Rightarrow \chi^2(\hat{\theta}) = 3.260$$

where the minimised parameter vector of Test Case 1 is

$$\hat{\theta}, \qquad \text{where} \begin{cases} S^2 &= -\infty, \\ \tau_e &= \infty, \\ R_{ex} &= 1.35 \text{ s}^{-1}. \end{cases}$$

Rather than the minimum being completely lost, an alternative hypothesis might be that the space has been distorted by the noise to such an extent that the failure lies not in the model but in the minimisation. The level of zoom together with the one hundred data point per dimension resolution of Figure 1 may be insufficient for a sharp and discrete minimum to be visible. Figure 2 dispels this hypothesis by zooming into the region of the space surrounding the true parameter values and by comparing this with the equivalent region in the noise-free space. The top three images are orthogonal views of the zoomed in space of Figure 1 which, as Test Case 1 originated from the

Random R_{ex} Grid, was constructed from noisy relaxation data. In comparison, the bottom three images are equivalent orthogonal views of exactly the same grid point as Test Case 1 except that the relaxation data originates from the noise-free Perfect R_{ex} Grid. The sole difference between the top and bottom model-free spaces is that the top space is simply the bottom space with noise added. A clear minimum surrounds the true parameter values indicated by the sphere in the noise-free space, minimisation in which returns exact copies of the true values with a χ^2 value of zero. In contrast, this distinct minimum in the noise-free data has been completely destroyed by noise in the randomised space of Test Case 1. The perturbation to the space is so severe that, rather than the sphere being located next to a minimum, it is positioned close to the start of the four-dimensional valley, represented by the funnels, which heads towards infinity in Figure 1. In this test case, the minimum of the noise-free data has been squeezed out of existence by the noise.

These Category One, lost minimum failures occur purely because of the randomisation of the relaxation data. Unlike the Category One, R_{ex} compensation failures, AIC model selection will pick these failed models in the majority of cases. After model selection, 2.65% of the Random R_{ex} Grid and 2.74% of the Random Double Motion Grid consists of these failed models. Although the percentages are low, these values are significant due to the catastrophic nature of the failure. In addition, certain combinations of motions in macromolecules may induce much higher failure rates. Because of the low χ^2 value of the failed models, no model selection scheme is able to avoid these models. Therefore failures will surface within the final results of a standard model-free analysis unless the models are eliminated prior to model selection.

A select example representative of the R_{ex} compensatory subset of failures is shown in Table 4. In this case, which will be referred to as Test Case 2, failure occurs solely in model *m5*. This Category One R_{ex} compensation test case originates from the Perfect R_{ex} Grid where the original model-free model was *m4* (Model 1.4). The grid point corresponding to the test case has parameter values of $\{S^2 = 0.931,\ \tau_e = 256$ ps, $R_{ex} = 27.043\ s^{-1}\}$. Perfect ^{15}N relaxation data was

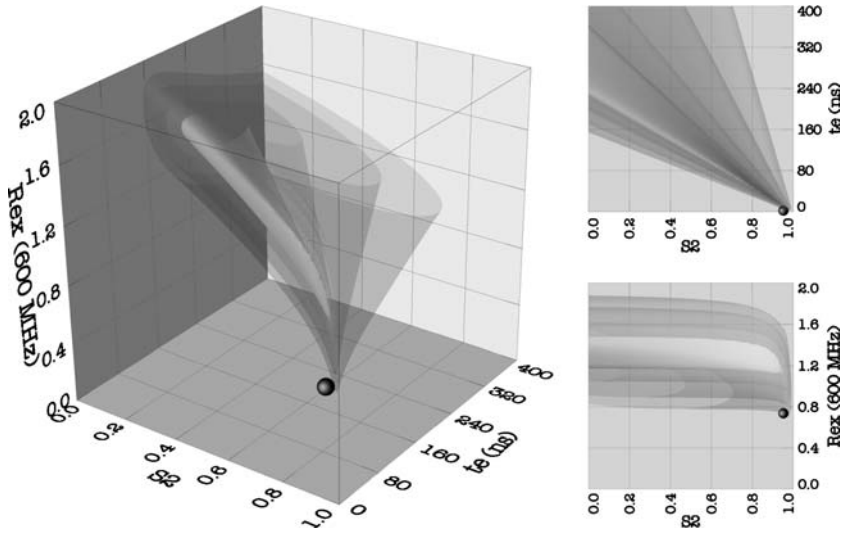

Figure 1. Test Case 1 isosurface representation of the three-dimensional model-free χ^2 space composed of the parameters $\{S^2, \tau_e, R_{ex}\}$. The true parameter values of $\{S^2 = 0.952, \tau_e = 2048$ ps, $R_{ex} = 0.739\ s^{-1}\}$ where used to generate randomised R_1, R_2, and NOE relaxation data at 600 and 500 MHz. These values are represented by the black sphere where the χ^2 value is 13.68. Due to the randomisation of the relaxation data and the application of constraints, the single minimum is located at $\{S^2 = 0, \tau_e = 279.145$ ns, $R_{ex} = 1.335\ s^{-1}\}$ where $\chi^2 = 3.27$. If the constraint of $S^2 \geq 0$ were to be lifted, τ_e would shoot to infinity. The χ^2 values of the four isosurfaces from outermost to innermost are 6, 5, 4, and 3.5 and the resolution of the space is one hundred data points per dimension.

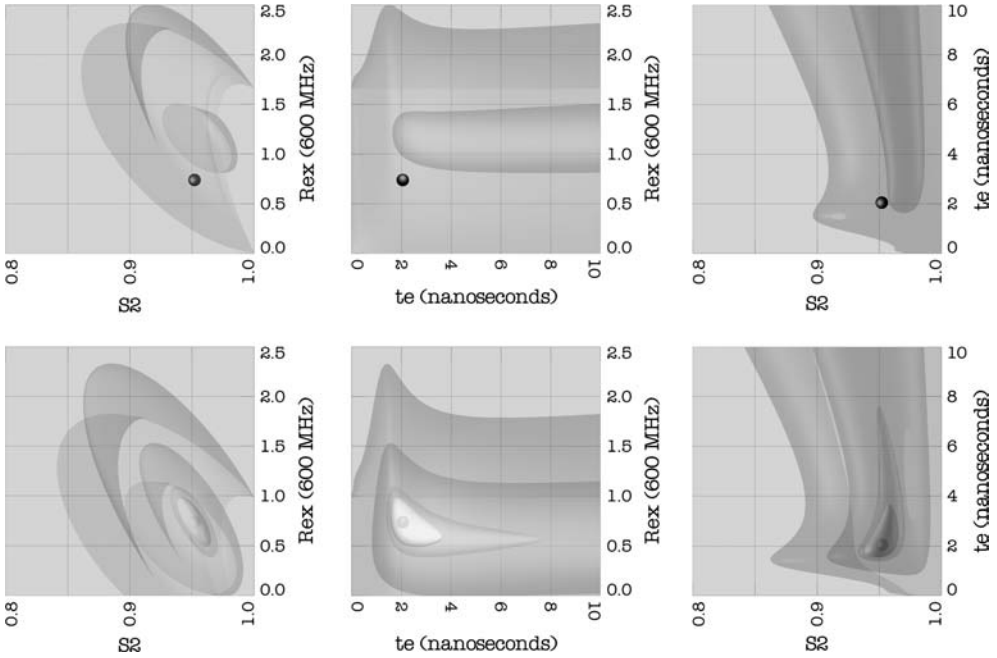

Figure 2. The lost minimum of Test Case 1. All images are of the three-dimensional $\{S^2, \tau_e, R_{ex}\}$ model-free χ^2 space. The true parameter values are $\{S^2 = 0.952, \tau_e = 2048$ ps, $R_{ex} = 0.739 s^{-1}\}$ and are represented by a black sphere. The top three images from left to right are the front view, side view, and top view of the χ^2 space of the randomised relaxation data from the Random R_{ex} Grid. These are simply a zoomed in view of Figure 1. In contrast, the bottom three images show the same angles of the χ^2 space but originate from the noise-free relaxation data of the Perfect R_{ex} Grid. The isosurfaces from outermost to innermost correspond to χ^2 values of 20, 5, 1, and 0.5 with the exception that the last two isosurfaces are not present in the randomised space. The minimum which is visible in the noise-free space has been obliterated in the randomised space. Instead of a minimum, the true parameter values are located next to the start of the funnel visible in Figure 1. The resolution of the space is one hundred data points per dimension.

used for the minimisation and Monte Carlo simulation error analysis, hence, the optimised parameters of *m4* returned exact copies of true parameter values with a final χ^2 value of zero. Although the S^2 value of models *m1*, *m2*, and *m3* is overestimated, the other parameter values are reasonably close to the true values. In contrast, the internal correlation time parameter τ_s in model *m5* has shot to the infeasible value of 83.67 μs. As this is four orders of magnitude greater than the global

Table 4. Test Case 2[a] model-free parameter values, χ^2 value, and AIC model selection criteria for the five minimised model free models

Model[b]	S^2	S_f^2	τ_e or τ_s (ps)	R_{ex} (s^{-1})	χ^2	AIC
m1	1.000 ± 0.007				1883.11	1885.11
m2	1.000 ± 0.133		$336.206 \pm 8.477e^4$		1883.11	1887.11
m3	0.978 ± 0.014			26.346 ± 0.643	34.56	38.56
m4	0.931 ± 0.012		256.000 ± 128.799	27.043 ± 0.633	0.00	6.00
m5	1.000 ± 0.398	1.000 ± 0.011	$8.367e^7 \pm 4.091e^7$		1883.11	1889.11
Truth	0.931		256	27.043	0.00	

[a]The relaxation data used in this case was perfect, noise-free data which can be back calculated from the true values. This case corresponds to the first category of failure whereby the internal correlation time parameter (τ_e, τ_f, or τ_s) shoots towards infinity. [b]The parameters of *m1* to *m5* are given in Models 1.1–1.5. The parameters of the model *Truth* correspond to the true, underlying values used in the creation of the perfect relaxation data.

correlation time of 10 ns, the model has clearly failed. Yet despite this failure, comparison of the χ^2 values of all models clearly demonstrates that, within this test case, the failed model *m5* will never be selected as the best model. The χ^2 value of *m4* is zero while the value of *m5* is 1883.11 thus, as both these models have the same number of parameters, model *m4* should always be chosen over *m5*

The two models, *m2* and *m5*, in which the Category One, R_{ex} compensatory subset of failures occurs (Table 2) are retaliating to the missing R_{ex} parameter by their internal correlation time shooting towards infinity. The combination of an R_{ex} contribution to the R_2 together with a model lacking the parameter results in an apparent increase of the spectral density function at zero frequency, $J(0)$, which would normally correspond to a slowing of correlation times. The combined effect of τ_m being fixed, the apparent increase in $J(0)$, and the invariance of high frequency spectral density values, $J(\omega)$ where $\omega \approx \omega_H$ means that, in certain cases, an increase in the internal correlation time parameter during minimisation can compensate. The result is a model which has failed, its fitted parameter values being clearly incorrect. Fortunately for this type of failure, AIC model selection very rarely selects these models with only five failed *m2* models and two failed *m5* models in the Random R_{ex} Grid of Table 2 having been selected which, as a percentage, is 0.265% of all data points. For the failed models which have not been selected, the internal correlation time is prevented from heading to higher values by S^2 hitting the one. Replacing the order parameter with the value of one in the original Lipari-Szabo model-free formulas (Lipari and Szabo, 1982a, b) results in a spectral density equation with no τ_e parameter, thus the value of the correlation time is meaningless. The final value reflects the curvature of the space prior to S^2 minimising to one. With S^2 at this position, τ_e can be changed to any value without affecting the minimised χ^2 value, hence the identical χ^2 values of models *m1*, *m2*, and *m5* in Table 4. In contrast, for the seven cases from Table 2 in which a failed model is selected, the S^2 value of these models is never one.

Category Two failures

The second category of failure occurs in spaces which contain both an internal correlation time parameter, τ_e, τ_f, or τ_s, together with a local τ_m parameter and is common in model-free analysis using models *tm1* (Model 2.1) to *tm8* (Model 2.8). As the original τ_m value for all data sets was set to 10 ns, Inequality 12 with the arbitrary time of $\epsilon_i = 50$ ns was used to judge failure. This is necessarily less than the constraint value, c_k, of 200 ns which was implemented during minimisation. Therefore any model with a local τ_m value greater than 50 ns is considered a failure. A summary of all Category Two failures is shown in Table 5. Comparison with the Category One failures in Table 2 reveals a very similar pattern of numbers between the two tables. The R_{ex} compensatory failures of Category One, which occur in models *m2* and *m5* in both the Perfect and Random R_{ex} Grids, also occur in Table 5. In contrast, the proportion of failure appears to be much higher when a local τ_m parameter is present in the model and, in addition, a number of these models are chosen by AIC model selection. Although less prominent than the first category, the lost minimum subclass of failures is also present within the table with model *tm4* failing in the Random R_{ex} Grid and model *tm5* failing in the Random Double Motion Grid. The proportion of these models being accepted by AIC model selection, though, is much less than that of the first category. A major

Table 5. The number of Category Two failures prior to and post model selection[a].

Model[b]	Perfect RG[c]	Random RG[c]	Perfect DMG[c]	Random DMG[c]
tm1	17 (0)	17 (0)	0 (0)	0 (0)
tm2	371 (1)	465 (30)	0 (0)	2 (1)
tm3	0 (0)	0 (0)	0 (0)	0 (0)
tm4	0 (0)	21 (9)	1 (0)	5 (0)
tm5	503 (0)	699 (9)	2 (0)	44 (2)

[a]Category Two failures are defined for models *tm1* to *tm8* and are when the local τ_m parameter heads towards infinity. Failure was judged as the local τ_m parameter being greater than 50 ns. In brackets are the number of failed models chosen by AIC model selection. [b]The parameters of *tm1* to *tm5* are given in Models 2.1–2.5. [c]The Perfect grid is an ensemble of noise-free relaxation data while the Random grid consists of the same data which has been randomised to add noise. The R_{ex} Grids, composed of the three dimensions $\{S^2, \tau_e, R_{ex}\}$ corresponding to model *m4*, are an ensemble of 2640 relaxation data sets. The Double Motion Grids, composed of the three dimensions $\{S_f^2, S_s^2, \tau_s\}$ corresponding to model *m5*, are an ensemble of 1936 relaxation data sets.

difference between Categories One and Two is that model *tm1*, which possesses a correlation time parameter, fails a number of times in both R_{ex} Grids.

A representative example of the R_{ex} compensatory subset of Category Two failures is shown in Table 6. This third test case is from the Random R_{ex} Grid, the randomised data of which was back calculated from model-free model *m4* with the parameter values of $\{S^2 = 0.97, \tau_e = 0.5$ ps, $R_{ex} = 5.46$ s$^{-1}\}$ and where the global correlation time was set to 10 ns. The randomised data and the errors used for Test Case 3 are shown in Table 1. Both models *tm2* (Model 2.2) and *tm5* (Model 2.5) have failed in this test case, as evidenced by their local τ_m parameter being stuck at the upper limit of 200 ns. In response, the τ_e or τ_s parameters have moved to the position of the global correlation time of 10 ns, while the S^2 parameter has shifted towards zero. Both models have identical parameter values, the additional parameter S_f^2 in *tm5* which has minimised to one collapses its model-free formula to that of *tm2* where τ_s and τ_e are one and the same parameter. Due to the low χ^2 values of both failed models, AIC model selection will choose *tm2* over all other models. Despite the S^2 values of one in the models which have not failed, their local τ_m values are close to the true value of 10 ns. For models *tm3* and *tm4*, τ_m is within errors to the true value. Additionally, both return estimates of R_{ex} almost within errors, the slight difference being because of the bias introduced by the S^2 value of one. Although not ideal, selection of one of these models will return reasonable parameter estimates. Model *tm4* would be the least desirable of the non-failed models as its internal correlation time is overestimated. Hence the best model for Test Case 3 would be *tm3* which can be extracted by using the combination of model elimination followed by AIC model selection.

The reason for model *tm2* failing within Test Case 3 is demonstrated in Figure 3. This diagram is the model-free χ^2 space which was mapped to a resolution of one hundred data points per dimension. The three dimensions are τ_m, S^2, and τ_e and the curvature of the space is represented by isosurfaces delineating regions of identical χ^2 values. The black sphere in the image is located at the true parameter values with the exemption of the R_{ex} parameter which does not exist within this space. The χ^2 value at the true parameter values is 300.17 and, from the curvature of the space, it is clear that there is no minimum close to this position. As in Test Case 1, the true values are located close to the start of a four-dimensional valley visible through the funnel shaped isosurfaces in the figure. Minimisation from this initial position will follow this highly curved valley and can be represented as

$$\min_{\theta} \chi^2(\theta) \Leftrightarrow \begin{cases} \tau_m \text{ (ns)}: & 10 & \rightarrow & 200, \\ S^2: & 0.970 & \rightarrow & 0.015, \\ \tau_e \text{ (ps)}: & 0.5 & \rightarrow & 10435. \end{cases}$$

where θ is the model-free parameter vector $\{\tau_m, S^2, \tau_e\}$. The minimised χ^2 value is 1.94. If the constraint of $\tau_m \leq 200$ ns is removed, a new minimum of

$$\min_{\theta \in \mathbb{R}^3} \chi^2(\theta) \Rightarrow \chi^2(\hat{\theta}) = 1.891$$

would be found where the fully minimised parameter vector is

$$\hat{\theta}, \quad \text{where} \begin{cases} \tau_m = \infty, \\ S^2 = 0, \\ \tau_e = 10.435 \text{ ns}. \end{cases}$$

The minimisation results associated with the R_{ex} compensatory subset of failures in Category Two are significantly different from those of Category One. This is due to the models in the second category possessing a local τ_m parameter whereas, in the first category, the global correlation time was fixed to 10 ns. Just as in the first category of failure, the R_2 values in both R_{ex} Grids are elevated due to the R_{ex} contribution to those relaxation data sets. For the models lacking the R_{ex} parameter, this elevation appears as an increase in $J(0)$ coupled with the other spectral density values, $J(\omega)$ where $\omega \neq 0$, not changing. As an increase in $J(0)$ is normally linked to an increase in one of the correlation times together with a compensatory decrease in the high frequency spectral density values, all the models lacking the R_{ex} parameter, *tm1*, *tm2*, and *tm5*, can compensate by increasing their local τ_m parameter. This is visible in Table 5 where, not only do models *tm2* and *tm5* fail more often when compared with Category One, but model *tm1* fails 17 times in both R_{ex} Grids as a compensation for the R_{ex} contribution to the R_2. In the first category of failure (Table 2), models *m2* and *m5* failed by allowing their internal correlation

Table 6. Test Case 3[a] model-free parameter values, χ^2 value, and AIC model selection criteria for the five minimised model free models.

Model[b]	S^2	S_f^2	τ_e or τ_s (ps)	R_{ex} (s^{-1})	Local τ_m[c] (ns)	χ^2	AIC
tm1	1.000 ± 0.007				11.784 ± 0.149	99.48	103.48
tm2	0.015 ± 0.268		10822.6 ± 1967.6		200.000 ± 88.205	1.94	7.94
tm3	1.000 ± 0.038			4.016 ± 1.267	10.471 ± 0.489	5.33	11.33
tm4	0.998 ± 0.104		459.2 ± 2331.7	4.045 ± 1.248	10.471 ± 8.636	5.30	13.30
tm5	0.015 ± 0.214	1.000 ± 0.083	10822.6 ± 1860.4		200.000 ± 75.686	1.94	9.94
Truth	0.97		0.5	5.46	10.0	300.17	9.94

[a]The randomised 600 and 500 MHz ^{15}N relaxation data used for the fitting of these parameters are shown in Table 1. This case corresponds to the second category of failure whereby the local τ_m parameter shoots towards infinity. [b]The parameters of *tm1* to *tm5* are given in Models 2.1–2.5. The parameters of the model *Truth* correspond to the true, underlying values used in the creation of the randomised relaxation data. [c]The local τ_m parameter was constrained within the limits $0 \leq \tau_m \leq 200$ ns.

time to shoot towards ∞ during minimisation, however, the final result was no different from model *m1* as their S^2 values minimised to one. In contrast, within Category Two models *tm1*, *tm2*, and *tm5* fail by allowing their local τ_m parameter to shoot towards ∞ and, in the majority of these failures, the internal correlation time minimises to the value of the global correlation time while the order parameter minimises to zero. This flexibility allows these models to minimise to lower χ^2 values and hence they sometimes filter through model selection.

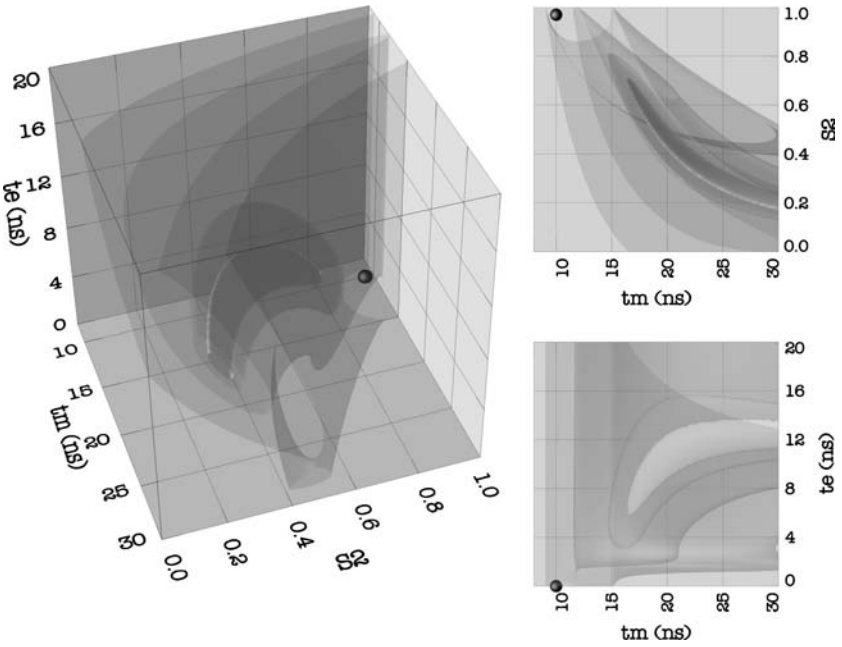

Figure 3. Test Case 3 isosurface representation of the three-dimensional model-free χ^2 space composed of the parameters $\{\tau_m, S^2, \tau_e\}$. The true parameter values of $\{\tau_m = 10$ ns, $S^2 = 0.97$, $\tau_e = 0.5$ ps, $R_{ex} = 5.46$ $s^{-1}\}$ where used to generate randomised R_1, R_2, and NOE relaxation data at 600 and 500 MHz. These values are represented by the black sphere where, without the R_{ex} parameter, the χ^2 value is 300.17. Due to randomisation and the fact that this space contains no R_{ex} dimension, the single minimum is located at infinity $\{\tau_m = \infty$ ns, $S^2 = 0$, $\tau_e = 10.435$ ns$\}$ where $\chi^2 = 1.89$. If an upper constraint of 200 ns is set for the parameter τ_m, the minimum shifts to $\{\tau_m = 200$ ns, $S^2 = 0.015$, $\tau_e = 10.822$ ns$\}$ where the minimum χ^2 value is 1.94. The χ^2 values of the four isosurfaces from outermost to innermost are 500, 100, 20, and 7 and the resolution of the space is one hundred data points per dimension. After minimisation and model selection, this is the model which is chosen to represent the dynamics of the system.

Category Three failures

By definition, the third category of failure occurs in model-free analysis using models *tm1* to *tm8*. The failure is defined as one of the internal correlation times heading towards ∞ while the local τ_m parameter stays at reasonable values. The satisfaction of Inequality 12 where ϵ_i is set to 50 ns together with the violation Inequality 10 where ϵ_i is set to 1.5 isolates the Category Three failures. As τ_m is a parameter of the model, the product $\epsilon_i \cdot \tau_m$ in Inequality 10 is flexible hence no fixed value of τ_e, τ_f, or τ_s will be eliminated. A summary of all Category Three failures is presented in Table 7. The total number of failures which fall under this third category is much lower than the first two categories and, importantly, almost none of these models are selected by AIC model selection. To understand failure, these models will be investigated nevertheless. Again the pattern within the table is very similar to the other two categories (Tables 2 and 5), with the distinctive R_{ex} compensatory failures in models *tm2* and *tm5* in both R_{ex} Grids, and the lost minimum failure of model *tm4* in the Random R_{ex} Grid and model *tm5* in the Random Double Motion Grid.

A representative example of the lost minimum failure in the Random Double Motion Grid is presented in Table 8 as the fourth test case. Here the original model was *m5* with parameter values of $\{S_f^2 = 0.970,\ S_s^2 = 0.970,\ \tau_s = 4096$ ps$\}$, thus $S^2 = 0.941$. As with the other lost minimum test cases, the true parameter values can be uncovered in the noise-free perfect grid with a χ^2 value of zero. Essentially, all models in the table return exactly the same results, all having identical χ^2 values. In models *tm2*, *tm3*, and *tm4*, the parameter τ_e is zero while R_{ex} is insignificant. Failure has occurred in model *tm5* where the parameter τ_s is 1.3 μs. Despite the failure being a lost minimum, in this case the symptoms of failure mirror those of the R_{ex} compensatory failures demonstrated in Test Cases 1 and 2 (Tables 3 and 4). The order parameter associated with τ_s is $S_s^2 = S^2/S_f^2 = 1$ hence τ_s, during minimisation, has moved towards infinity but can no longer move as its associated order parameter has hit one.

From the Category Three table (Table 7), another subclass of failure which is not present in the first two categories is evident. These appear as a failure of model *tm4* in the Perfect R_{ex} Grid and model *tm5* in the Perfect Double Motion Grid. Rather than these being a failure of the model-free model, they are in fact minimisation failures. An example is from the Perfect Double Motion Grid where the original model *m5* parameters are $\{S_f^2 = 0.388,\ S_s^2 = 0.001,\ \tau_s = 2048$ ps$\}$. Instead of finding the true minimum where the χ^2 value is zero, the minimised parameter values are $\{\tau_m = 1.73$ ns, $S_f^2 = 0.388,\ S_s^2 = 0.963,\ \tau_s = 3460.427$ ps$\}$ where the χ^2 value is 0.00132. This example was judged a failure because τ_s violated Inequality 10 where $\epsilon_i \cdot \tau_m = 2594$ ps. A simple test to differentiate minimisation versus model failure was to initialise minimisation from the true values, the result of which was only one iteration of optimisation returning the true values with a χ^2 value of $1.11e^{-27}$. The reason minimisation failed was because of the highly convoluted nature of the space induced by the unrealistically low S^2 value together with the grid search unable to find a position close to the true values. The Category Three failures are of no concern, however, as AIC model selection almost never picks these failures.

Table 7. The number of Category Three failures prior to and post model selection[a].

Model[b]	Perfect RG[c]	Random RG[c]	Perfect DMG[c]	Random DMG[c]
tm1	0 (0)	0 (0)	0 (0)	0 (0)
tm2	25 (1)	16 (0)	1 (0)	1 (0)
tm3	0 (0)	0 (0)	0 (0)	0 (0)
tm4	36 (0)	17 (0)	1 (0)	1 (0)
tm5	22 (0)	13 (0)	36 (0)	39 (0)

[a]Category Three failures are defined for models *tm1* to *tm8* and are when the internal correlation time parameter τ_e, τ_f, or τ_s heads towards infinity. Failure was judged as the internal correlation time being 1.5 times greater than the variable local τ_m parameter. In brackets are the number of failed models chosen by AIC model selection. [b]The parameters of *tm1* to *tm5* are given in Models 2.1–2.5. [c]The Perfect grid is an ensemble of noise-free relaxation data while the Random grid consists of the same data which has been randomised to add noise. The R_{ex} Grids, composed of the three dimensions $\{S^2, \tau_e, R_{ex}\}$ corresponding to model *m4*, are an ensemble of 2640 relaxation data sets. The Double Motion Grids, composed of the three dimensions $\{S_f^2, S_s^2, \tau_s\}$ corresponding to model *m5*, are an ensemble of 1936 relaxation data sets.

Model elimination

One possible solution to the failed model problem is to attempt to prevent the model from failing by

Table 8. Test Case 4[a] model-free parameter values, χ^2 value, and AIC model selection criteria for the five minimised model free models.

Model[b]	S^2	S_f^2	τ_e or τ_s (ps)	R_{ex} (s^{-1})	Local τ_m (ns)	χ^2	AIC
tm1	0.973 ± 0.010				9.592 ± 0.118	0.77	4.77
tm2	0.973 ± 0.010		0.000 ± 112.990		9.592 ± 0.114	0.77	6.77
tm3	0.972 ± 0.030			0.012 ± 0.871	9.588 ± 0.346	0.77	6.77
tm4	0.972 ± 0.026		0.000 ± 124.757	0.012 ± 0.786	9.588 ± 0.323	0.77	8.77
tm5	0.973 ± 0.380	0.973 ± 0.073	$1.387e^6 \pm 2.017e^9$		$9.592 \pm 1.076e^3$	0.77	8.77
Truth	0.941	0.970	4096		10.0	2.94	

[a]The randomised 600 and 500 MHz ^{15}N relaxation data used for the fitting of these parameters are shown in Table 1. This case corresponds to the third category of failure whereby the internal correlation time parameter (τ_e, τ_f, or τ_s) shoots towards infinity while the local τ_m parameter remains at reasonable values. [b]The parameters of *tm1* to *tm5* are given in Models 2.1–2.5. The parameters of the model *Truth* correspond to the true, underlying values used in the creation of the randomised relaxation data.

placing an upper constraint on the global and internal correlation times. When this strategy is implemented, the result is that the model still fails but minimisation cannot reach the unconstrained minimum. The correlation times will then be stuck at the upper limit and the χ^2 value will be elevated in comparison with the unconstrained value. The change in χ^2 value will, in some cases, solve the problem as the likelihood of model selection choosing one of the other models is increased. This approach is, nevertheless, non-selective and does not guarantee that the final model-free results will be free from failed models. Although constraints can be used to keep parameter values within reasonable bounds, this artificially creates the appearance of reliable parameter values. Yet these parameters will be far from the true values and therefore a failed constrained model with misleading parameter values will appear in the final results. The application of constraints can never cause a failed model to recover by having parameter values close to the true values, as it is the curvature of the model-free space itself which is intrinsically linked to model failure.

The constraint $\tau_e, \tau_f, \tau_s \leq 2\tau_m$, which is hard coded into the program relax, was applied during minimisation of Test Case 3. Although the limit has been dropped for Test Cases 1, 2, and 4 for the express purpose of investigating the reasons for model failure, this constraint is complementary to the model elimination inequality $\tau_e, \tau_f, \tau_s \leq 1.5\tau_m$. While the constraint can be used to aid in optimisation, the model elimination rule is still necessary for the removal of failed models. Hence both the constraint and model elimination should be used together when processing the relaxation data. The two disparate applications of inequalities are very much complementary to each other within model-free data analysis.

It should be noted that a failed model is very different from a model in which minimisation has failed. In the first instance the minimum is located at unreasonable parameter values whereas in the second the minimum has not been found, in which case constraints may help. Considering that only synthetic relaxation data was used to search for cases of failure, other categories of failed model-free models may exist when the true underlying dynamics is much more complex than models *m1* to *m8*. However, assuming that these failures exist, they cannot currently be identified and may not be as trivial to discern.

The only solution for preventing model failure is to collect better NMR spectra with less noise or to collect additional data. These measures will improve the curvature of the model-free space and possibly diminish the number of lost minimum subset of failures. Yet this is not guaranteed and the minimum may not be uncovered, hence, failed models may still be present in the final results. In the case of the R_{ex} compensatory subset of failures, no solution exists for their prevention as the failure is due to the model missing a R_{ex} parameter. As model failure in model-free analysis is inevitable, the addition of a model elimination step using Inequalities 10 and 12 between model-free minimisation and model-free model selection will, with certainty, weed out the failed models from the final model-free results.

Monte Carlo simulation failure

Model failure has been identified within model-free data analysis, manifesting itself as when one of the correlation times heads towards infinity. This problem surfaces not only after the initial minimisation of the model-free models but it can also plague error analysis by affecting individual Monte Carlo simulations. Fundamentally, each of the n Monte Carlo simulations is an independent model-free data set. They all germinate from the same synthetic relaxation data which is back calculated from the minimised model-free parameter values of the selected model. The segregation originates from the n discrete randomisations of the synthetic relaxation data which creates n distinct relaxation data sets. Each data set is minimised separately using the same methodology as for the real relaxation data. As the failure of models is intricately linked to the relaxation data itself, the n randomised simulations which mirror the data used in the randomised R_{ex} and Double Motion Grids are thus susceptible to the same failures. Only the lost minimum subset of failures surfaces within Monte Carlo simulations as the R_{ex} compensatory subset of failures are precluded by the nature of simulation construction. The elevated R_2 values are dealt with prior to error analysis by the model elimination and AIC model selection data analysis steps and, therefore, the final model used for the Monte Carlo simulations should include an R_{ex} parameter.

An extreme example illustrating the significance of simulation failure is if just one single simulation out of a million has a τ_e parameter which minimises to infinity, the standard deviation of τ_e will then be infinity. In reality, the standard deviation should be much lower – the single outlier causes a catastrophic failure of error analysis. A demonstration of the failure of Monte Carlo simulations is displayed in Figure 4. The original relaxation data used in this example, which will be called Test Case 5, is from the Perfect R_{ex} Grid where the original model is *m4*. The true model-free parameter values, which are found to machine precision by χ^2 minimisation, are $\{S^2 = 0.931$, $\tau_e = 2048$ ps, $R_{ex} = 1.644$ s$^{-1}\}$. These values are, for comparison, very similar to those of Test Case 1 (Table 3 and Figures 1, 2). Five thousand Monte Carlo simulations were minimised to generate the figure. While the fourth dimension, χ^2, of Test Case 1 could be mapped as a continuous surface in a three-dimensional space, the simulations are independent, each mapping to a different position within this space. Consequently, the simulations within the four-dimensional $\{S^2, \tau_e, R_{ex}, \chi^2\}$ model-free space are best represented visually by the six orthogonal two-dimensional projections of that space (Figure 4). Despite the completely different circumstances, another example portraying Monte Carlo simulations in this manner has been made by Andrec et al. (1999) however, in that case, no failures have occurred. To further enhance the visualisation of this space, the τ_e dimension is plotted on a log scale. In the figure, the top three projections can be used to reconstruct a three-dimensional space excluding the χ^2 values. Comparison with the complementary three bottom images of Figure 2, which is the perfect, noise-free χ^2 space of Test Case 1, unequivocally demonstrates that the positional probability of the Monte Carlo simulation is intrinsically coupled to the curvature of the χ^2 space of the original relaxation data. The contours of equal χ^2 value in Figure 2, represented by isosurfaces, clearly matches the density of the Monte Carlo simulations in Figure 4.

The same rules for model elimination can be used for culling failed Monte Carlo simulations. In Test Case 5 (Figure 4), Inequality 10 with ϵ_i set to 1.5 was used to part the simulations. As the global correlation time was 10 ns, any simulations with a τ_e value of greater than 15 ns was removed. These are represented in the figure as open circles whereas the accepted simulations are presented as crosses. The arbitrary nature of the line drawn through the model-free space used to segregate failure from non-failure is visible in the projections possessing a τ_e dimension. However, the projections with a S^2 dimension illustrate the accuracy of ϵ_i set to 1.5 – only open circles fall outside of the well defined S^2 distribution. Without eliminating the failed simulations, the model-free parameter errors are quite large (Table 9) with the τ_e error being considerably greater than the value of the parameter itself. After removal of the failed simulations using Inequality 10, the parameter errors are a much more reliable description of the system. The bias introduced into the error analysis by not removing these outliers is embodied in an order of magnitude overestimation of the S^2 and τ_e standard deviations (Table 9).

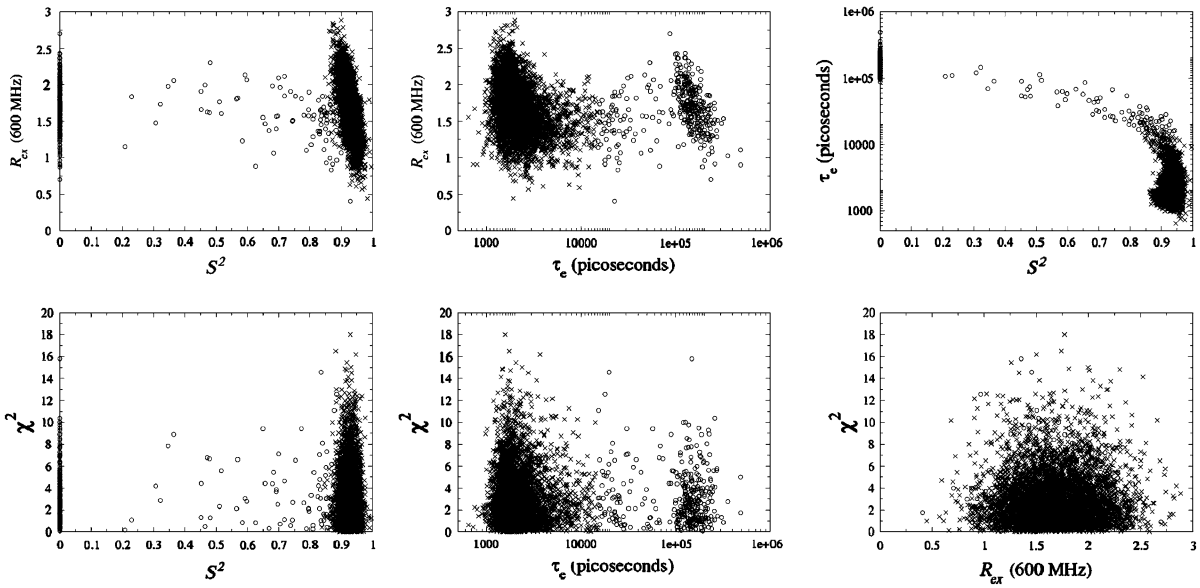

Figure 4. The results of five thousand Monte Carlo simulations represented by the six orthogonal projections of the four-dimensional space consisting of the dimensions S^2, τ_e, R_{ex}, and χ^2. The true parameter values used to back calculate the noise-free relaxation data of Test Case 5 are $\{S^2 = 0.931$, $\tau_e = 2048$ ps, $R_{ex} = 1.644$ s$^{-1}\}$. Minimisation of the noise-free data returned parameter values within machine precision to the true values with a chi-squared value of $3.21e^{-28}$. For the minimisation of the model-free parameters of both the relaxation data set and the simulations, no upper limit on τ_e was used. Each cross or circle corresponds to one Monte Carlo simulation. Using Inequality 10 with ϵ_i set to 1.5, as the global correlation time was set to 10 ns, simulations with a τ_e value greater than 15 ns were eliminated. These are represented by circles while all other simulations are represented by crosses. The empirical rule for elimination can be seen to be completely arbitrary. The projection of all simulations onto the χ^2 axis is the χ^2 distribution and, from these plots, it can be seen that the χ^2 values of the failed simulations are evenly distributed throughout this distribution.

The failure of Monte Carlo simulations has been previously identified and their removal attempted by the trimming of the upper and lower tails of the χ^2 distribution of simulations prior to calculating the standard deviation of the parameters. This process is implemented within the Modelfree program (Palmer et al., 1991; Mandel et al., 1995). In Figure 4, the one-dimensional projection of the three bottom images onto the χ^2 axis, each returning the same projection, is the distribution of χ^2 values for all 5000 simulations. The failed simulations are evenly dispersed throughout this distribution. When 10% of the simulations are culled by pruning the tails of the χ^2 distribution, the bias in the parameter errors are slightly increased in comparison with the errors calculated using all simulations (Table 9). Although an upper limit on τ_e may cause some of the χ^2 values of the failed simulations to be elevated, this does not guarantee that the χ^2 values of these simulations will be clustered at the upper

Table 9. The model-free parameter values of the Test Case 5 model *m4* together with the errors calculated from the five thousand Monte Carlo simulations.

Data set	S^2	τ_e (ps)	R_{ex} (s^{-1})
True values[a]	0.931	2048.00	1.644
Minimised values[b]	0.931	2048.00	1.644
Errors (all simulations)[c]	0.210	38403.47	0.342
Errors (elimination)[d]	0.019	1624.20	0.342
Errors (10% trimming)[e]	0.212	38785.36	0.341

[a]The true values of Test Case 5 used to back calculate the original noise-free relaxation data. [b]Minimisation returned the true values to within machine precision with a χ^2 value of zero. [c]All Monte Carlo simulations have been used in the calculation of parameter standard deviations. [d]Exclusion of the failed Monte Carlo simulations, judged by the τ_e parameter being greater than 15 ns, from the calculation of parameter standard deviations. [e]Exclusion of the upper and lower tails of the χ^2 distribution of Monte Carlo simulations from the calculation of parameter standard deviations.

tail of the χ^2 distribution. Therefore the selective elimination of failed simulations should be implemented over the non-selective trimming of simulations using the distribution of χ^2 values.

Conclusion

Model failure has been identified as a phenomenon which plagues model-free analysis and is manifested through one of the correlation time parameters heading towards infinity. Dependent on the set of model-free models employed as well as which correlation time parameter tends to infinity, these failures can be classified into one of three major categories. The first category of failure surfaces during standard model-free analysis using models *m1* (Model 1.1) to *m8* (Model 1.8) and is characterised by the internal correlation time, either τ_e, τ_f, or τ_s, shooting to infinity (τ_e, τ_f, or $\tau_s \rightarrow \infty$). The second and third categories both occur in model-free analysis using models *tm1* (Model 2.1) to *tm8* (Model 2.8) where the global diffusion tensor is replaced by the local τ_m parameter. The second category is defined as the local $\tau_m \rightarrow \infty$ whereas within the third category the local τ_m parameter remains close to the true value while τ_e, τ_f, or $\tau_s \rightarrow \infty$. These three categories can be further subdivided into the lost minimum subset of failures whereby the minimum present in the equivalent noise-free data has been distorted to such an extent that it no longer exists or the R_{ex} compensatory subset of failures whereby the correlation time is compensating for a missing R_{ex} parameter.

The lost minimum subdivision of failures have also been observed to affect individual Monte Carlo simulations. The n simulations are equivalent to n independent model-free minimisations. As these have been randomised using the experimental errors, the simulations are analogous to the randomised data of the R_{ex} and Double Motion Grids which were used to identify cases of model failure. If the small number of Monte Carlo simulation failures are not removed prior to calculating the model-free parameter standard deviations, these outliers cause a catastrophic overestimation of the errors.

Model-free data analysis is purely modelling incorporating the mathematical field of optimisation, specifically χ^2 minimisation (Palmer et al., 1991; Mandel et al., 1995; Orekhov et al., 1995; Fushman et al., 1997), and the statistical field of model selection to choose between a number of preset models (Mandel et al., 1995; d'Auvergne and Gooley, 2003; Chen et al., 2004). Between these two steps, the mathematical modelling concept of model validation should be implemented. The result of not removing the failed models prior to model selection is that the final results of a model-free analysis will likely include models with parameter values far from the truth. Although the alternative models selected after elimination may be missing certain parameters, the values of those which remain are reasonably close to the truth. The removal of failed model-free models and Monte Carlo simulations can be performed by testing the correlation time parameter against the linear Inequalities 10 and 12. These simple empirical rules will selectively eliminate the failed models and simulations. The use of either constraints or the trimming of the upper and lower tails of a χ^2 distribution does not solve the problem as these techniques are non-selective. Nevertheless, the constraints should be implemented to simplify optimisation. The sequence of data analysis steps which will yield the most accurate model-free results is, in order, model-free minimisation, model-free model elimination, and finally model-free model selection. In addition, elimination of failed Monte Carlo simulations is essential for the error analysis of the final model-free results.

References

Andrec, M., Montelione, G.T. and Levy, R.M. (1999) *J. Magn. Reson.*, **139**, 408–421.

Bellomo, N. and Preziosi, L. (1994) *Modelling Mathematical Methods and Scientific Computation, CRC Mathematical Modelling Series*, CRC Press, Boca Raton, FL.

Chen, J., Brooks, C.L. and Wright, P.E. (2004) *J. Biomol. NMR*, **29**, 243–257.

Clore, G.M., Szabo, A., Bax, A., Kay, L.E., Driscoll, P.C. and Gronenborn, A.M. (1990) *J. Am. Chem. Soc.*, **112**, 4989–4991.

d'Auvergne, E.J. and Gooley, P.R. (2003) *J. Biomol. NMR*, **25**, 25–39.

Fushman, D., Cahill, S. and Cowburn, D. (1997) *J. Mol. Biol.*, **266**, 173–194.

Korzhnev, D.M., Billeter, M., Arseniev, A.S. and Orehkov, V.Yu. (2001) *Prog. NMR Spectrosc.*, **38**, 197–266.

Lipari, G. and Szabo, A. (1982a) *J. Am. Chem. Soc.*, **104**, 4546–4559.

Lipari, G. and Szabo, A. (1982b) *J. Am. Chem. Soc.*, **104**, 4559–4570.
Mandel, A.M., Akke, M. and Palmer, A.G. (1995) *J. Mol. Biol.*, **246**, 144–163.
Orekhov, V.Yu., Korzhnev, D.M., Diercks, T., Kessler, H. and Arseniev, A.S. (1999) *J. Biomol. NMR*, **14**, 345–356.
Orekhov, V.Yu., Nolde, D.E, Golovanov, A.P., Korzhnev, D.M. and Arseniev, A.S. (1995) *Appl. Magn. Reson.*, **9**, 581–588.
Palmer, A.G., Rance, M. and Wright, P.E. (1991) *J. Am. Chem. Soc.*, **113**, 4371–4380.
Zhuravleva, A.V., Korzhnev, D.M., Kupce, E., Arseniev, A.S., Billeter, M. and Orekhov, V.Yu. (2004) *J. Mol. Biol.*, **342**, 1599–1611.