



HAL
open science

Calibration in Machine Learning Uncertainty Quantification: beyond consistency to target adaptivity

Pascal Pernot

► **To cite this version:**

Pascal Pernot. Calibration in Machine Learning Uncertainty Quantification: beyond consistency to target adaptivity. *APL Machine Learning*, 2023, 1 (4), pp.046121. 10.1063/5.0174943 . hal-04356591

HAL Id: hal-04356591

<https://hal.science/hal-04356591>

Submitted on 20 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Calibration in Machine Learning Uncertainty Quantification: beyond consistency to target adaptivity

Pascal PERNOT ¹

*Institut de Chimie Physique, UMR8000 CNRS,
Université Paris-Saclay, 91405 Orsay, France^{a)}*

Reliable uncertainty quantification (UQ) in machine learning (ML) regression tasks is becoming the focus of many studies in materials and chemical science. It is now well understood that *average* calibration is insufficient, and most studies implement additional methods testing the *conditional* calibration with respect to *uncertainty*, i.e. *consistency*. Consistency is assessed mostly by so-called reliability diagrams. There exists however another way beyond average calibration, which is *conditional* calibration with respect to *input features*, i.e. *adaptivity*. In practice, adaptivity is the main concern of the final users of a ML-UQ method, seeking for the reliability of predictions and uncertainties for any point in features space. This article aims to show that consistency and adaptivity are complementary validation targets, and that a good consistency does not imply a good adaptivity. An integrated validation framework is proposed and illustrated on a representative example.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/1.50174943

^{a)}Electronic mail: pascal.pernot@cnrs.fr

I. INTRODUCTION

The quest for trust or confidence in the predictions of data-based algorithms¹⁻⁴ has led to a profusion of uncertainty quantification (UQ) methods in machine learning (ML).⁵⁻¹⁹ However, not all of these UQ methods provide uncertainties that can be relied upon,^{20,21} notably if, as in metrology, one expects uncertainty to inform us on a range of plausible values for a predicted property.^{22,23}

In pre-ML computational chemistry, UQ metrics consisted essentially in *standard uncertainty*, i.e. the standard deviation of the distribution of plausible values (a *variance-based* metric), or *expanded uncertainty*, i.e. the half-range of a prediction interval, typically at the 95 % level (an *interval-based* metric).^{23,24} The advent of ML methods provided UQ metrics beyond this standard setup, for instance distances in feature or latent space^{13,25,26} or the Δ -metric,²⁷ which have no direct statistical or probabilistic meaning. These metrics might however be converted to variance-based metrics by *post hoc* recalibration methods such as temperature scaling²⁸⁻³⁰ and isotonic regression³¹, or to interval-based metrics by conformal inference.³²⁻³⁵ Nevertheless, all UQ metrics need to be validated to ensure that they are adapted to their intended use. In this study, I focus on the reliability of variance-based UQ metrics for the prediction of properties at the individual level.³⁶

The validation of UQ metrics is based on the concept of *calibration*. A handful of validation methods exist that explore more or less complementary aspects of calibration. A trio of methods seems to have recently taken the center stage: the *reliability diagram*^{28,30}, the *calibration curve*²⁹ and the *confidence curve*^{37,38}. They implement three different approaches to calibration which are not necessarily independent, but it is essential to realize that they do not cover the full spectrum of calibration requirements. In particular, none of these methods addresses the essential reliability of predicted uncertainties with respect to the input features, sometimes called *individual calibration*^{39,40}.

A. Scope and limitations of the study

The aim of this article is to propose a complete validation framework for variance-based UQ metrics, based on the concept of *conditional* calibration and its complementary aspects of *consistency* (conditional calibration with respect to uncertainty) and *adaptivity* (conditional calibration with respect to input features).

It is well known that average calibration is not sufficient to establish the reliability of ML-UQ predictions. This study goes one step further and is designed to alert ML users

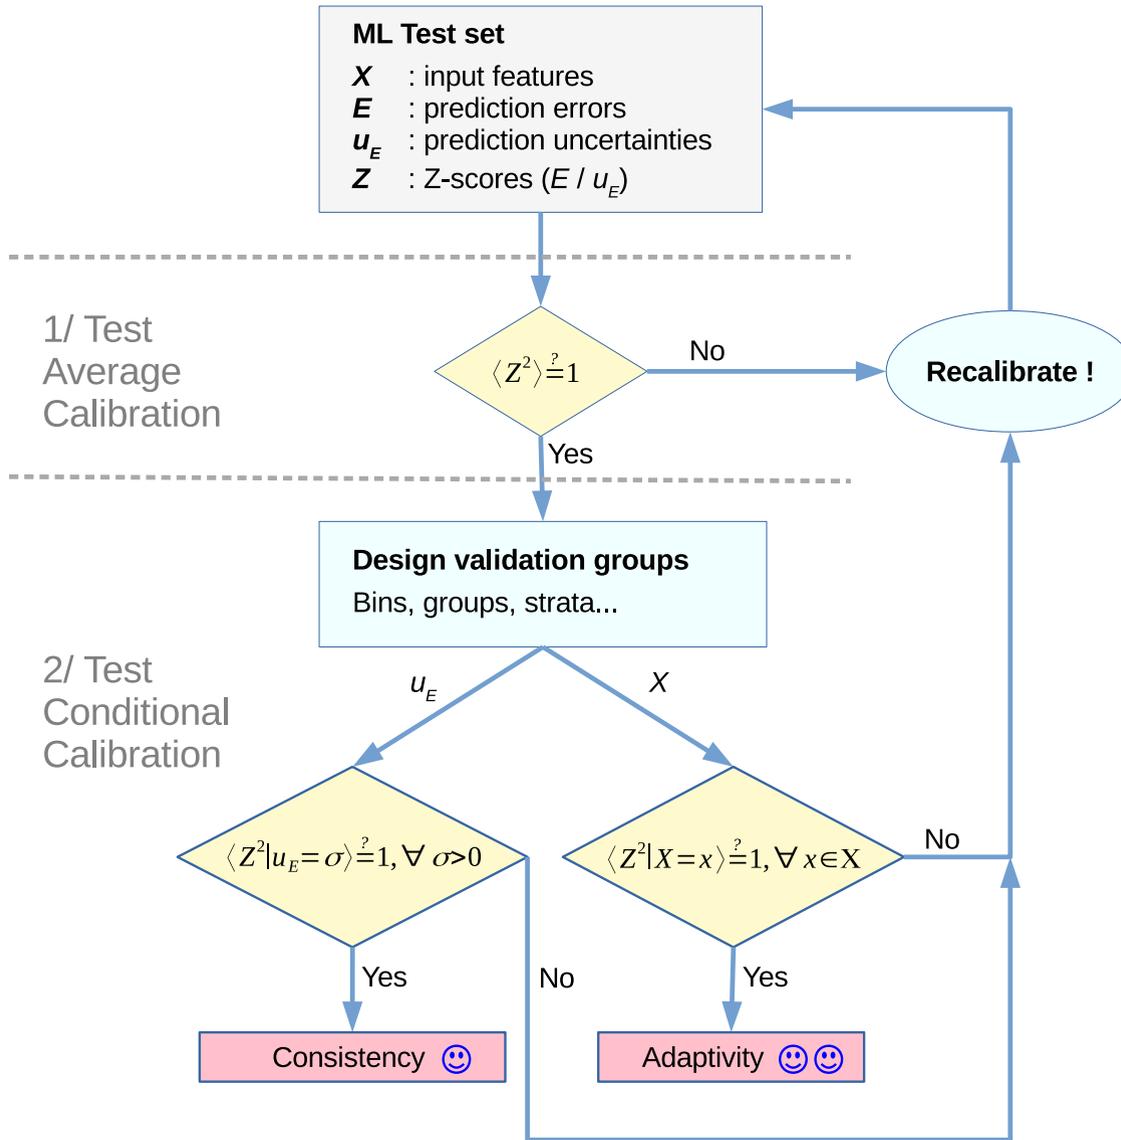


Figure 1. Flowchart of the z-scores-based validation framework.

that conditional calibration with respect to uncertainty (consistency), as commonly tested by reliability diagrams^{28,30,31,38,41–43}, does not guarantee individual calibration. In order to approach individual calibration, it is necessary to ensure conditional calibration with respect to input features (adaptivity). As reliability diagrams are not designed to deal with adaptivity, a more convenient statistical framework dealing homogeneously with consistency and adaptivity is proposed. The corresponding workflow is illustrated in Fig. 1.

Note that for the sake of brevity, I present here only methods for variance-based UQ metrics, but the approach can be directly transposed to interval-based metrics⁴⁴. Also, this

Notation	Definition
ACF	Auto-Correlation Function
AGC	Adversarial Group Calibration
CI	Confidence Interval (at the 95% level)
E	Prediction error
$f_{v,x}$	Fraction of valid intervals for conditioning variable x
MSE	Mean Squared Error: $\langle E^2 \rangle$
(L)RCE	(Local) Relative Calibration Error: $(RMV - RMSE)/RMV$
(L)ZM	(Local) Z-scores Mean: $\langle Z \rangle$
(L)ZMS	(Local) Z-scores Mean Squared: $\langle Z^2 \rangle$
RMSE	Root Mean Squared Error: \sqrt{RMSE}
RMV	Root Mean Variance: $\sqrt{\langle u_E^2 \rangle}$
u_E	Prediction uncertainty
X, X_i	Input feature
Z	Z-score: E/u_E

Table I. Main acronyms and notations used in this study.

study does not offer advice nor recipes on how to achieve good conditional calibration.

B. Structure of the article

The next section (Sect. II) introduces the notations and theoretical elements. The main notations and acronyms are summarized in Table I. The validation methods are presented in Sect. III and applied to a computational chemistry oriented example. The main conclusions are presented in Sect. IV.

II. VALIDATION OF VARIANCE-BASED UQ METRICS

The validation of variance-based UQ metrics requires at minimum a set of predicted values $V = \{V_i\}_{i=1}^M$, the corresponding uncertainties $u_V = \{u_V\}_{i=1}^M$, and reference data to compare with $R = \{R_i\}_{i=1}^M$ (with their uncertainties $u_R = \{u_{R_i}\}_{i=1}^M$, when relevant). From these, one estimates *prediction errors* $E = R - V$ and *prediction uncertainties* $u_E = (u_R^2 + u_V^2)^{1/2}$. These data enable to test average calibration (Sect. II A) and *consistency* as conditional calibration with respect to uncertainty (Sect. II B). An additional set of *input features* or adequate proxies

$X_j = \{X_{j,i}\}_{i=1}^M$ is required for a full validation setup including conditional calibration with respect to inputs or *adaptivity* (Sect. II B).

A. Average calibration

The validation of prediction uncertainty u_E can be based on the requirement that it correctly quantifies the *dispersion* of prediction errors E .^{21,45} Following the metrological definition of uncertainty, this is valid for so-called adequate models, i.e. models with negligible systematic or model errors. For models with non-negligible inadequacy levels, as can be expected, for instance, from ML methods below the interpolation threshold⁴⁶, prediction uncertainty is often designed to cover also for model errors.⁴⁷ In such cases, validation should account for both bias and dispersion components of the errors. I will not discuss here the problems of reporting such uncertainties for actionable predictions or risk assessment, but will mostly focus on a self-consistent setup for the validation of UQ calibration where one seeks a reliable estimation of the amplitude of errors.

In such conditions, the basis for validation is to require that the mean squared error (MSE) is close to the *mean variance* over the validation dataset⁴⁷

$$\langle E^2 \rangle \simeq \langle u_E^2 \rangle \quad (1)$$

However, this formula ignores the essential one-to-one pairing of errors and uncertainties, and a more stringent approach is based on z-scores ($Z = E/u_E$), using the condition

$$\langle Z^2 \rangle \simeq 1 \quad (2)$$

which is related to the Birge ratio⁴⁸ for the validation of the residuals of least-squares fit²¹.

Remarks.

- If E and u_E are obtained as the means and standard deviations of small ensembles of predictions (e.g. with less than 30 elements) these formulas have to be adapted, and hypotheses need then to be made on the error distributions for these small ensembles²¹. For a normal generative distribution of errors, the distribution of the mean of n values (ensemble size) is a Student's- t distribution with $\nu = n - 1$ degrees of freedom, and one should have $\langle Z^2 \rangle - \langle Z \rangle^2 \simeq \nu / (\nu - 2)$.²¹
- Unbiasedness is not an essential part of calibration, but it is a highly desirable property for predictions and z-scores and will be systematically considered as a test of prediction quality, i.e.

$$\langle Z \rangle \simeq 0 \quad (3)$$

- For unbiased z-scores, one recovers the variance test proposed previously by Pernot²¹

$$\text{Var}(Z) \simeq 1 \quad (4)$$

which can be superseded by Eq. 2.

The satisfaction of Eq. 2 validates *average calibration*, which is a necessary requirement, but does not guarantee the reliability of individual uncertainties, as average calibration can be satisfied by a compensation of under- and over-estimated values of u_E . The next section presents an approach to more local validation techniques.

▷ **Example.** In a recent article, Busk *et al.*³¹ extended a message passing neural network in order to predict properties of molecules and materials with a calibrated probabilistic predictive distribution. A *post hoc* isotonic regression on data unseen during the training was used to ensure calibration and was applied to materials science datasets. I consider here the QM9 test dataset, which consists of M=13 885 atomization energies (V, u_V), reference values (R), and molecular formulas.

Average calibration for this dataset is satisfactory, with $\langle Z^2 \rangle = 0.96(2) \simeq 1$, and z-scores are unbiased in average, with $\langle Z \rangle = 0.0082(83) \simeq 0$. □

3. Individual, conditional and local calibration

The best calibration one could ideally achieve is *individual* calibration, a condition where one is confident that uncertainty is correctly calibrated for any individual prediction. The formulation of individual calibration for probabilistic forecasters by Chung *et al.*³⁹, led them to formalize it as *conditional calibration in input features space*. In practice (i.e. for finite size datasets), individual calibration has been shown to be unreachable⁴⁰, and an alternative is to consider a discretized form as *local* or *group* calibration⁴⁹. This is reflected in the practical estimation of conditional statistics by data binning or grouping³⁹. In a similar spirit, *conditional coverage* with respect to input features was proposed by Vovk³² to assess the *adaptivity*³³ of conformal predictors.

For variance-based UQ metrics, Levi *et al.*^{30,42} proposed an approach based on *conditional calibration in uncertainty space*, namely

$$\langle E^2 | u_E = \sigma \rangle \simeq \sigma^2, \forall \sigma > 0 \quad (5)$$

which is the basis of the popular *reliability diagrams*²⁸ or *RMSE vs RMV* plots, also called *calibration diagrams*⁴¹, *error-based calibration plots*³⁸, *RvE* plots⁵⁰, or *RMSE vs. RMV* curves⁴³.

Levi *et al.* claim that, assuming that each uncertainty value occurs only once in the dataset, their method “captures the desired meaning of calibration, i.e., for each individual example, one can correctly predict the expected mistake”. In practice, the unicity assumption faces two major difficulties: (1) some datasets are *stratified*, with several occurrences of the same uncertainty value⁵¹, and (2) the practical implementation of conditional calibration requires to group data to estimate the mean squared error (MSE), breaking the one-to-one correspondence between the tested uncertainties and errors, as mentioned above for average calibration.

In consequence, conditional calibration based on Eq. 5 is not sufficient to validate calibration at the individual level. To go further, one should consider other conditioning variables besides u_E , notably input features or variables of interest for the end-user of a ML model, as proposed for probabilistic forecasters³⁹ and conformal predictors^{32,33}.

Building on the the works of Levi *et al.*⁴², Pernot²¹ and Angelopoulos *et al.*³³ about conditional calibration, I propose here to distinguish two calibration targets (besides average calibration), namely *consistency* as the conditional calibration with respect to prediction uncertainty, and *adaptivity* as the conditional calibration with respect to input features:

- **Consistency** is a special case of conditional calibration, in the sense that it involves only E and u_E . Using the z-scores statistics introduced for average calibration, one can define consistency by the following equation

$$\langle Z^2 | u_E = \sigma \rangle \simeq 1, \forall \sigma > 0 \quad (6)$$

Consistency is related to the metrological consistency of measurements⁵².

- **Adaptivity** is also conveniently formulated with z-scores as

$$\langle Z^2 | X = x \rangle \simeq 1, \forall x \in \mathcal{X} \quad (7)$$

where \mathcal{X} is the ensemble of values accessible to X . Adaptivity involves more information than consistency (X , E , and u_E).

Unless there is a monotonous transformation between u_E and X , consistency and adaptivity are distinct calibration targets, and a good consistency does not augur of a good adaptivity and vice-versa, so that both should be assessed. Note that *tightness*, as introduced earlier by Pernot²¹, covers both consistency and adaptivity.

Average calibration is a necessary condition to reach consistency or adaptivity. In fact, consistency/adaptivity expressed as conditional calibration should imply average calibration, but the splitting of the data into subsets makes that the power of individual consis-

tency/adaptivity tests is smaller than for the full validation set. It is therefore better to test average calibration separately, notably for small validation datasets.

Most methods used to this day for the validation of variance-based UQ metrics in chemical/materials sciences ML studies involve only E and u_E (reliability diagrams, calibration curves, confidence curves...) ³⁸. Adaptivity can thus be considered as a blind spot in UQ validation, despite its necessity to achieve reliable UQ at the molecule-specific level advocated by Reiher ³⁶.

III. VALIDATION METHODS

This section presents z-scores-based methods to assess and validate consistency and adaptivity. An alternative formulation, based on relative calibration errors, is also proposed in Sect. III B 4.

A. Homoscedasticity plots of z-scores

A simple way to estimate consistency is to plot the z-scores Z as a function of u_E ²¹. The dispersion of Z should be homogeneous along u_E (homoscedasticity) and, ideally, symmetric around $Z = 0$ (unbiasedness). In areas where the z-scores are biased, if any, one should observe a larger dispersion. This might not be easy to appreciate visually, and *running statistics* can be superimposed to the data cloud such as the *mean*, to be compared with $Z = 0$, and *mean squares* to be compared with the $Z = 1$ line.

This plot is easily extended to any variable X other than u_E and can be directly applied to the visual appreciation of adaptivity. Note that in the present context, the Z vs u_E plot is preferable to the E vs u_E plot used in other studies ^{13,21,25}, as it offers a consistent representation for both consistency and adaptivity estimation.

For cases where consistency/adaptivity cannot be frankly rejected on the basis of the shape or scale of this data cloud, it is necessary to perform more quantitative tests as presented below. One should not conclude on good consistency/adaptivity based solely on this kind of plot.

▷ **Example, continued.** The molecular mass (X_1 ; in Dalton (Da)) and fraction of heteroatoms (X_2 ; unitless) are generated from the molecular formulas of the QM9 dataset, and used as proxies for input features. They are practically uncorrelated between themselves and weakly

	X_2	$ E $	u_E
X_1	0.12	0.10	0.03
X_2	-	0.12	0.30
$ E $	-	-	0.32

Table II. Rank correlation coefficients for the $\{X_1, X_2, |E|, u_E\}$ set.

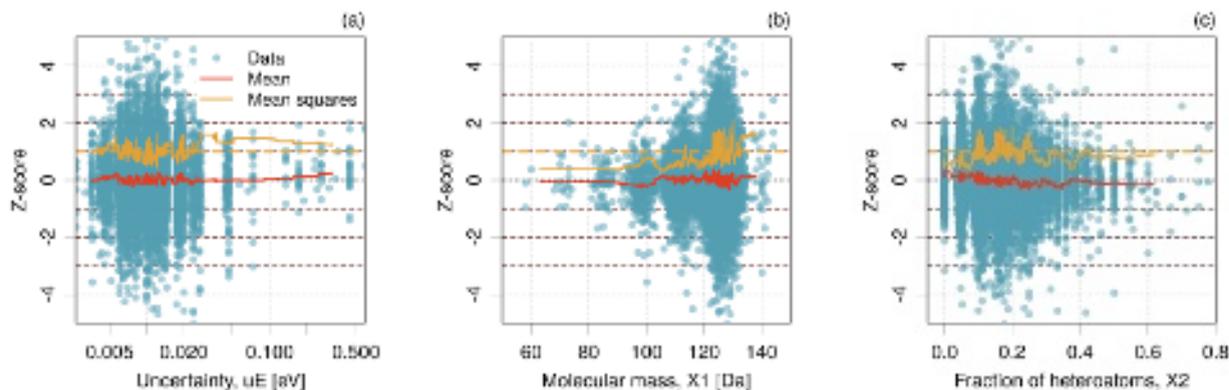


Figure 2. QM9 dataset: z-scores vs. uncertainty (a), molecular mass (b) and the fraction of heteroatoms (c). Running statistics (mean $\langle Z \rangle$ in red and mean squares $\langle Z^2 \rangle$ in orange) are estimated for a sliding window of size $M/100$.

correlated with $|E|$ and u_E (Table II). The dataset can thus be tested for consistency and adaptivity.

The homoscedasticity of z-scores for the QM9 dataset is estimated against u_E , X_1 and X_2 (Fig. 2). One sees in Fig. 2(a) that the data points are fairly symmetrically dispersed (mean; red line) and that the running mean squares (orange line) follows rather closely the $Z = 1$ line, up to $u_E \simeq 0.02$ eV, after which it lies at higher values. However, this concerns a small population (980 points) and the problem could be due to the data sparsity in this uncertainty range.

The “Z vs X_1 ” plot in Fig. 2(b) enables to check if calibration is homogeneous in molecular mass space. The running mean does not deviate notably from 0 (except around $X_1 \simeq 100$ Da, with a correlated increase in $\langle Z^2 \rangle$). The shape of the running mean squares line, erring towards small $\langle Z^2 \rangle$ values, indicates that uncertainties are probably overestimated for masses smaller than the main mass cluster (around 125-130 Da) evolving to a slight underestimation above this peak. This trend hints at a lack of adaptivity. A similar plot is shown for X_2 , the fraction of heteroatoms [Fig. 2(c)] where the running mean presents a weak but systematic trend from positive to negative values. Besides, the z-scores are under-dispersed for

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset. PLEASE DO NOT DISTRIBUTE OR IDENTIFY THIS ARTICLE AS: DOI: 10.1063/1.507494

molecules with low heteroatoms fractions (below 0.1), after which the running mean squares line presents notable oscillations around the $Z = 1$ reference line and seem to stabilize above $X_2 = 0.4$, where the data are sparse.

In this dataset, stratification of the conditioning variables is notable. For instance, the set of uncertainties u_E contains only 138 distinct numerical values, a fact which can be attributed to recalibration by a step-wise isotonic regression function. But stratification might also occur independently of any algorithm: X_1 contains 398 unique values, and X_2 is strongly stratified, with only 76 values. Stratification should be taken into account when binning these variables (see Sect. III B 3).

From these three plots, one gets the impression that calibration is rather good at the core of the dataset (where the density of data is highest), but more problematic in the margins. A more quantitative analysis of these features is desirable, but one might already conclude that adaptivity is not reached. \square

B. Local calibration

Conditional calibration in uncertainty space as formulated in Eq. 5 is often tested in the literature by reliability diagrams based on groups defined as uncertainty bins³⁰. This representation does not adapt conveniently to other grouping schemes. In contrast, the approach based on z-scores (Eqns. 6, 7), besides its interest evoked for average calibration, offers a uniform treatment for all conditioning variables and is used preferentially in this study. For readers more familiar with the use of calibration errors, an alternative formulation based on Local Relative Calibration Error is proposed in Sect. III B 4.

1. Local Z-Mean and Z-Mean-Squares analysis

Testing for consistency is based on a binning of the data according to increasing uncertainties. A Local Z Variance (LZV) analysis was introduced by Pernot⁴⁵ as a method to test local calibration: for each bin, one estimates $\text{Var}(Z)$ and compares it to 1. In the present framework, the LZV analysis is adapted to account for the possibility of accepting significant deviations of local $\langle Z \rangle$ values from 0, and one will be using the Local Z Mean Squares (LZMS) statistic, based on Eq. 6. A Local Z Mean statistic can also be used to check the local unbiasedness of z-scores.

Assessment of a LZMS analysis is based on two criteria, the deviation of the LZMS values from 1 and the homogeneity of their distribution along the conditioning variable.

The maximal admissible deviations depend on the bin size and error distribution. For instance, one should expect larger deviations from errors and uncertainties obtained as statistical summaries of small ensembles than from errors and uncertainties describing a normal distribution. Unless the error model is well known and controlled, which might not be the case for post-hoc calibration methods, it is impossible to define a threshold to LZMS values for validation purpose. It is therefore necessary to estimate confidence intervals (CIs) on the LZMS values to test their consistency with the target value.

For $\langle Z \rangle$, the standard formula based on the quantiles of the Student's- t distribution provides intervals with satisfactory coverage, even for small samples and non-normal distributions with finite variance. The case of $\langle Z^2 \rangle$ is more difficult, as standard formulas fare poorly when one deviates from the standard normality of the z -scores. The same problem was observed for $Var(Z)$,⁴⁵ and a convergence and power study concluded that the most reliable approach was to use bootstrapping^{53,54} with samples of at least 100 points. In such conditions, the effective coverage of 95 % CIs reaches at least 90 % for $Var(Z)$ and $\langle Z^2 \rangle$. To achieve a 95% coverage with a Student's distribution of z -scores, 1000 points per bin are required, which might limit the resolution of the local analysis. When in doubt, the LZMS analysis can be performed with several bin sizes to assess its reliability. When reliable CIs are obtained, the proportion of valid intervals, i.e. those covering the target value, can be used as a validation metric (Sect. III B 2).

The homogeneity of the distribution of the LZMS values along the conditioning variable can often be appreciated visually (any cluster showing systematic deviation from the target represents a local calibration problem). However, the graph might sometimes be crowded, and the auto-correlation function (ACF) of the LZMS statistics might help to detect the presence of unsuitable serial correlations.

2. Validation metrics

Calibration metrics^{1,55} are widely used in the ML-UQ literature: for instance, metrics have been designed for calibration curves^{8,29}, reliability diagrams⁴², and confidence curves³⁸. These metrics are generally used to compare and rank UQ methods, but they do not provide a validation setup accounting for the statistical fluctuations due to finite-sized datasets or bins numbers. It has been shown recently⁵⁶, that the *expected normalized calibration error* (ENCE)^{38,42,43} cannot be used directly as a validation metric. The calculation of reference values for those metrics is an option introduced recently^{57,58}, but being based on a proba-

bilistic model, it requires the choice of a probability distribution for the errors which might complicate the diagnostic.

Here, let us take advantage of the availability of confidence intervals on the local statistics of the LZM and LZMS methods as the basis for a validation metric. For a perfectly calibrated dataset, the fraction of binned statistics with a CI containing the target value should be close to the coverage probability of the CIs. Namely, about 95% of the binned $\langle Z^2 \rangle$ values should have 95% CIs containing the target value. Let us denote this fraction of validated intervals by $f_{v,ZMS}$. In practice, one should not expect to recover exactly 95%, and a CI for $f_{v,ZMS}$ has to be estimated from the binomial distribution to account for the limited number of bins⁴⁵.

3. Binning/grouping strategies

A sensitive point for the LZMS analysis is the choice of a binning scheme. The bin size should be small enough to get as close as possible to individual calibration testing and to provide information on the localization of any miscalibrated area, but also large enough to ensure a reasonable power for statistical estimation and testing of binned statistics. Moreover, stratification of the dataset, if present, has also to be considered.

a. Effect of equal-size binning for stratified conditioning variables. The equal-size binning scheme is a standard approach implemented for instance in reliability diagrams³⁰. One has to be aware that it does not account for the possible stratification of the conditioning variable. It was shown recently that bin-based statistics in such conditions are affected by the order of the data in the analyzed dataset.⁵¹ This effect is unavoidable and its impact on the statistics should be checked, for instance by repeated estimation of the binned statistics for randomly reordered datasets.

Note that getting a good estimate of $f_{v,ZMS}$ requires contradictory conditions, i.e. reliable confidence intervals for the binned statistics, therefore a number of points per bin as large as possible, but also a number of bins as large as possible. A good balance is obtained by choosing the number of bins as the square root of the dataset size ($N = M^{1/2}$). The use of this statistic should therefore preferably be reserved to large datasets, with more than 10^4 points.

b. Stratified binning. For notably stratified conditioning variables, a binning scheme preserving the strata might be more appropriate than equal-size binning, as it avoids the splitting of strata into arbitrary bins. However, many strata might have sizes too small to

enable reliable statistics. Instead of rejecting these low-counts strata, one can merge them with their neighbors. I use here an iterative algorithm where any small stratum (typically less than 100 points) is merged with the smallest of its neighbors. The value and counts of the strata are updated according to the relative counts of the merged strata. This simple merging is iterated until no small stratum is left. The result is not affected by data ordering. An inconvenience is that one does not have the control of the number of bins, which might get too low for a reliable estimation of the $f_{v,ZMS}$ validation statistic.

c. Choice of conditioning variables or groups for adaptivity. Although using one or several input features as conditioning variable is the most direct way to test adaptivity, it might not always be practical, for instance when input features are strings, graphs or images. In such cases, one might use dimension reduction algorithms such as t-SNE^{25,59} or UMAP⁶⁰ in order to define relevant groups. One might also use proxy variables, latent variables, or even the predicted property value V . Using V answers to the question: *are uncertainties reliable over the full range of predictions?* A problem with V is that it is potentially strongly correlated with E which might lead to spurious features in the LZMS analysis. For the complementarity of consistency and adaptivity tests, it is better to use X variables that are not strongly correlated to E and u_E . If there is no sensible way to define a conditioning variable, one might consider adversarial group validation.

d. Adversarial groups. One can avoid to choose conditioning variables by designing random groups to be tested for calibration, in the spirit of *adversarial group calibration* (AGC). In AGC, the largest calibration error is estimated over a set of random samples of a given size, for sizes varying from a small fraction of the dataset to the full dataset.^{39,40,61} This approach is mostly used to compare datasets, but does not provide a validation setup. As exposed above, even for fully calibrated datasets, the amplitude of calibration errors depends also on the group size and error distribution, which makes comparisons difficult for datasets with unknown or different distributions.

It is possible to design an AG *validation* method based of the $f_{v,ZMS}$ metric. Preliminary test of this approach revealed three main limitations that make it unpractical: (1) if a dataset presents localized calibration issues, there is a low probability to randomly sample groups revealing this problem, and one will get an overly optimistic diagnostic; (2) random groups are not interpretable, and there is therefore very few to learn about the origins of miscalibration; and (3) the computer time for the repeated estimation of converged bootstrap-based CIs might be prohibitive considering the small information return.

Another option for validation is to use a conventional AGC curve and build a probabilis-

tic reference AGC curve, as suggested previously for confidence curves⁵⁷. In contrast with the latter case, this probabilistic AGC reference curve is very sensitive to the choice of a probability distribution for the generated errors, which might lead to ambiguous diagnostics.

Because of these difficulties, AG validation has not been retained for the present study. More generally, it has nevertheless to be kept as an option when the design of adequate conditioning variables is problematic. Further research to design a robust AGC reference curve is needed.

▷ **Example, continued.** The unbiasedness and consistency of the QM9 validation set are tested by performing a LZM/LZMS analysis in u_E space with 100 equi-sized bins. The corresponding f_v validation statistics are presented in Fig. 4.

One sees for $\langle Z \rangle$ [Fig. 3(a)] that there are no outstanding deviations of the binned statistic, except for the rightmost point (in red), in agreement with the observation on the previous z-scores plots [Fig. 2(a)]. The fraction of valid intervals (in blue) is high ($f_{v,ZM} = 0.97$) and in statistical agreement with its target value of 0.95. For $\langle Z^2 \rangle$ [Fig. 3(d)], one observes slightly more deviant bins, with $f_{v,ZMS} = 0.86$, but no major trend, as confirmed by the ACF analysis [Fig. 3(g)]. Consistency is therefore not perfect but without strong local miscalibration. As the uncertainties are stratified, one should evaluate the impact of data ordering on these statistics.

For this, $f_{v,ZM}$ and $f_{v,ZMS}$ have been simulated for 1000 random samplings of the data order and summarized by their mean value and 95 % confidence intervals [red square in Fig. 4(a)]. To account for the finite number of bins, each value has also been perturbed by binomial noise [orange diamonds in Fig. 4(a)]. In the present case, the dispersion due to data reordering is small when compared to the binomial uncertainty for both statistics. The data reordering uncertainty is not sufficient to get an overlap of the confidence interval for $f_{v,ZMS}$ with the target coverage (0.95), validating the conclusions of the nominal LZMS analysis.

Adaptivity is tested using the same protocol. For the molecular mass [Fig. 3(b,e,g)], the fraction of biased bins reaches 12 % ($f_{v,ZM} = 0.88$) and the fraction of deviant bins for $\langle Z^2 \rangle$ is about 40 % ($f_{v,ZMS} = 0.6$). There is a strong predominance of deviant bins for the masses below 120 Da, where the small values of the statistic point to overestimated uncertainties. In consequence, the ACF of the LZMS series presents a slow decay, to be compared with the one obtained for u_E . If one accepts that there is no strong bias of Z in this area, values of $\langle Z^2 \rangle$ around 0.5 can be interpreted as an excess factor of $1/\sqrt{0.5} \simeq 1.4$ for the uncertainties. Here again, despite the notable stratification of the molecular masses, the ordering of the data has not a strong impact on the f_v statistics [Fig. 4(b)].

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset. PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0174943

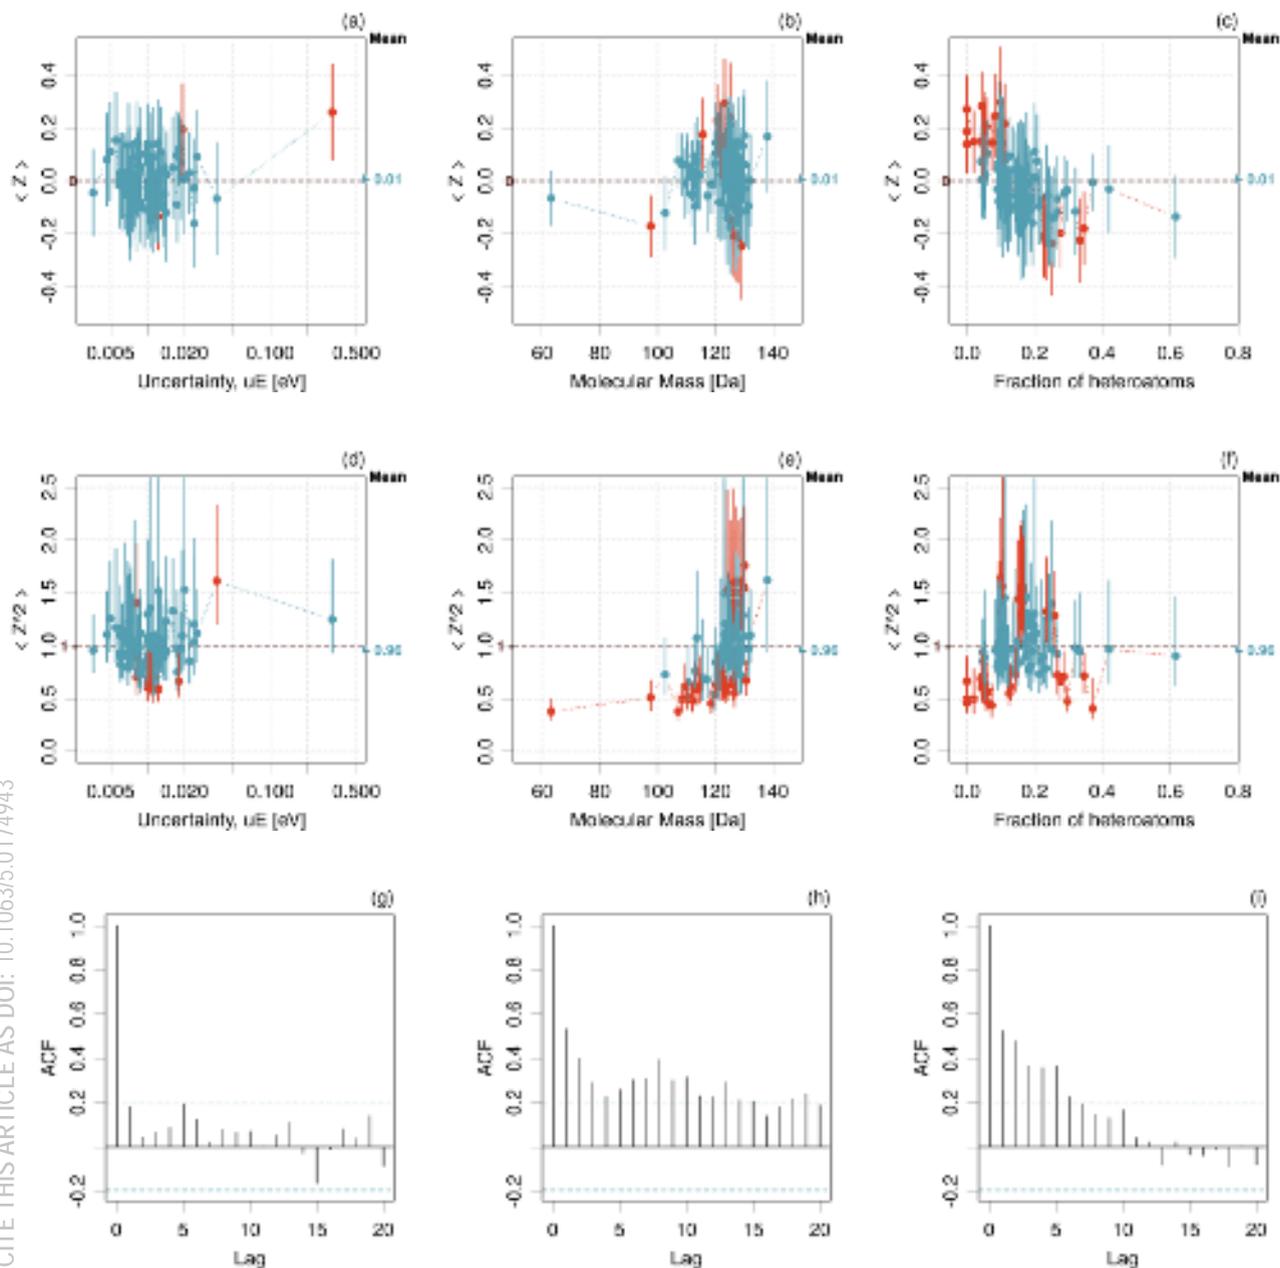


Figure 3. QM9 validation dataset. Consistency and adaptivity validation plots based on 100 equal-size bins: LZM and LZMS analyses and ACF of LZMS vs. u_E (a, d, g), molecular mass X_1 (b, e, h) and fraction of heteroatoms X_2 (c, f, i). For the LZM and LZMS analyses (a-f), the red symbols depict confidence intervals that do not contain the target statistic (0.0 for $\langle Z \rangle$; 1.0 for $\langle Z^2 \rangle$), and the mean statistic (for the whole dataset) is reported in the right margin, with the same color code as for the local statistics. The corresponding f_v statistics are reported in Fig. 4.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset. PLEASE DO NOT DISTRIBUTE THIS ARTICLE. DOI: 10.1063/1.5174183

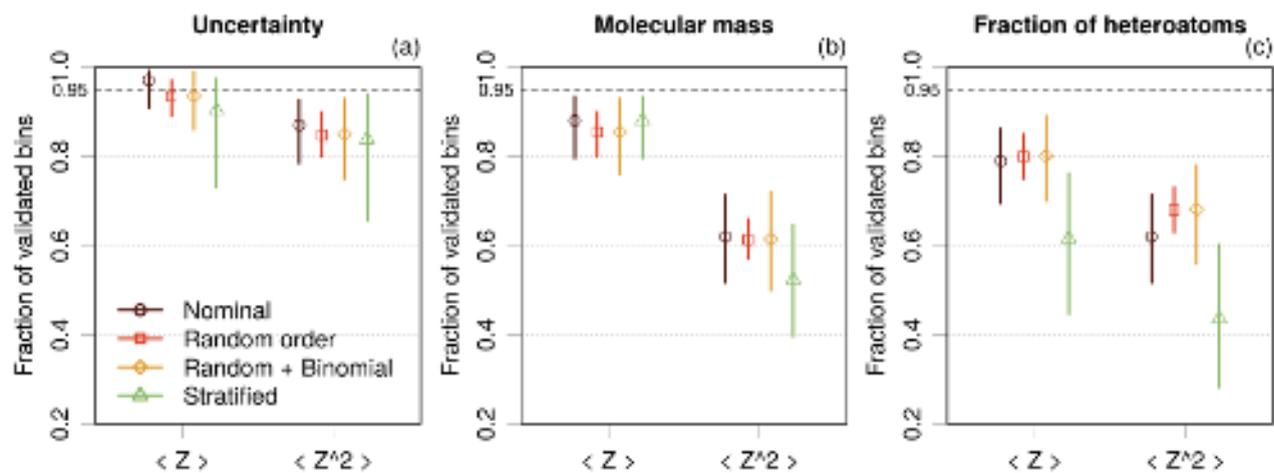


Figure 4. Fraction of validated bins for $\langle Z \rangle$ (left) and $\langle Z^2 \rangle$ (right) according to three conditioning variables (a-c). The error bars depict 95 % confidence intervals. The fractions should ideally be compatible with the 0.95 target (horizontal dashed line). The “Nominal” values (black circles) result from equal-sized binning with 100 bins of the dataset and the error bars are estimated from a binomial distribution. They summarize the LZM and LZMS analyses reported in Fig. 3. The “Random order” values (red squares) display the mean and 95 % confidence interval for a random ordering of the dataset (based on 1000 permutations). The “Random + Binomial” values (orange diamonds) combine the binomial uncertainty with the previous values. The “Stratified” values (green triangles) are the statistics for the binning scheme based on the preservation of strata, with binomial uncertainty. They summarize the LZM and LZMS analyses reported in Fig. 5.

For the fraction of heteroatoms [Fig. 3(c,f,i)], the LZM analysis displays the same trend from positive to negative $\langle Z \rangle$ values as observed in Fig. 2(c), with a sub-optimal fraction of valid bins ($f_{v,ZM} = 0.80$). The LZMS analysis reveals clusters of deviant bins at several spots along the X_2 axis, which is reflected in a slowly decreasing ACF. The fraction of valid bins is small ($f_{v,ZMS} = 0.62$). Here again, despite the strong stratification of X_2 , the $f_{v,ZMS}$ statistic is not strongly affected by the reordering perturbation [Fig. 4(c)].

The LZM/LZMS analysis based on stratified binning with a minimum of 100 points per bin is presented in Fig. 5, and the f_v values are also reported in Fig. 4. This representation is less crowded than the equal-size binning and provides essentially the same conclusions. One notes more severe values of the f_v statistic for the adaptivity analysis, with larger error bars due to a smaller number of bins.

All diagnostics based on E and u_E conclude therefore to a good calibration and an acceptable consistency. The main feature revealed by this analysis is the lack of adaptivity seen

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE DO NOT REMOVE THIS INFORMATION FROM THE MANUSCRIPT. DOI: 10.6028/AIP.MPL.A.1174943

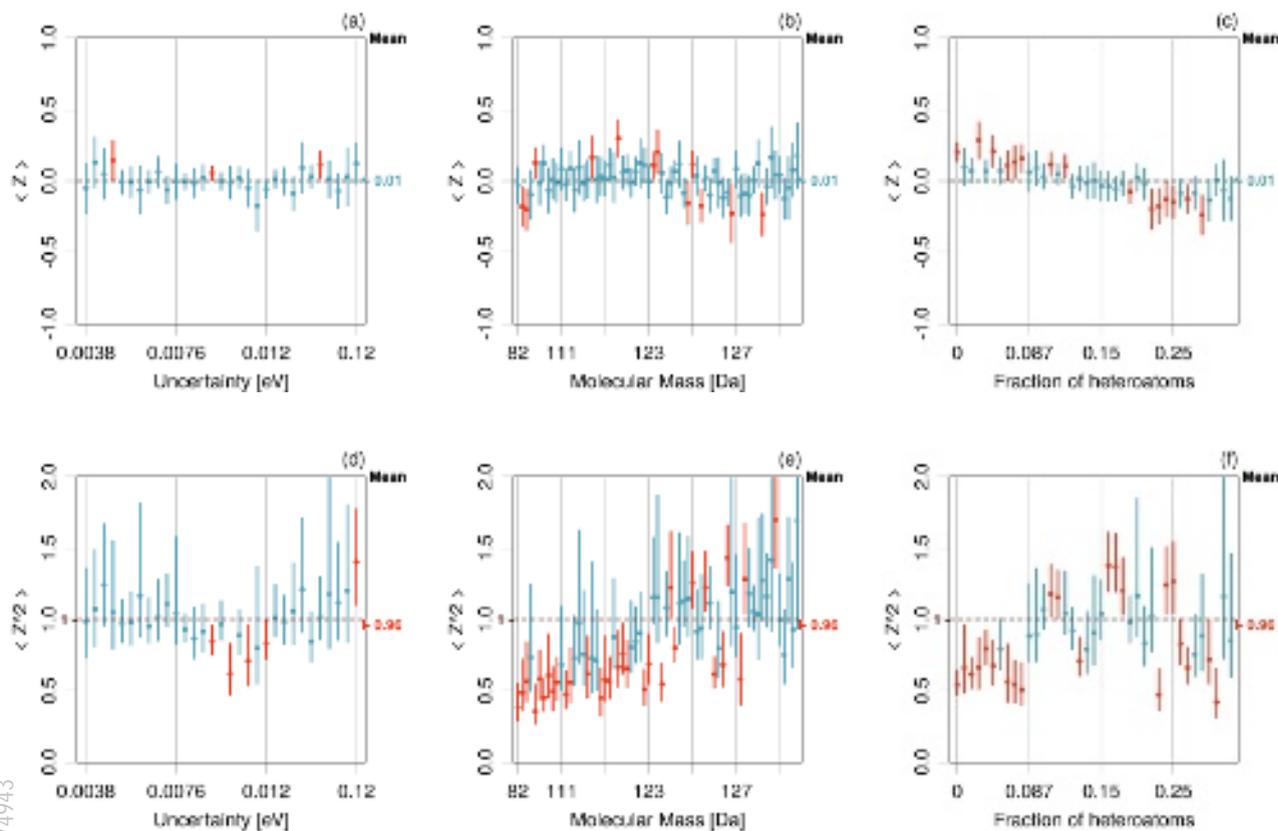


Figure 5. QM9 validation dataset. LZM, LZMS analyses vs. u_E (a, d), molecular mass (b, e) and fraction of heteroatoms (c, f). The data have been aggregated to get a minimum of 100 points per stratum. The red symbols depict confidence intervals that do not contain the target statistic (0.0 for $\langle Z \rangle$; 1.0 for $\langle Z^2 \rangle$). The mean statistic (over the whole dataset) is reported in the right margin, with the same color code as for the local statistics. The corresponding f_v statistics are reported in Fig. 4.

by the LZMS analysis for both molecular mass and fraction of heteroatoms. A major trend is a significant underestimation of the quality of predictions for the lighter molecules in the QM9 dataset (below 120 Da) and also for those with a small fraction of heteroatoms (below 0.1). □

4. Alternative approach: the Local Relative Calibration Error

Deriving from the logic behind reliability diagrams, a popular measure to assess the error in calibration is the Expected Normalized Calibration Error (ENCE)³⁰, which averages the absolute relative calibration errors over the bins

$$ENCE = \frac{1}{N} \sum_{i=1}^N |RCE_i| \quad (8)$$

where N is the number of bins, RCE_i is the Relative Calibration Error in bin i

$$RCE_i = \frac{RMV_i - RMSE_i}{RMV_i} \quad (9)$$

$RMSE_i$ is the root mean squared error for bin i , and RMV_i is the root mean variance (u_E^2) in bin i .

In the usual applications of ENCE, the bins are based on uncertainty, so that the ENCE is a measure of consistency. However, Eq. 8 is valid for any binning scheme, and therefore the ENCE can also be used as an adaptivity measure.

In this context, the binned RCE offers an alternative to the z-scores formulation and can be used to establish conditional equations similar to Eqns. 6-7

$$(RCE|u_E = \sigma) \simeq 0, \forall \sigma > 0 \quad (10)$$

and

$$(RCE|X = x) \simeq 0, \forall x \in \mathcal{X} \quad (11)$$

This defines the Local RCE (LRCE) analysis that can be implemented through data binning according to any conditioning variable, as for the LZMS analysis.

This formulation could be more appealing to users familiar with the ENCE, despite the underlying problem mentioned for average calibration that the $RMSE$ and RMV values are insensitive to the pairing of errors and uncertainties. The ZMS approach is therefore more robust. However, the most important goal at the present stage of ML-UQ development is for practitioners to assess adaptivity, be it by LRCE or LZMS. Besides, a LRCE analysis could be more consistent with existing ENCE-based toolboxes, such as the Uncertainty Toolbox⁶¹.

IV. CONCLUSIONS

The concept of conditional calibration enables to define two aspects of local calibration: consistency, which assesses the reliability of UQ metrics across the range of uncertainty values, and adaptivity, which assesses the reliability of UQ metrics across the range of input features. As the validation of individual calibration is practically impossible, one has to rely on validation methods based on local or group calibration, making consistency and adaptivity complementary validation targets.

Consistency and adaptivity can be tested by binned statistics such as the mean of squared z-scores $\langle Z^2 \rangle$, leading to the LZMS analysis. Bins with large deviations from the target

value (typically 1) and groups of adjacent bins with similar deviations reveal local calibration errors. The LZMS analysis enables to test conditional calibration for any conditioning variable, giving access to both consistency and adaptivity validation. An alternative formulation based on local relative calibration errors (LRCE) could also be considered. A validation metric f_v based on the proportion of bins with the confidence interval of a statistic containing its target value was proposed. The focus of this study is on variance-based UQ metrics, but this validation framework can easily be extended to interval-based UQ metrics^{21,45}.

These methods were applied to a representative example issued from a recent study by Busk *et al.*³¹ about atomization energies from the QM9 dataset, revealing a good average calibration, a slightly sub-optimal consistency, and a problematic adaptivity, either in the molecular mass space or the heteroatoms fraction space. This dataset presents several sources of stratification, and it was shown that the uncertainty due to the interplay of equal-size binning with data ordering expected for stratified conditioning variables is not dominant for the statistics considered here. An alternative strata-based binning LZMS approach led to similar diagnostics, with the inconvenience of a larger uncertainty on the validation statistics due to the smaller numbers of bins.

Up to now, ML-UQ validation studies in chemical and materials sciences are mainly focused on consistency. This covers somehow the concerns of ML-UQ designers who want the reliability of all uncertainties, either small or large. It was shown however that a positive consistency diagnostic does not augur of a positive adaptivity diagnostic, and therefore that a good consistency does not imply a good individual calibration. There is therefore a strong need that adaptivity be also systematically considered in ML-UQ studies, notably for final users, who expect the reliability of uncertainty for individual predictions, throughout the input features space.

ACKNOWLEDGMENTS

I warmly thank J. Busk for providing me the data for the running example of this study.

AUTHOR DECLARATIONS

Conflict of Interest

The author has no conflicts to disclose.

CODE AND DATA AVAILABILITY

The code and data to reproduce the results of this article are available at https://github.com/ppernot/2023_Adaptivity/releases/tag/v1.1 and at Zenodo⁶². The R,⁶³ ErrViewLib package implements the validation functions used in the present study, under version ErrViewLib-v1.7.3 (<https://github.com/ppernot/ErrViewLib/releases/tag/v1.7.3>), also available at Zenodo⁶⁴.

REFERENCES

- ¹G. Vishwakarma, A. Sonpal, and J. Hachmann. [Metrics for Benchmarking and Uncertainty Quantification: Quality, Applicability, and Best Practices for Machine Learning in Chemistry](#). *Trends in Chemistry*, 3:146–156, 2021.
- ²C. Gruich, V. Madhavan, Y. Wang, and B. Goldsmith. [Clarifying Trust of Materials Property Predictions using Neural Networks with Distribution-Specific Uncertainty Quantification](#). *Mach. Learn.: Sci. Technol.*, 4:025019, 2023.
- ³E. Heid, C. J. McGill, F. H. Vermeire, and W. H. Green. [Characterizing Uncertainty in Machine Learning for Chemistry](#). *J. Chem. Inf. Model.*, 63:4012–4029, 2023.
- ⁴S. B. Torrisi, M. Z. Bazant, A. E. Cohen, M. G. Cho, J. S. Hummelshøj, L. Hung, G. Kamat, A. Khajeh, A. Kolluru, X. Lei, H. Ling, J. H. Montoya, T. Mueller, A. Palizhati, B. A. Paren, B. Phan, J. Pietryga, E. Sandraz, D. Schweigert, Y. Shao-Horn, A. Trewartha, R. Zhu, D. Zhuang, and S. Sun. [Materials cartography: A forward-looking perspective on materials representation and devising better maps](#). *APL Mach. Learn.*, 1:020901, 2023.
- ⁵T. Pearce, A. Brintrup, M. Zaki, and A. Neely. [High-Quality Prediction Intervals for Deep Learning: A Distribution-Free, Ensembled Approach](#). In *International Conference on Machine Learning*, pages 4075–4084. PMLR, 2018. URL: <https://proceedings.mlr.press/v80/pearce18a.html>.
- ⁶F. Musil, M. J. Willatt, M. A. Langovoy, and M. Ceriotti. [Fast and accurate uncertainty estimation in chemical machine learning](#). *J. Chem. Theory Comput.*, 15:906–915, 2019.
- ⁷L. Hirschfeld, K. Swanson, K. Yang, R. Barzilay, and C. W. Coley. [Uncertainty Quantification Using Neural Networks for Molecular Property Prediction](#). *J. Chem. Inf. Model.*, 60:3770–3780, 2020.
- ⁸K. Tran, W. Neiswanger, J. Yoon, Q. Zhang, E. Xing, and Z. W. Ulissi. [Methods for comparing uncertainty quantifications for material property predictions](#). *Mach. Learn.: Sci. Technol.*, 1:025006, 2020.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset. PLEASE CITE THIS ARTICLE AS DOI: 10.1063/1.50174943

- ⁹M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, and S. Nahavandi. [A review of uncertainty quantification in deep learning: Techniques, applications and challenges](#). *Inform. Fusion*, 76:243–297, 2021.
- ¹⁰J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, M. Shahzad, W. Yang, R. Bamler, and X. X. Zhu. [A Survey of Uncertainty in Deep Neural Networks](#). *arXiv:2107.03342*, July 2021.
- ¹¹M. Tynes, W. Gao, D. J. Burrill, E. R. Batista, D. Perez, P. Yang, and N. Lubbers. [Pairwise difference regression: A machine learning meta-algorithm for improved prediction and uncertainty quantification in chemical search](#). *J. Chem. Inf. Model.*, 61:3846–3857, 2021. PMID: 34347460.
- ¹²E. Zelikman, C. Healy, S. Zhou, and A. Avati. [CRUDE: Calibrating Regression Uncertainty Distributions Empirically](#). *arXiv:2005.12496*, March 2021.
- ¹³Y. Hu, J. Musielewicz, Z. W. Ulissi, and A. J. Medford. [Robust and scalable uncertainty estimation with conformal prediction for machine-learned interatomic potentials](#). *Mach. Learn.: Sci. Technol.*, 3:045028, November 2022.
- ¹⁴D. Varivoda, R. Dong, S. S. Omeel, and J. Hu. [Materials Property Prediction with Uncertainty Quantification: A Benchmark Study](#). *arXiv:2211.02235*, November 2022.
- ¹⁵M. Battaglia, E. Comi, T. Stadelmann, R. Hiestand, B. Ruhstaller, and E. Knapp. [Deep ensemble inverse model for image-based estimation of solar cell parameters](#). *APL Mach. Learn.*, 1:036108, 2023.
- ¹⁶J. Busk, M. N. Schmidt, O. Winther, T. Vegge, and P. B. Jørgensen. [Graph Neural Network Interatomic Potential Ensembles with Calibrated Aleatoric and Epistemic Uncertainty on Energy and Forces](#). *arXiv:2305.16325*, May 2023.
- ¹⁷W. He and Z. Jiang. [A Survey on Uncertainty Quantification Methods for Deep Neural Networks: An Uncertainty Source Perspective](#). *arXiv:2302.13425*, February 2023.
- ¹⁸T. Mohanty, K. S. R. Chandran, and T. D. Sparks. [Machine learning guided optimal composition selection of niobium alloys for high temperature applications](#). *APL Mach. Learn.*, 1:036102, 2023.
- ¹⁹T. Tohme, K. Vanslette, and K. Youcef-Toumi. [Reliable neural networks for regression uncertainty estimation](#). *Reliab. Eng. Syst. Saf.*, 229:108811, 2023.
- ²⁰Y. Liu, M. Pagliardini, T. Chavdarova, and S. U. Stich. [The Peril of Popular Deep Learning Uncertainty Estimation Methods](#). *arXiv:2112.05000*, 2021.

- ²¹P. Pernot. [Prediction uncertainty validation for computational chemists](#). *J. Chem. Phys.*, 157:144103, 2022.
- ²²BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP, and OIML. [Evaluation of measurement data - Guide to the expression of uncertainty in measurement \(GUM\)](#). Technical Report 100:2008, Joint Committee for Guides in Metrology, JCGM, 2008. URL: http://www.bipm.org/utils/common/documents/jcgm/JCGM_100_2008_F.pdf.
- ²³K. K. Irikura, R. D. Johnson, and R. N. Kacker. [Uncertainty associated with virtual measurements from computational quantum chemistry models](#). *Metrologia*, 41:369–375, 2004.
- ²⁴B. Ruscic, R. E. Pinzon, M. L. Morton, G. von Laszewski, S. J. Bittner, S. G. Nijssure, K. A. Amin, M. Minkoff, and A. F. Wagner. [Introduction to Active Thermochemical Tables: Several "key" enthalpies of formation revisited](#). *J. Phys. Chem. A*, 108:9979–9997, 2004.
- ²⁵J. P. Janet, C. Duan, T. Yang, A. Nandy, and H. J. Kulik. [A quantitative uncertainty metric controls error in neural network-driven chemical discovery](#). *Chem. Sci.*, 10:7913–7922, 2019.
- ²⁶E. Hüllermeier and W. Waegeman. [Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods](#). *Mach. Learn.*, 110:457–506, 2021.
- ²⁷V. Korolev, I. Nevolin, and P. Protsenko. [A universal similarity based approach for predictive uncertainty quantification in materials science](#). *Sci. Rep.*, 12:1–10, 2022.
- ²⁸C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. [On Calibration of Modern Neural Networks](#). In *International Conference on Machine Learning*, pages 1321–1330. 2017. URL: <https://proceedings.mlr.press/v70/guo17a.html>.
- ²⁹V. Kuleshov, N. Fenner, and S. Ermon. [Accurate uncertainties for deep learning using calibrated regression](#). In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2796–2804. PMLR, 10–15 Jul 2018. URL: <https://proceedings.mlr.press/v80/kuleshov18a.html>.
- ³⁰D. Levi, L. Gispan, N. Giladi, and E. Fetaya. [Evaluating and Calibrating Uncertainty Prediction in Regression Tasks](#). *Sensors*, 22:5540, 2022.
- ³¹J. Busk, P. B. Jørgensen, A. Bhowmik, M. N. Schmidt, O. Winther, and T. Vegge. [Calibrated uncertainty for molecular property prediction using ensembles of message passing neural networks](#). *Mach. Learn.: Sci. Technol.*, 3:015012, 2022.
- ³²V. Vovk. [Conditional validity of inductive conformal predictors](#). In S. C. H. Hoi and W. Buntine, editors, *Proceedings of the Asian Conference on Machine Learning*, volume 25 of *Proceedings of Machine Learning Research*, pages 475–490, Singapore Management University, Singapore, 04–06 Nov 2012. PMLR. URL: <https://proceedings.mlr.press/v25/>

[vovk12.html](#).

- ³³A. N. Angelopoulos and S. Bates. [A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification](#). *arXiv:2107.07511*, July 2021.
- ³⁴M. Cauchois, S. Gupta, and J. C. Duchi. [Knowing what you know: valid and validated confidence sets in multiclass and multilabel prediction](#). *J. Mach. Learn. Res.*, 22:3681–3722, 2021.
- ³⁵S. Feldman, S. Bates, and Y. Romano. [Improving Conditional Coverage via Orthogonal Quantile Regression](#). *arXiv:2106.00394*, 2021.
- ³⁶M. Reiher. [Molecule-specific uncertainty quantification in quantum chemical studies](#). *Isr. J. Chem.*, 62:e202100101, 2022.
- ³⁷E. Ilg, O. Çiçek, S. Galesso, A. Klein, O. Makansi, F. Hutter, and T. Brox. [Uncertainty Estimates and Multi-hypotheses Networks for Optical Flow](#). In *Computer Vision – ECCV 2018*, pages 677–693. Springer, Cham, Switzerland, 2018.
- ³⁸G. Scalia, C. A. Grambow, B. Pernici, Y.-P. Li, and W. H. Green. [Evaluating scalable uncertainty estimation methods for deep learning-based molecular property prediction](#). *J. Chem. Inf. Model.*, 60:2697–2717, 2020.
- ³⁹Y. Chung, W. Neiswanger, I. Char, and J. Schneider. [Beyond pinball loss: Quantile methods for calibrated uncertainty quantification](#). *arXiv:2011.09588*, 2020.
- ⁴⁰S. Zhao, T. Ma, and S. Ermon. [Individual Calibration with Randomized Forecasting](#). *arXiv:2006.10288*, June 2020.
- ⁴¹M.-H. Laves, S. Ihler, J. F. Fast, L. A. Kahrs, and T. Ortmaier. [Well-calibrated regression uncertainty in medical imaging with deep learning](#). In T. Arbel, I. Ben Ayed, M. de Bruijne, M. Descoteaux, H. Lombaert, and C. Pal, editors, *Proceedings of the Third Conference on Medical Imaging with Deep Learning*, volume 121 of *Proceedings of Machine Learning Research*, pages 393–412. PMLR, 06–08 Jul 2020. URL: <https://proceedings.mlr.press/v121/laves20a.html>.
- ⁴²D. Levi, L. Gispan, N. Giladi, and E. Fetaya. [Evaluating and Calibrating Uncertainty Prediction in Regression Tasks](#). *arXiv:1905.11659v3*, February 2020.
- ⁴³L. I. Vazquez-Salazar, E. D. Boittier, and M. Meuwly. [Uncertainty quantification for predictions of atomistic neural networks](#). *arXiv:2207.06916*, July 2022.
- ⁴⁴P. Pernot. [Validation of uncertainty quantification metrics: a primer based on the consistency and adaptivity concepts](#). *arXiv:2303.07170*, 2023. [arXiv:2303.07170](#).
- ⁴⁵P. Pernot. [The long road to calibrated prediction uncertainty in computational chemistry](#). *J. Chem. Phys.*, 156:114109, 2022.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset. PLEASE CITE THIS ARTICLE AS DOI:10.1063/1.5017413

- ⁴⁶Y. Dar, V. Muthukumar, and R. G. Baraniuk. [A Farewell to the Bias-Variance Tradeoff? An Overview of the Theory of Overparameterized Machine Learning](#). *arXiv:2109.02355*, September 2021.
- ⁴⁷P. Pernot and F. Cailliez. [A critical review of statistical calibration/prediction models handling data inconsistency and model inadequacy](#). *AIChE J.*, 63:4642–4665, 2017.
- ⁴⁸R. T. Birge. [The calculation of errors by the method of least squares](#). *Phys. Rev.*, 40:207–227, 1932.
- ⁴⁹R. Luo, A. Bhatnagar, Y. Bai, S. Zhao, H. Wang, C. Xiong, S. Savarese, S. Ermon, E. Schmerling, and M. Pavone. [Local Calibration: Metrics and Recalibration](#). *arXiv:2102.10809*, February 2021.
- ⁵⁰G. Palmer, S. Du, A. Politowicz, J. P. Emory, X. Yang, A. Gautam, G. Gupta, Z. Li, R. Jacobs, and D. Morgan. [Calibration after bootstrap for accurate uncertainty quantification in regression models](#). *npj Comput. Mater.*, 8:1–9, 2022.
- ⁵¹P. Pernot. [Stratification of uncertainties recalibrated by isotonic regression and its impact on calibration error statistics](#). *arXiv:2306.05180*, June 2023.
- ⁵²R. N. Kacker, R. Kessel, and K.-D. Sommer. [Assessing differences between results determined according to the guide to the expression of uncertainty in measurement](#). *J. Res. Nat. Inst. Stand. Technol.*, 115(6):453, 2010.
- ⁵³B. Efron. [Bootstrap Methods: Another Look at the Jackknife](#). *Ann. Stat.*, 7:1–26, January 1979.
- ⁵⁴B. Efron and R. Tibshirani. [Statistical data analysis in the computer age](#). *Science*, 253:390–395, 1991.
- ⁵⁵K. A. Maupin, L. P. Swiler, and N. W. Porter. [Validation Metrics for Deterministic and Probabilistic Data](#). *J. Verif. Validation Uncertainty Quantif.*, 3:031002, 2018.
- ⁵⁶P. Pernot. [Properties of the ENCE and other MAD-based calibration metrics](#). *arXiv:2305.11905*, May 2023.
- ⁵⁷P. Pernot. [Confidence curves for UQ validation: probabilistic reference vs. oracle](#). *arXiv:2206.15272*, June 2022.
- ⁵⁸M. H. Rasmussen, C. Duan, H. J. Kulik, and J. H. Jensen. [Uncertain of uncertainties? A comparison of uncertainty quantification metrics for chemical data sets](#). *ChemRxiv*, September 2023.
- ⁵⁹L. van der Maaten and G. Hinton. [Visualizing data using t-sne](#). *J Mach Learn Res*, 9(8):2579–2605, 2008. URL: <http://jmlr.org/papers/v9/vandermaaten08a.html>.

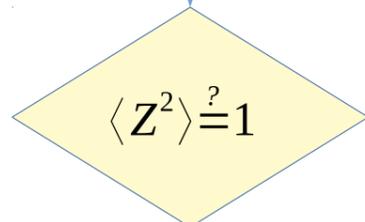
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0174943

- ⁶⁰L. McInnes, J. Healy, and J. Melville. [UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction](#). *arXiv:1802.03426*, February 2018.
- ⁶¹Y. Chung, I. Char, H. Guo, J. Schneider, and W. Neiswanger. [Uncertainty toolbox: an open-source library for assessing, visualizing, and improving uncertainty quantification](#). *arXiv:2109.10254*, September 2021.
- ⁶²P. Pernot. [Codes and data for the reproduction of the results of the present paper](#), 2023.
- ⁶³R Core Team. [R: A Language and Environment for Statistical Computing](#). R Foundation for Statistical Computing, Vienna, Austria, 2019. URL: <http://www.R-project.org/>.
- ⁶⁴P. Pernot. [Errviewlib-v1.7.3](#), 2023.

ML Test set
 X : input features
 E : prediction errors
 u_E : prediction uncertainties
 Z : Z-scores (E / u_E)

1/ Test
Average
Calibration



No

Recalibrate !

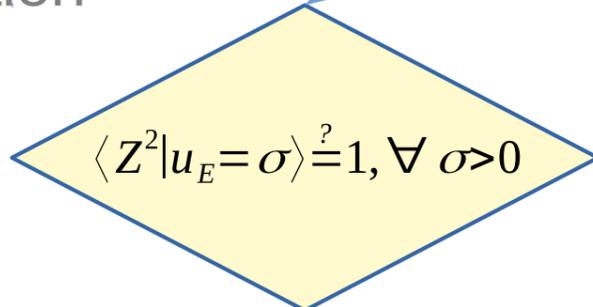
Yes

2/ Test
Conditional
Calibration

Design validation groups
Bins, groups, strata...

u_E

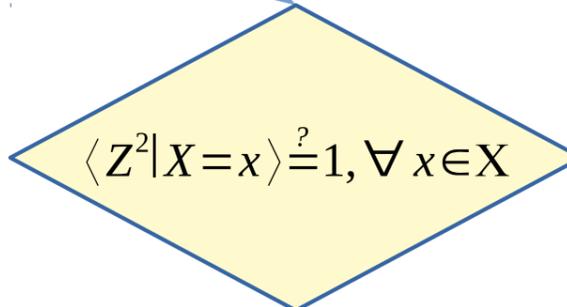
X



Yes

Consistency 😊

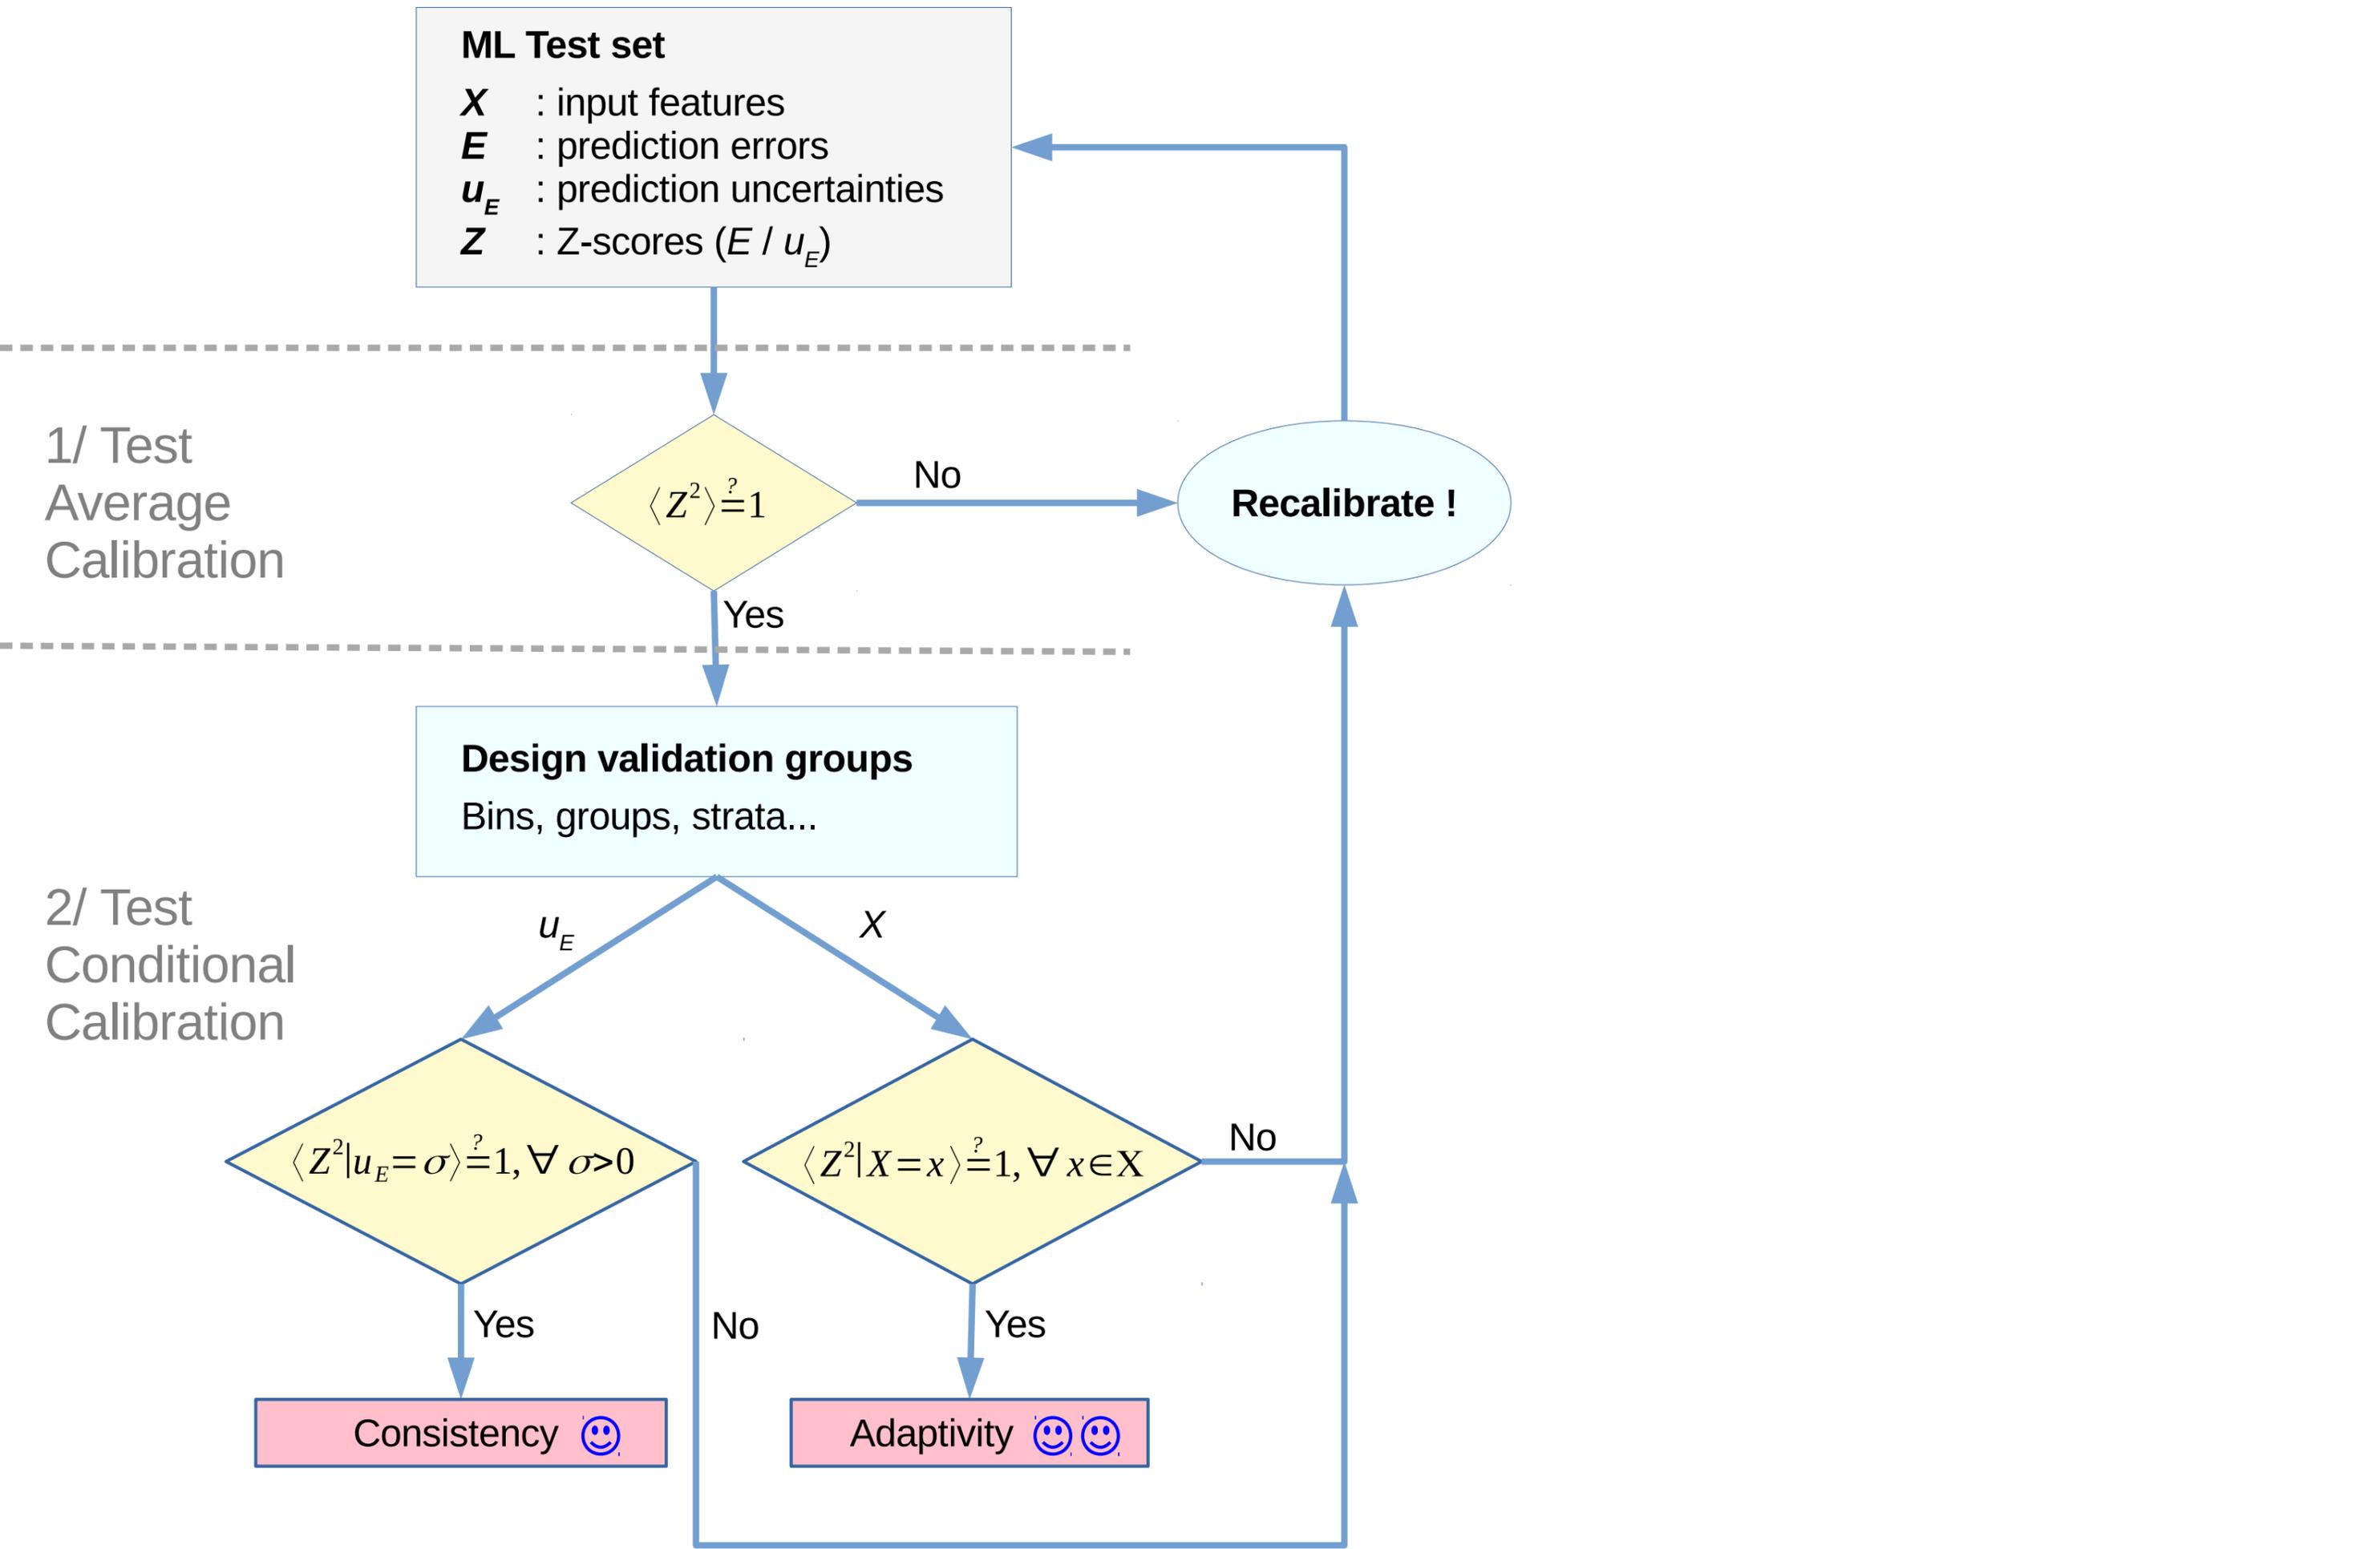
No



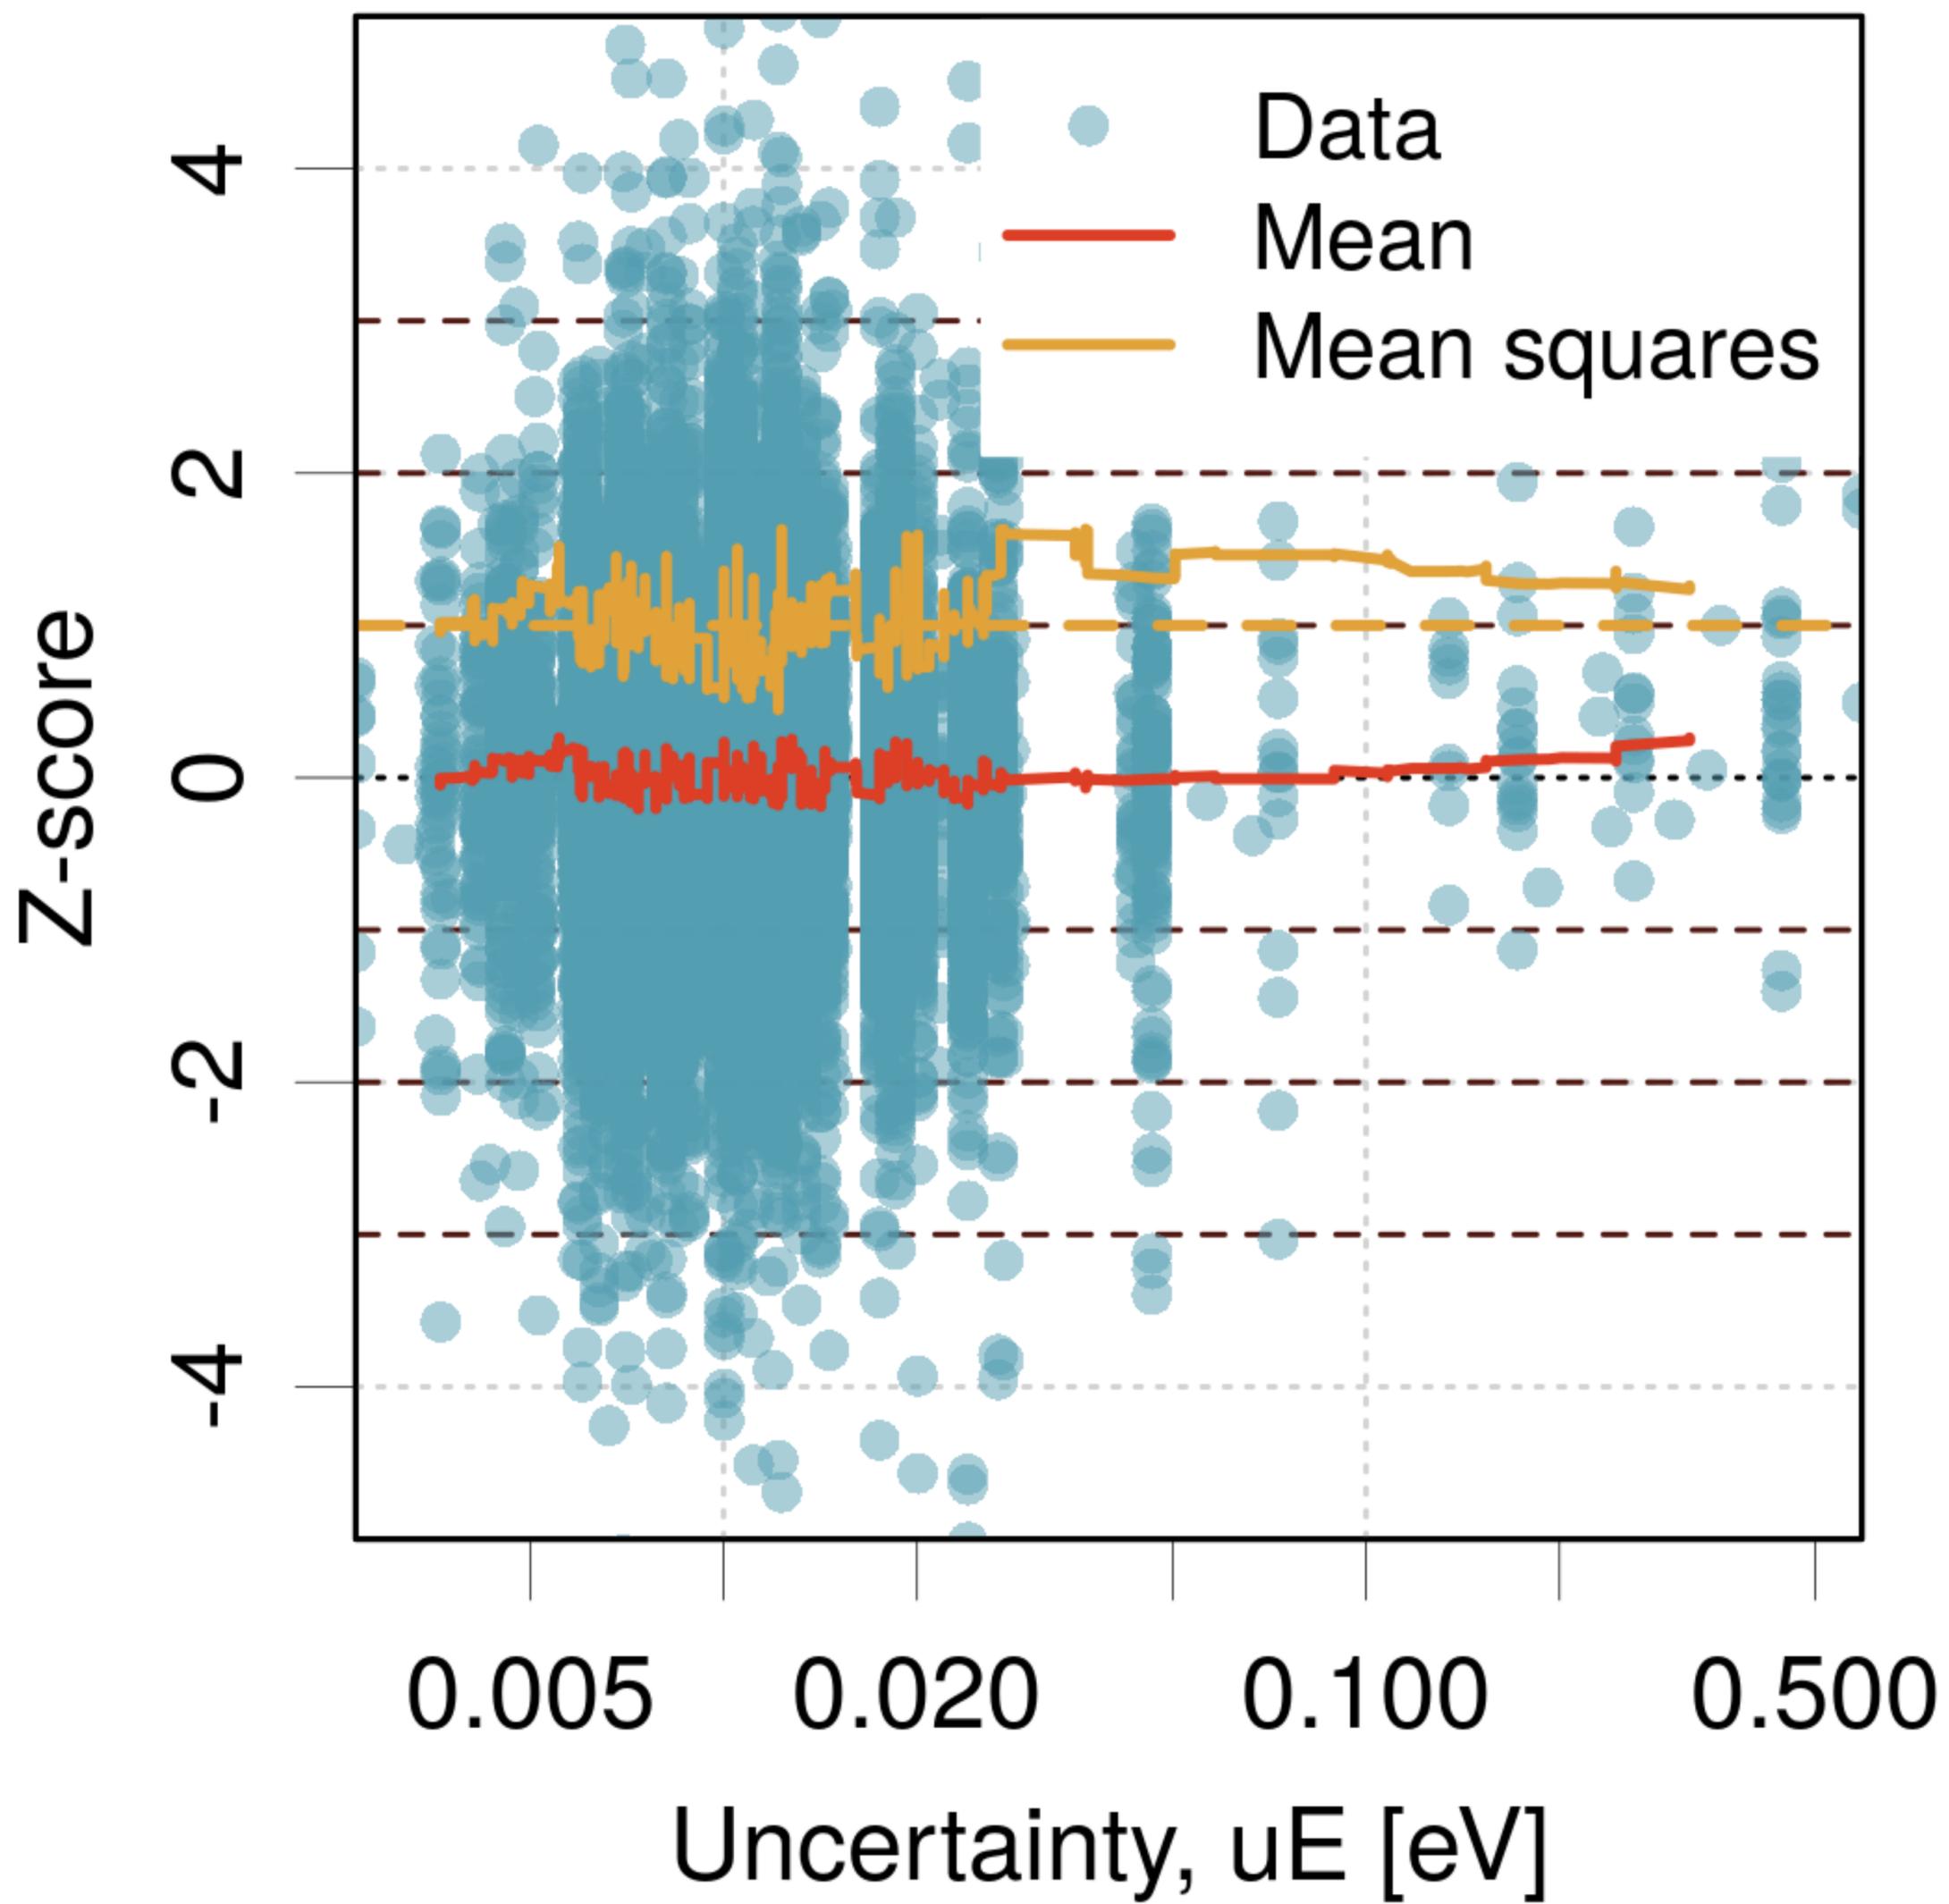
No

Adaptivity 😊😊

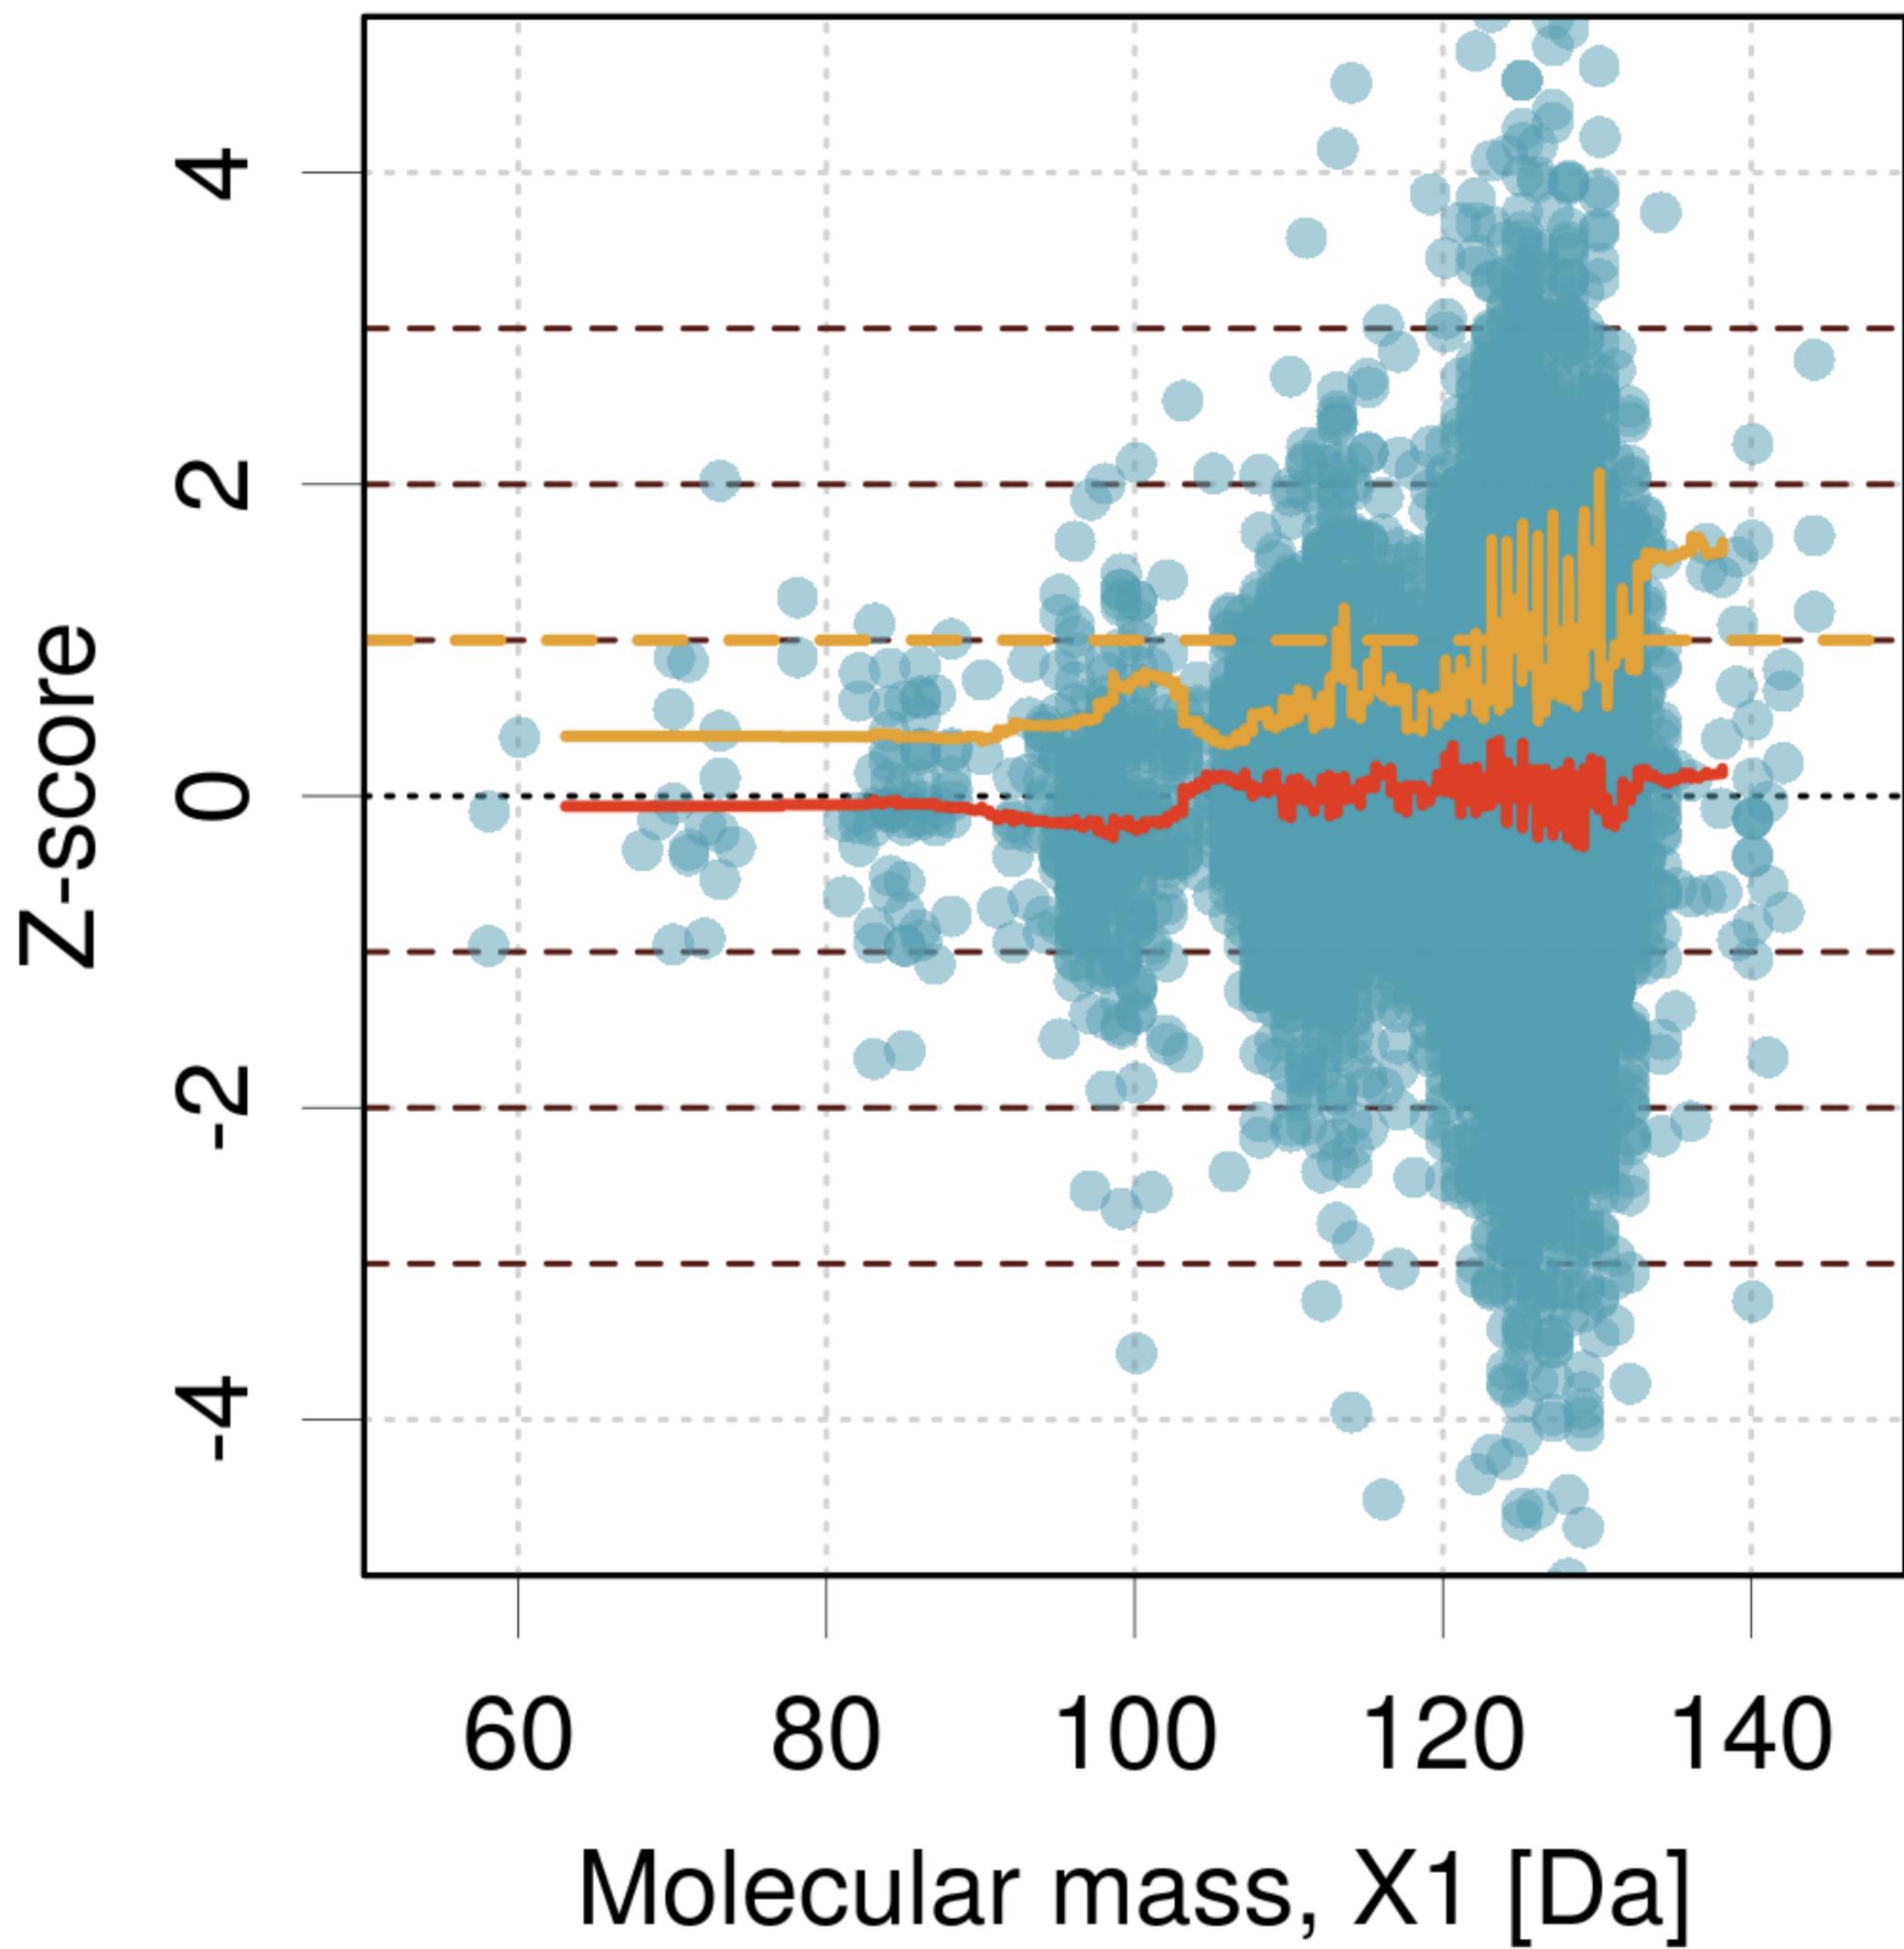
Yes



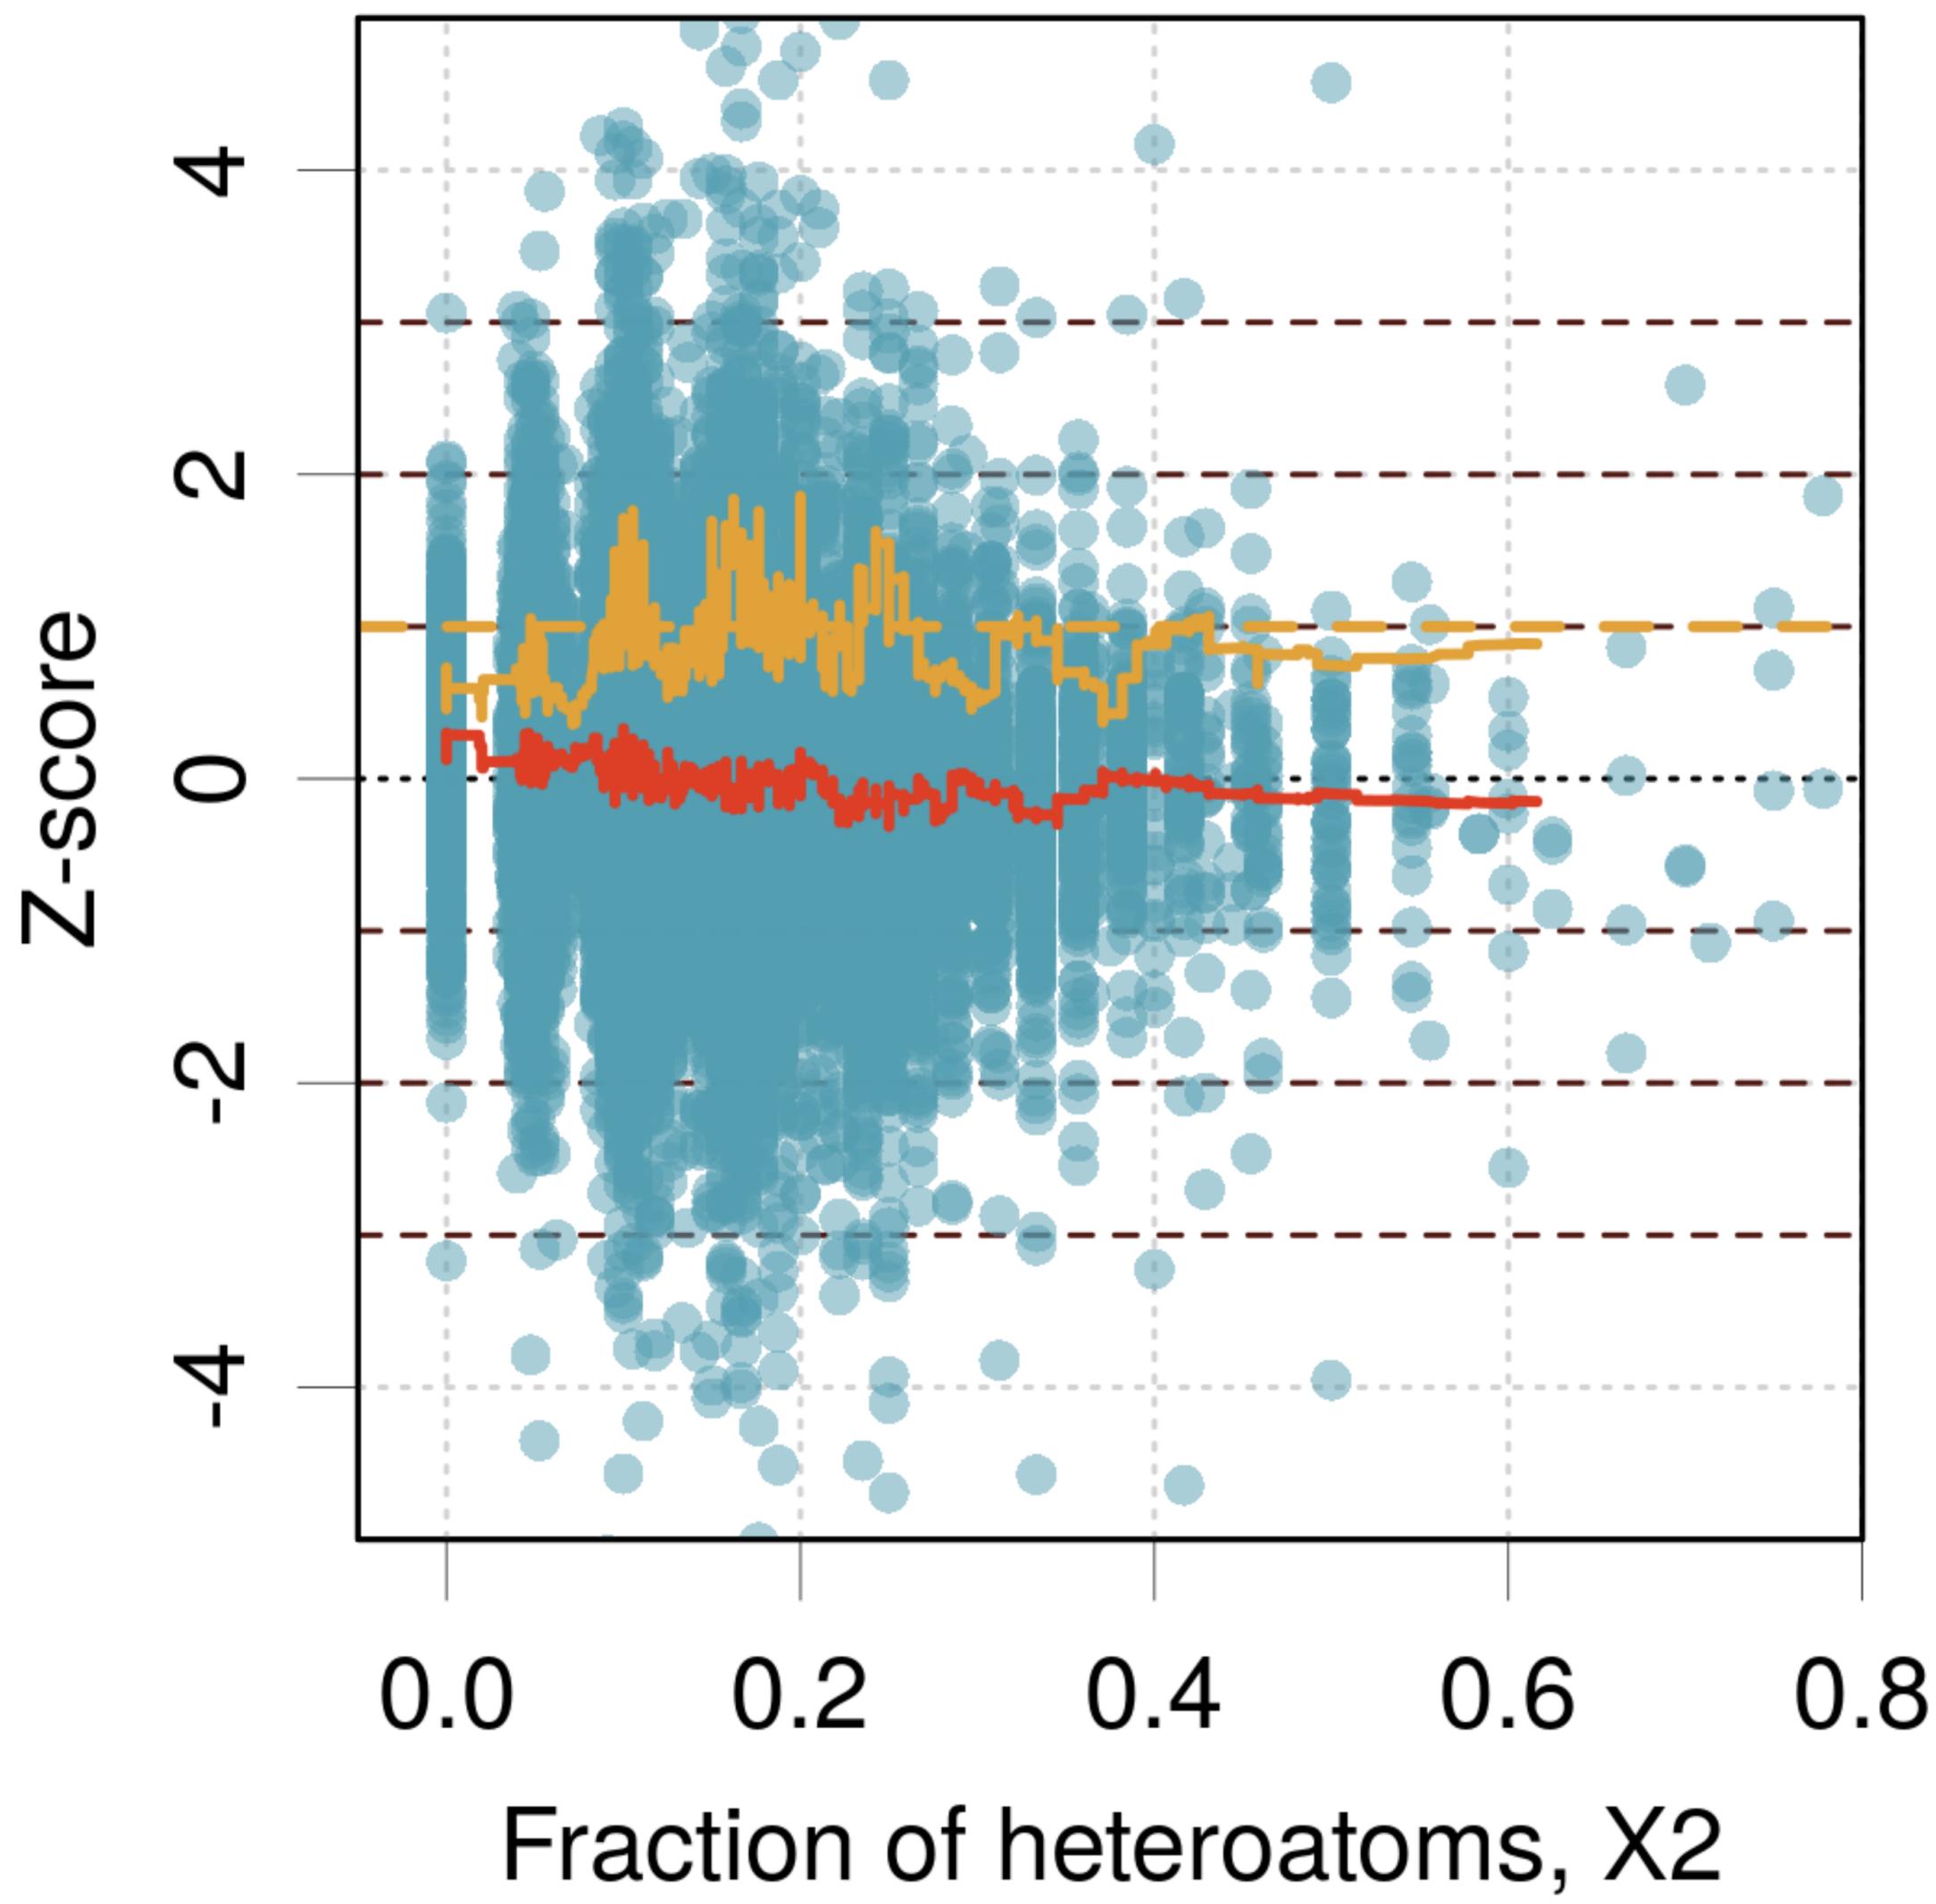
(a)

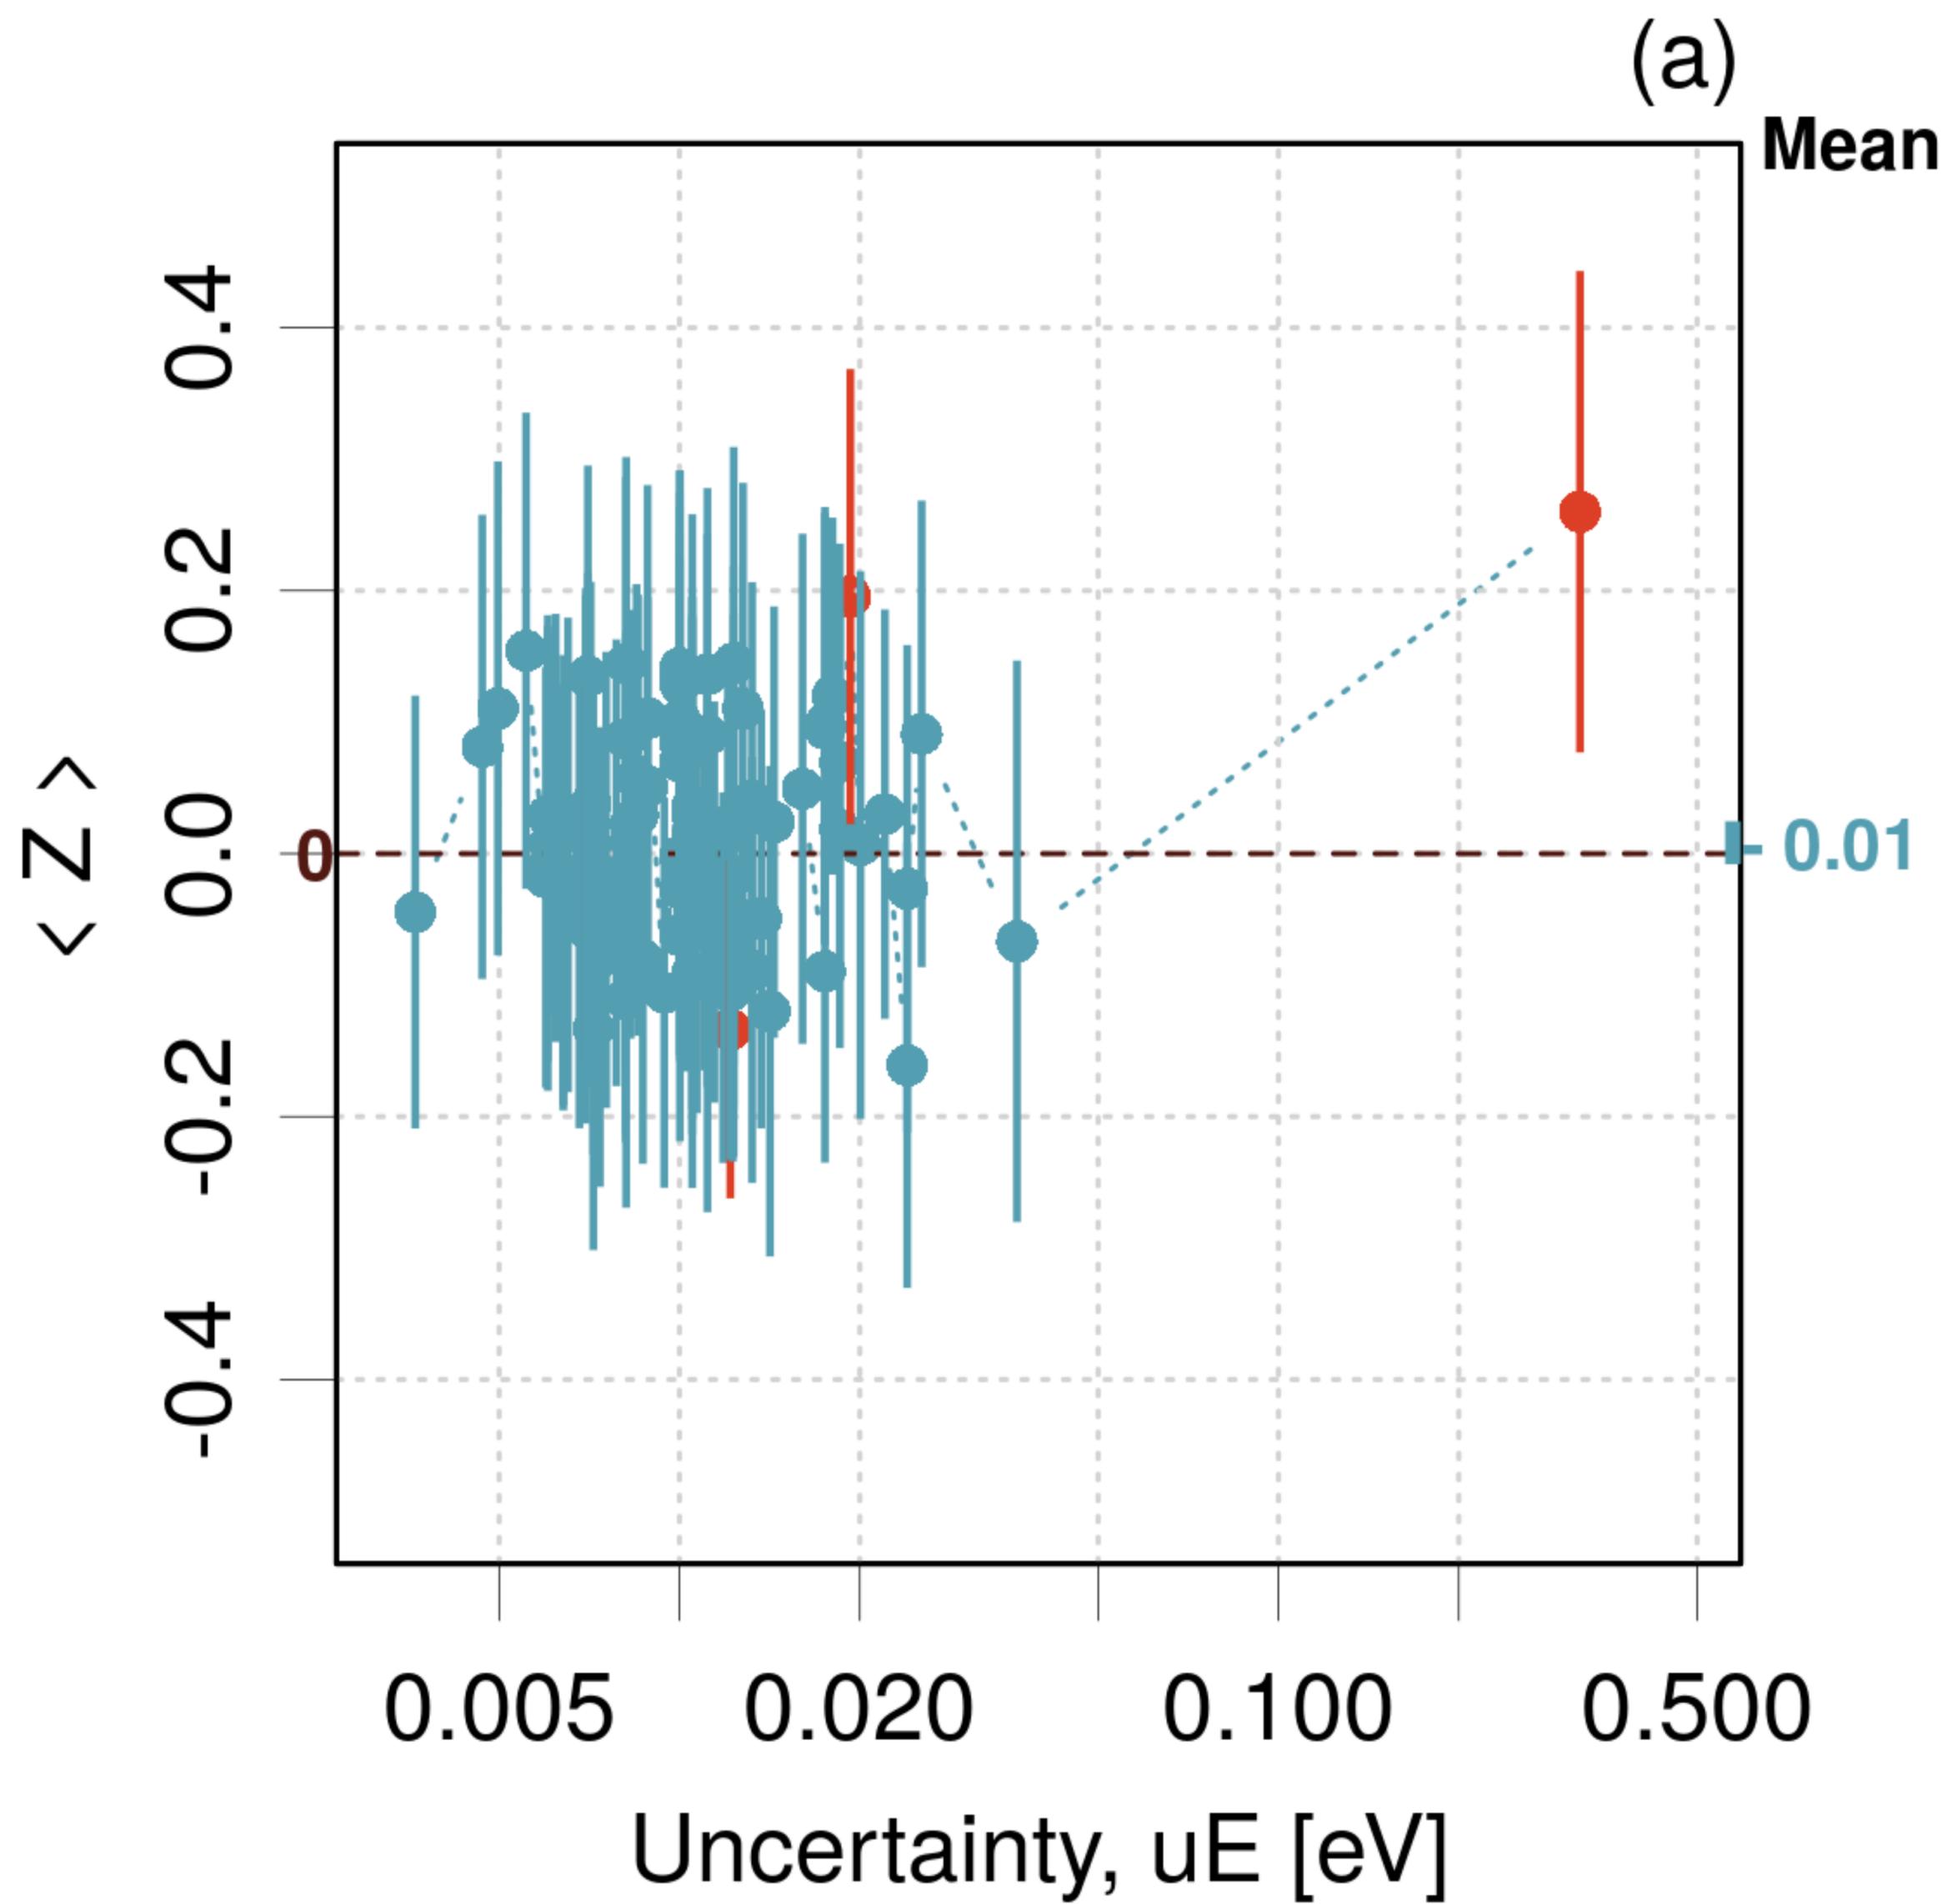


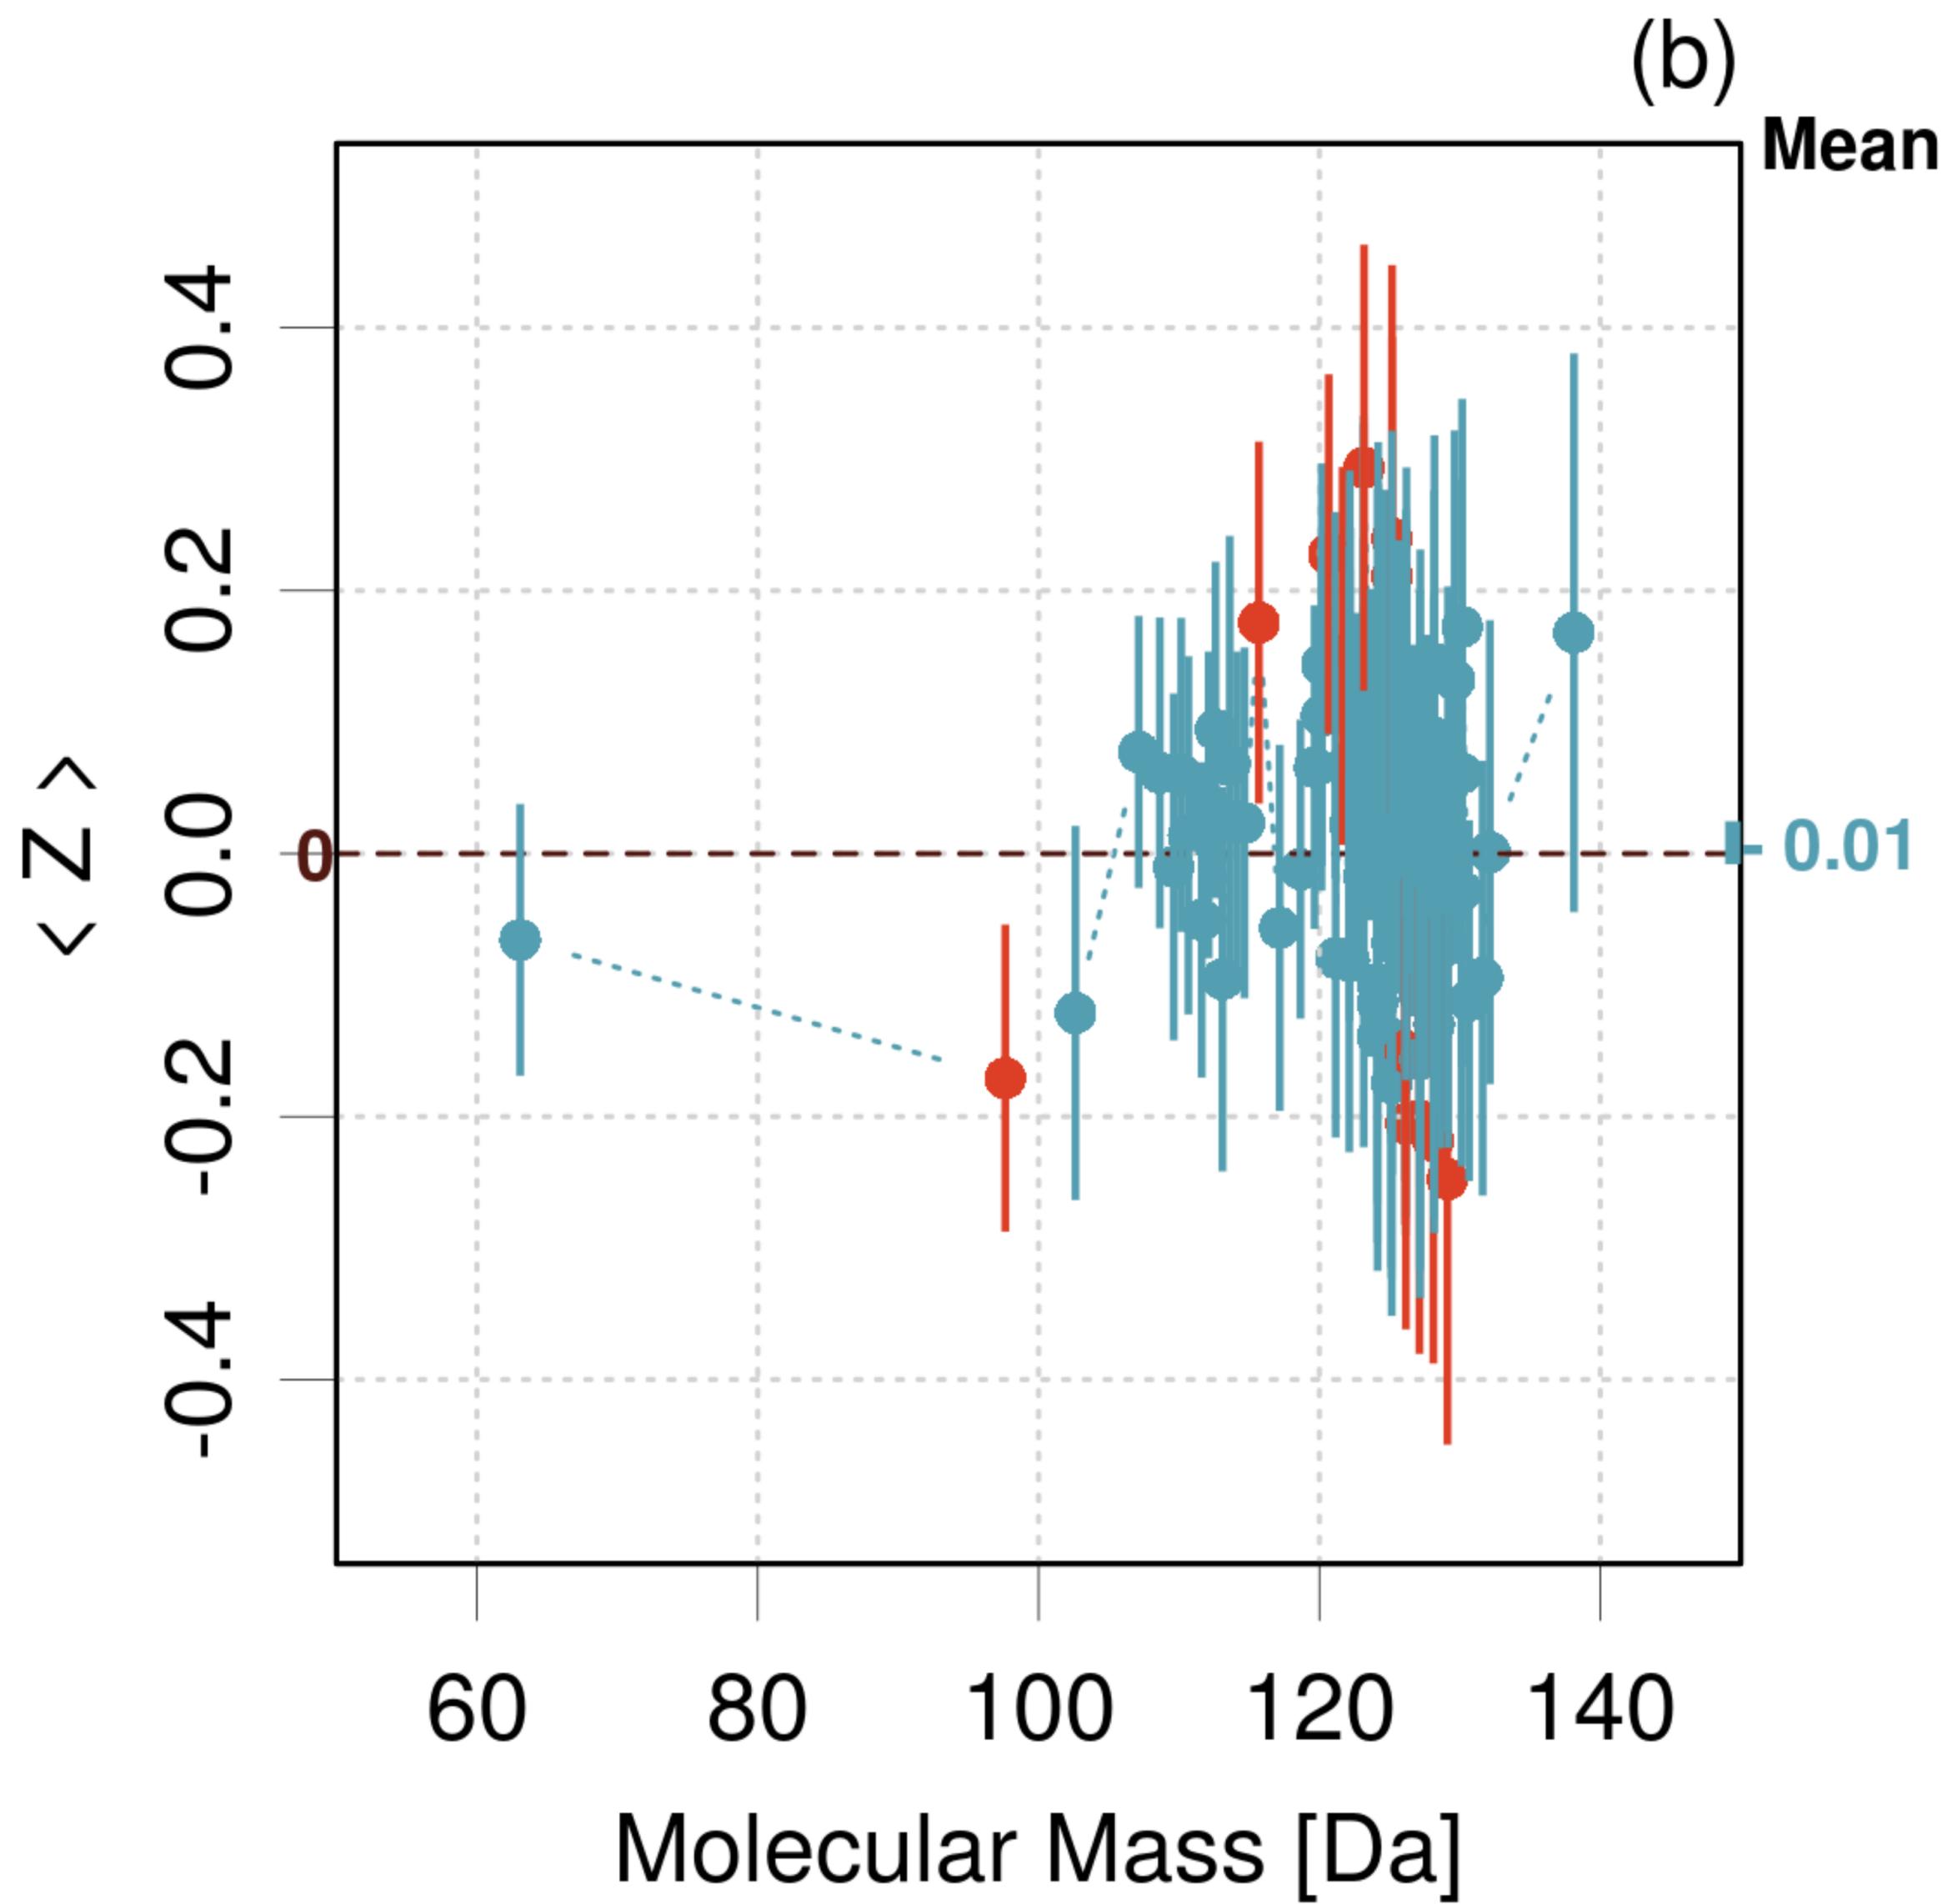
(b)

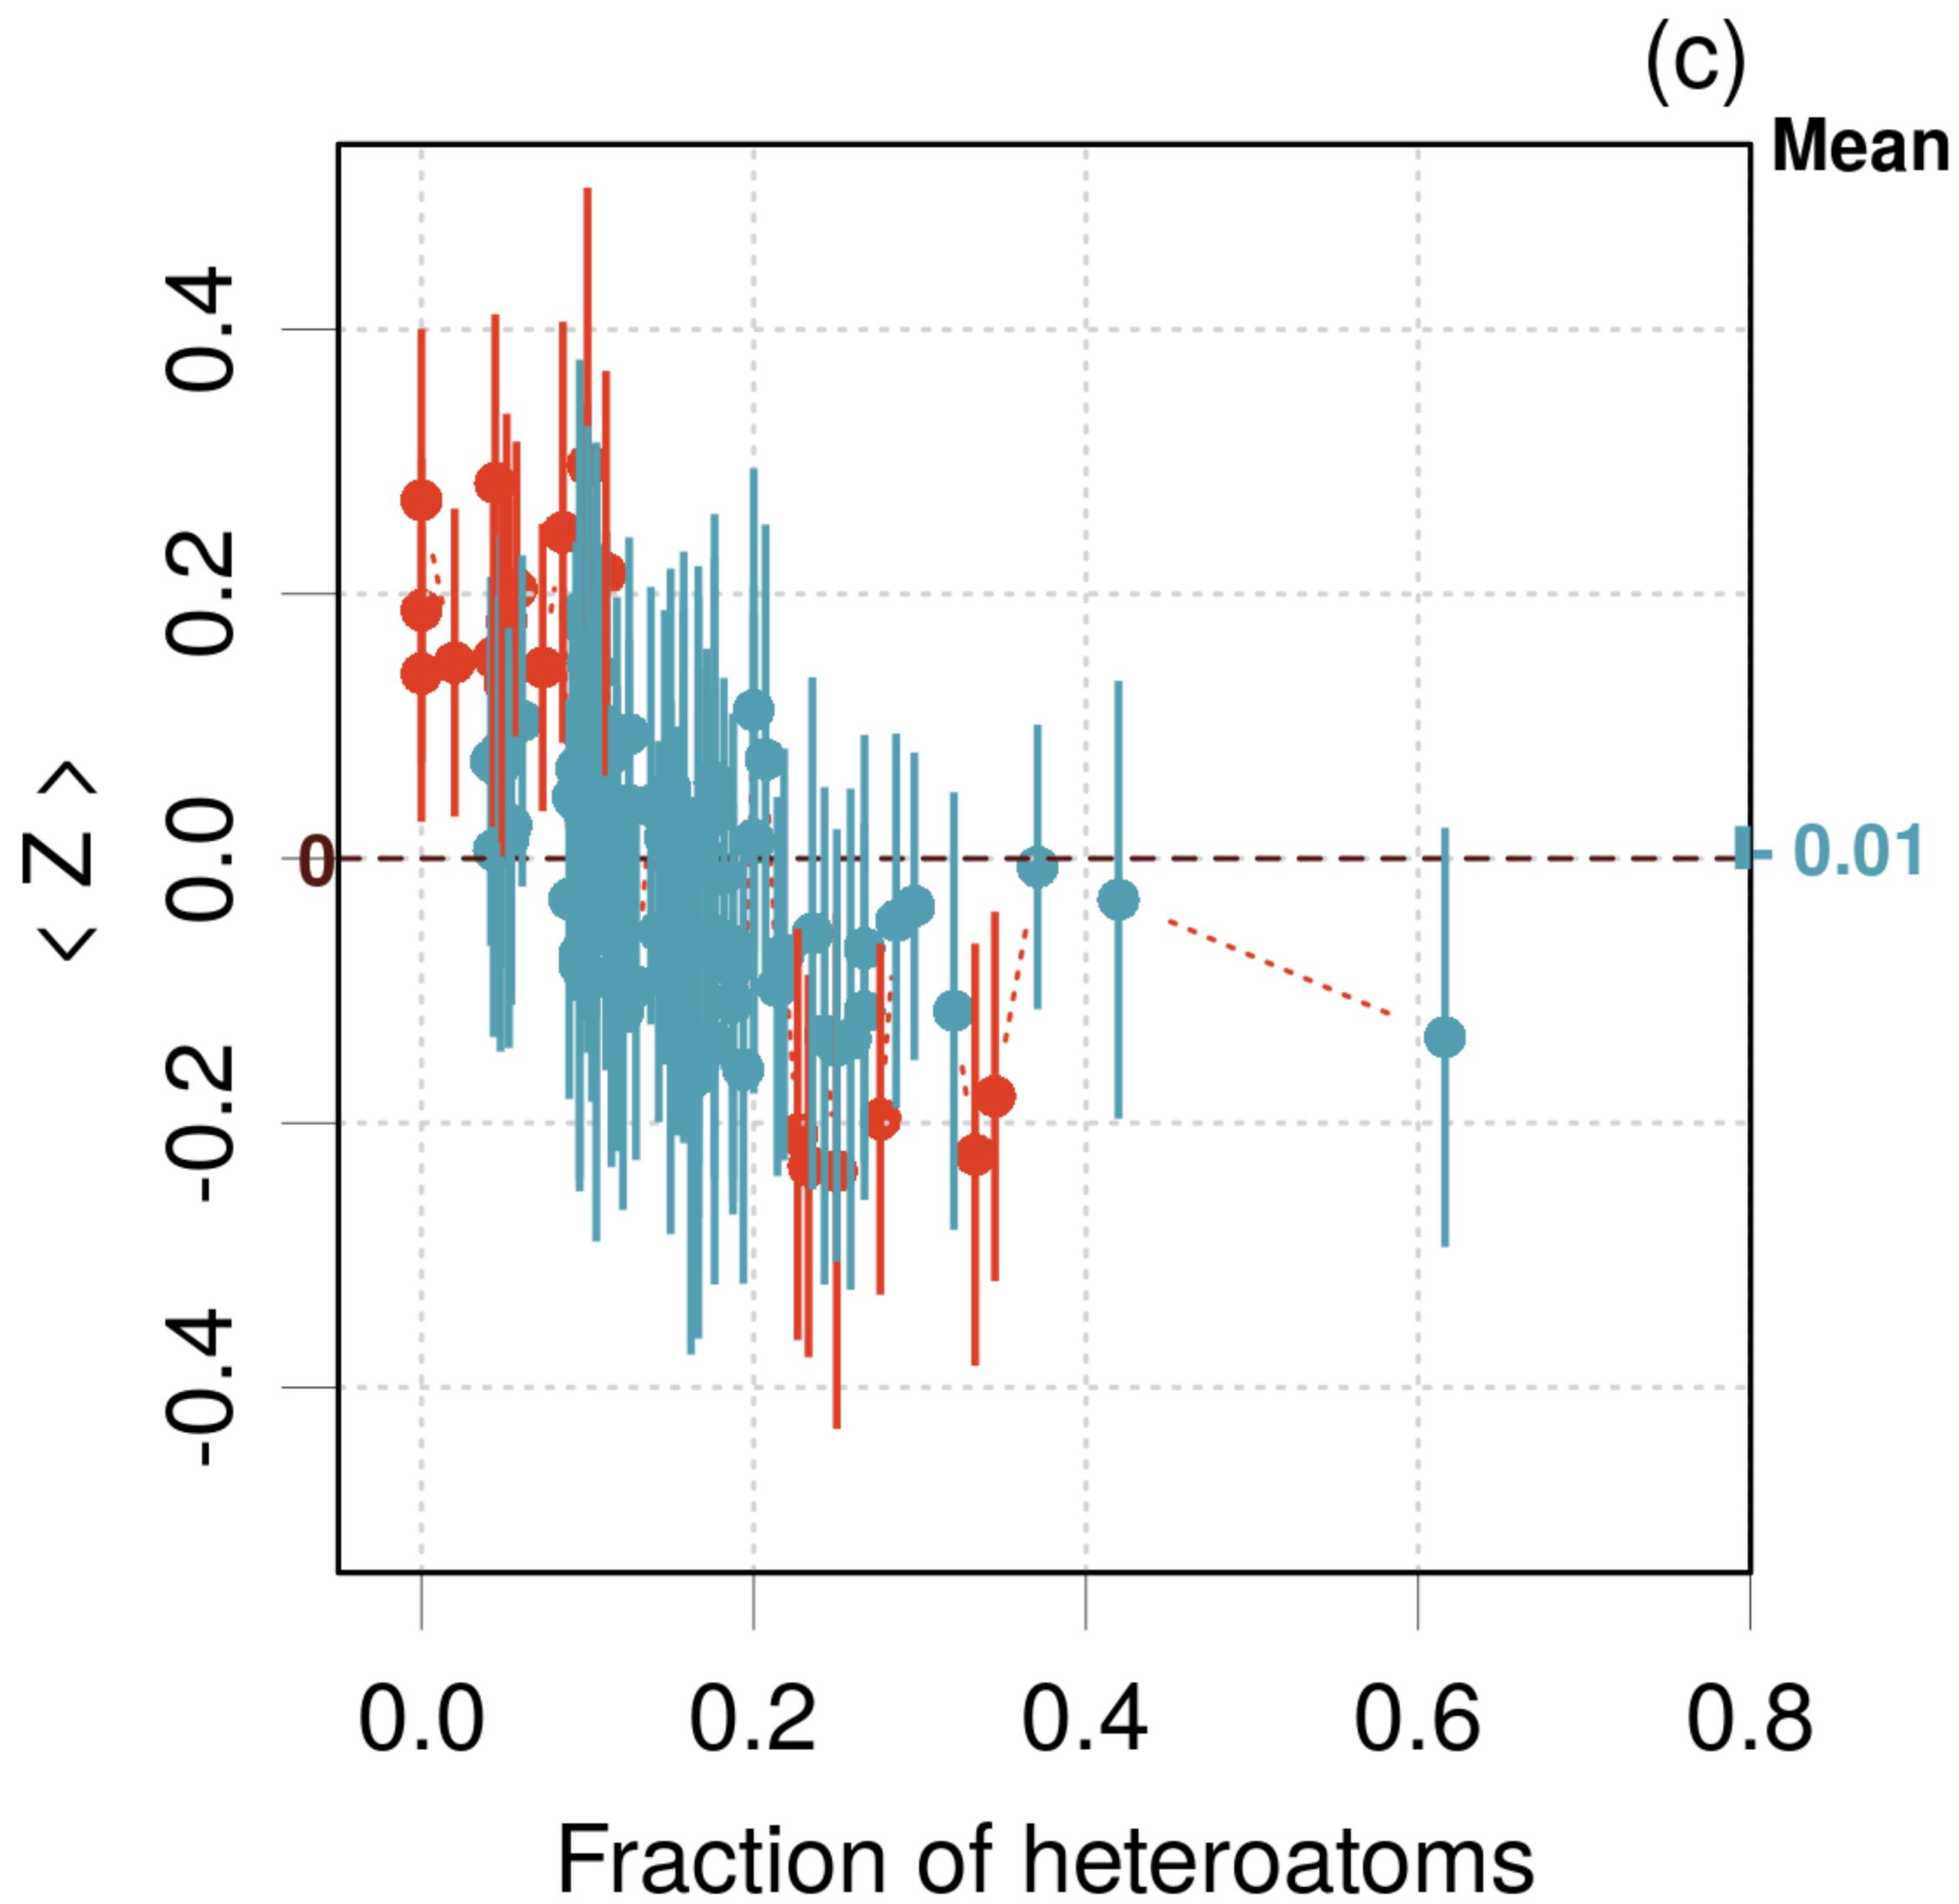


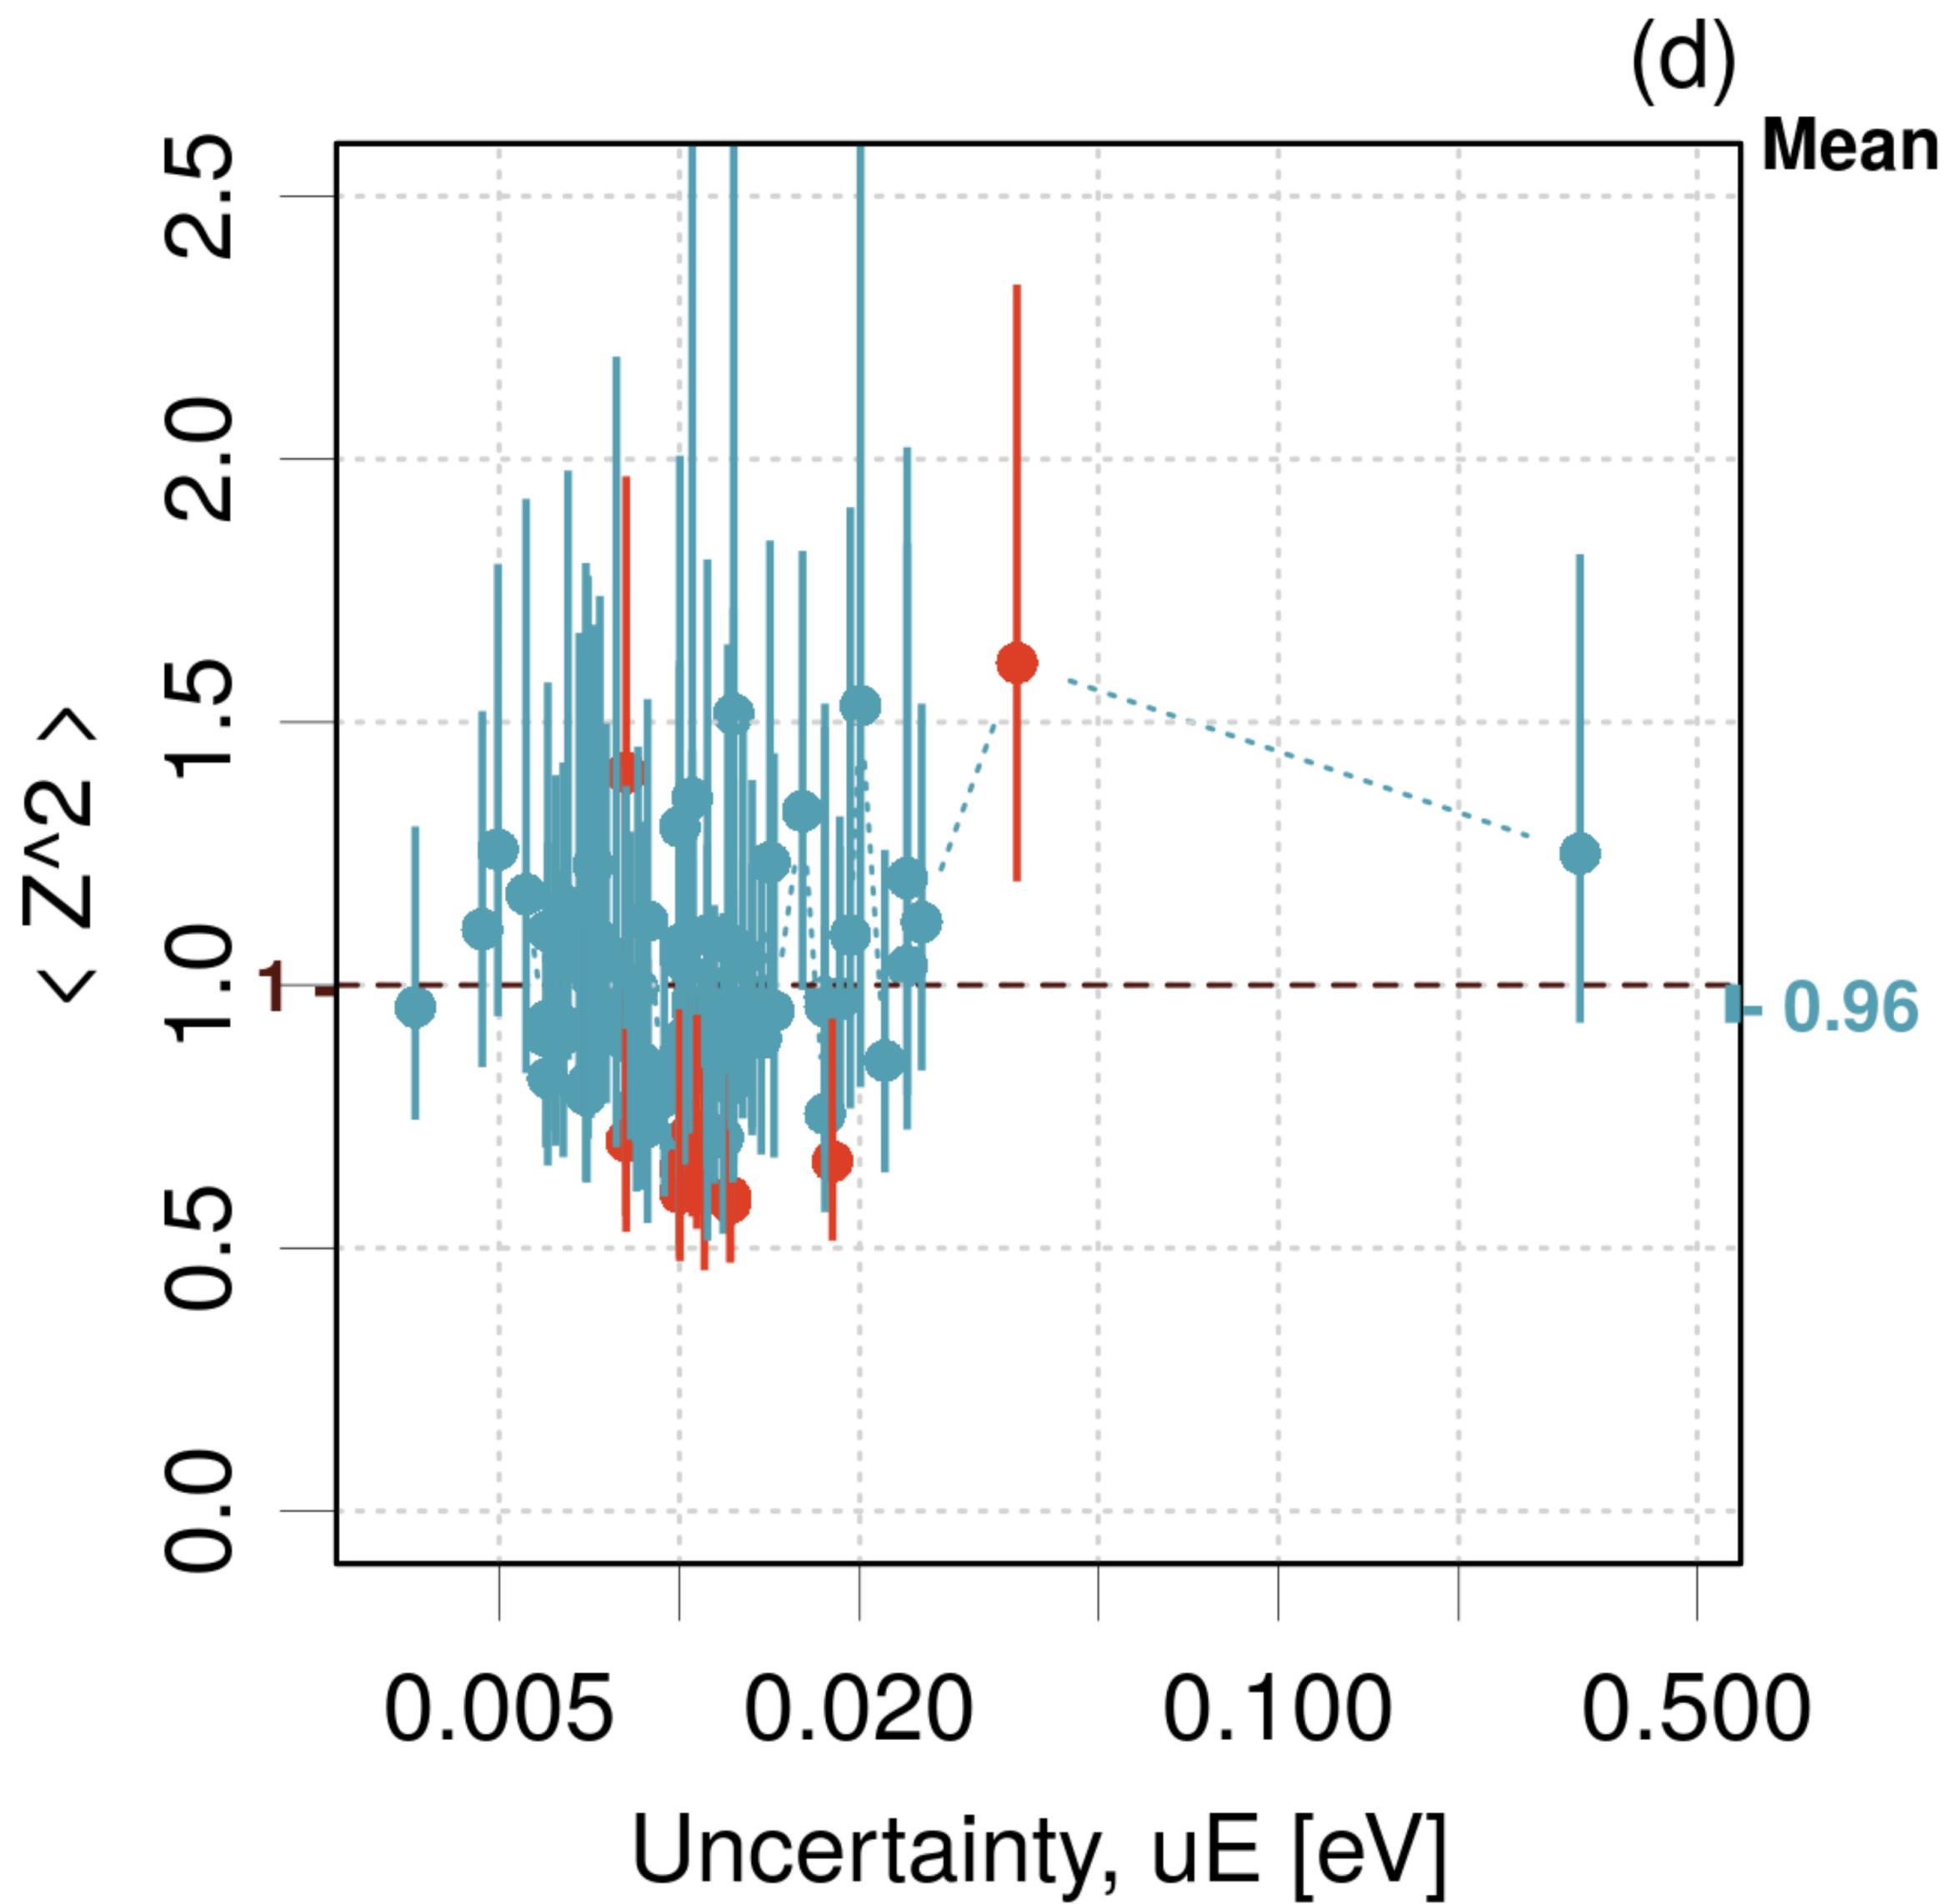
(c)

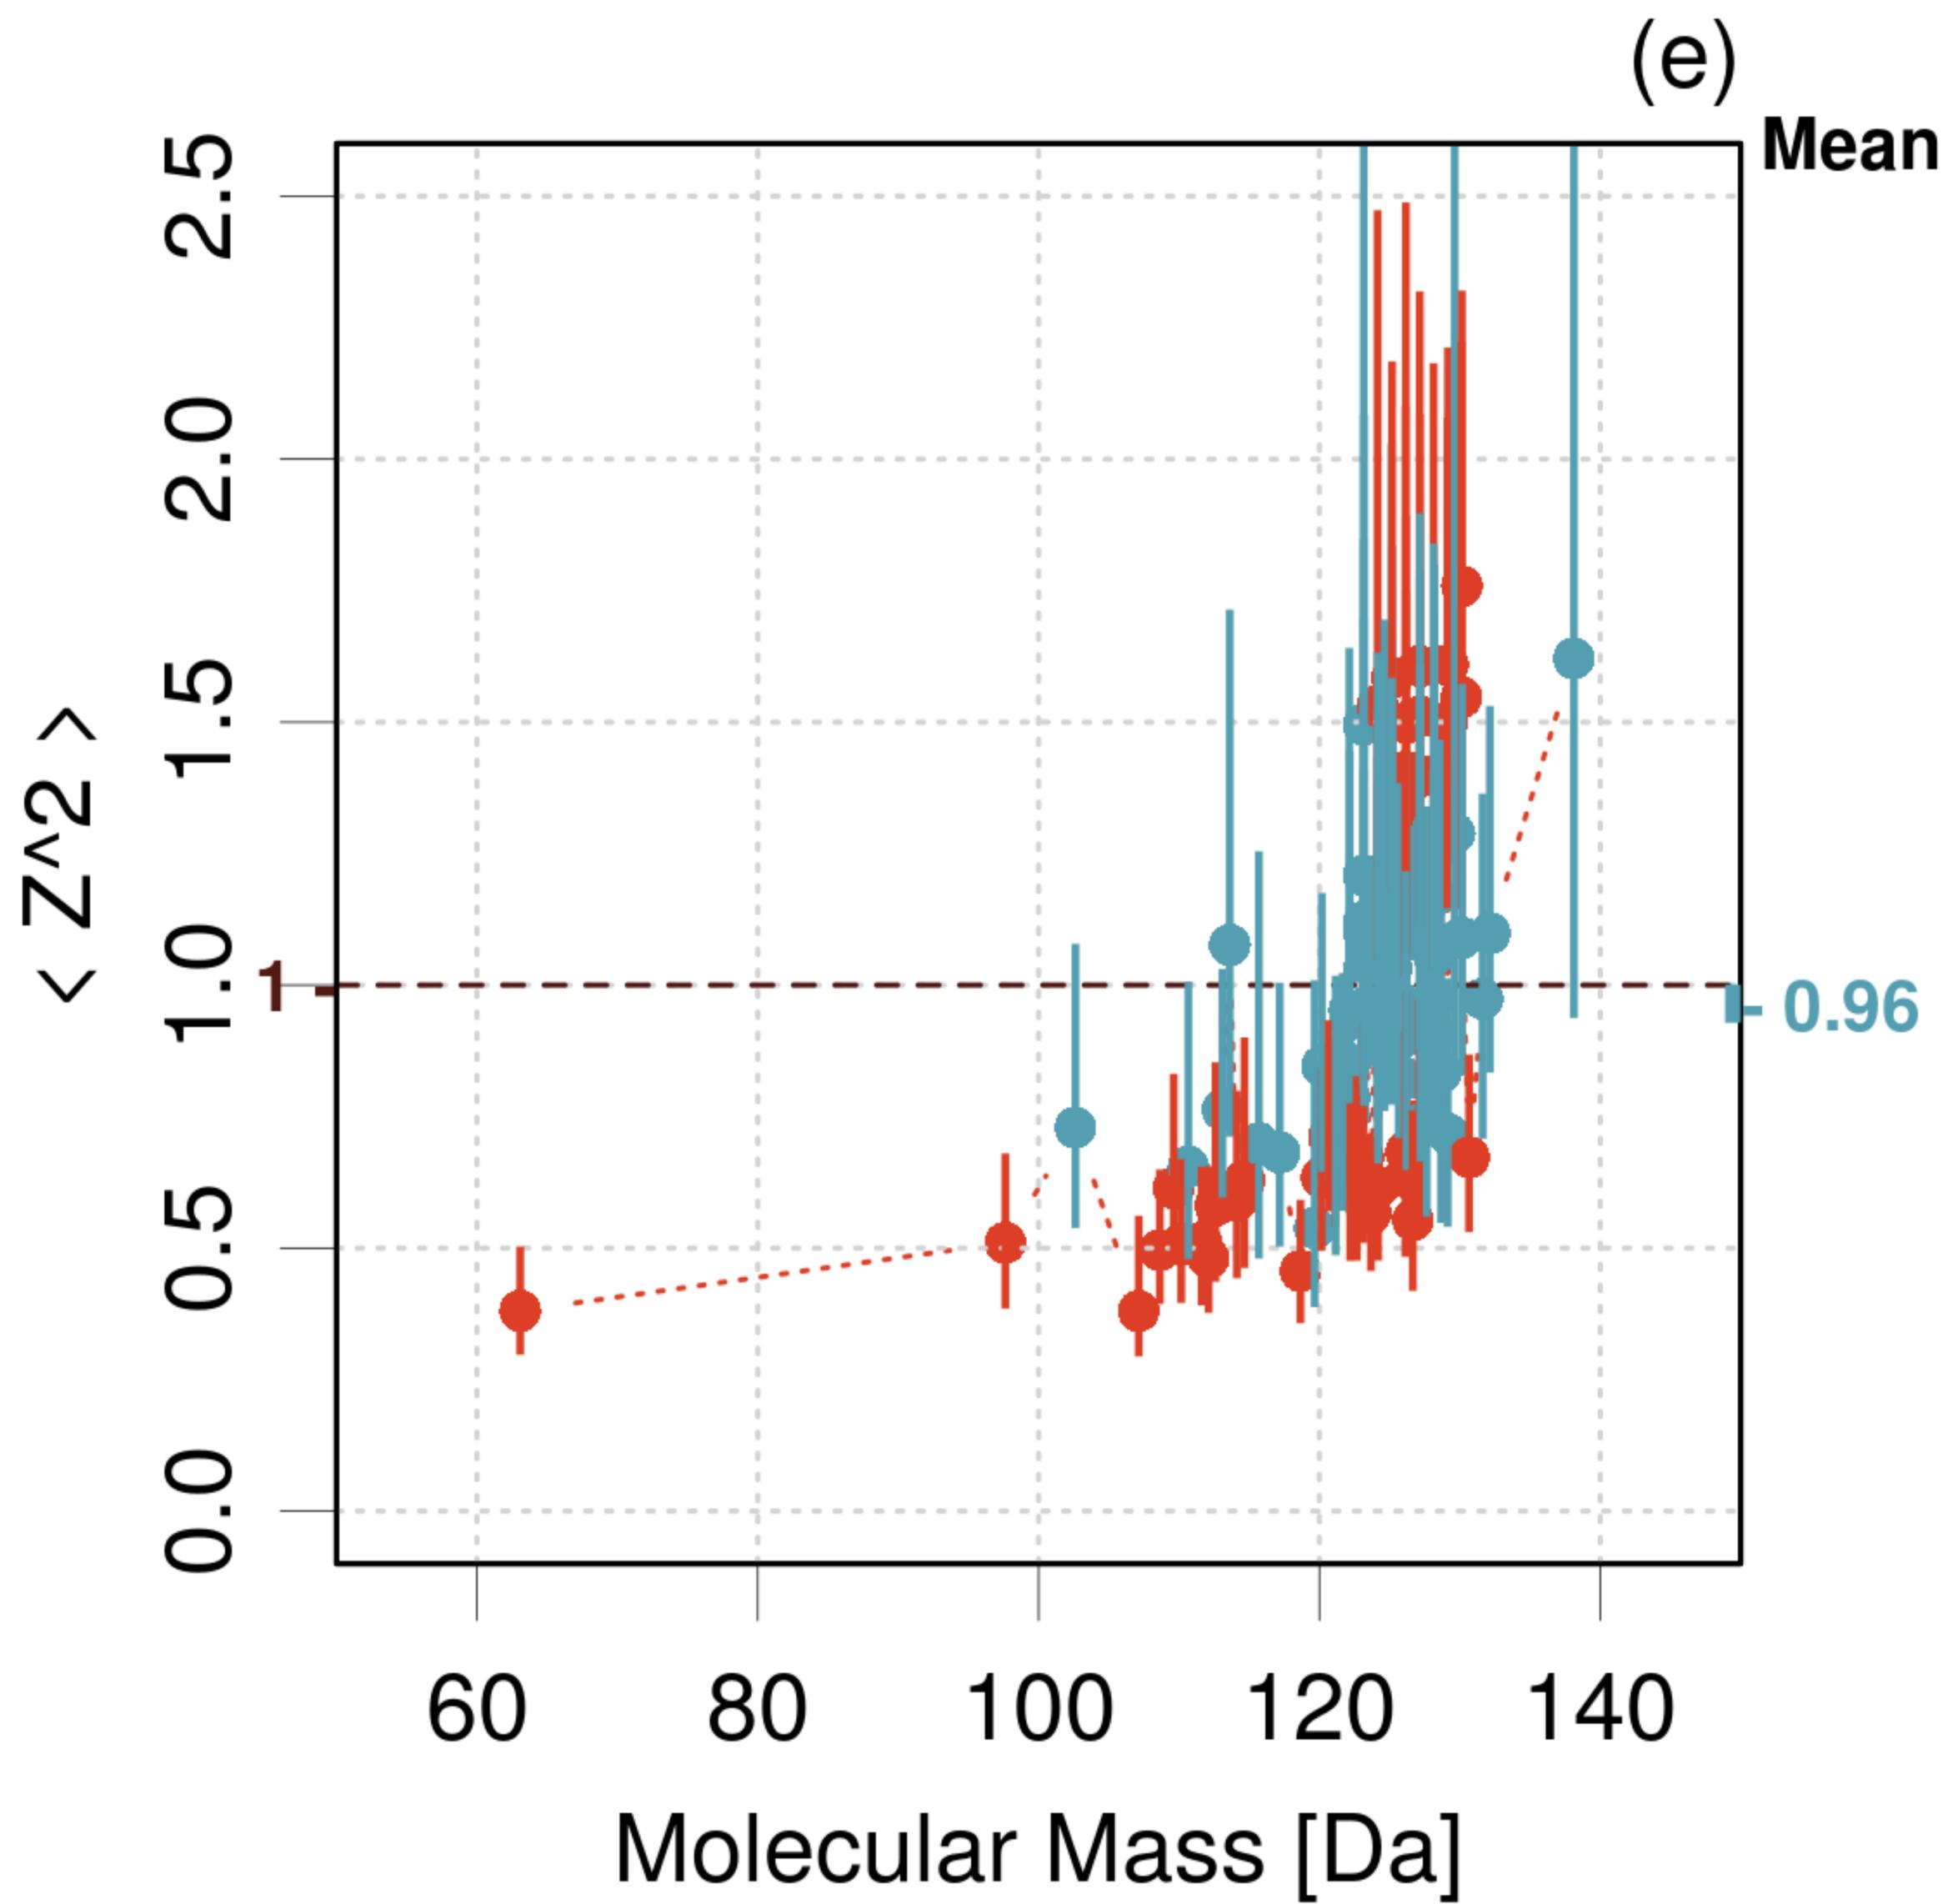


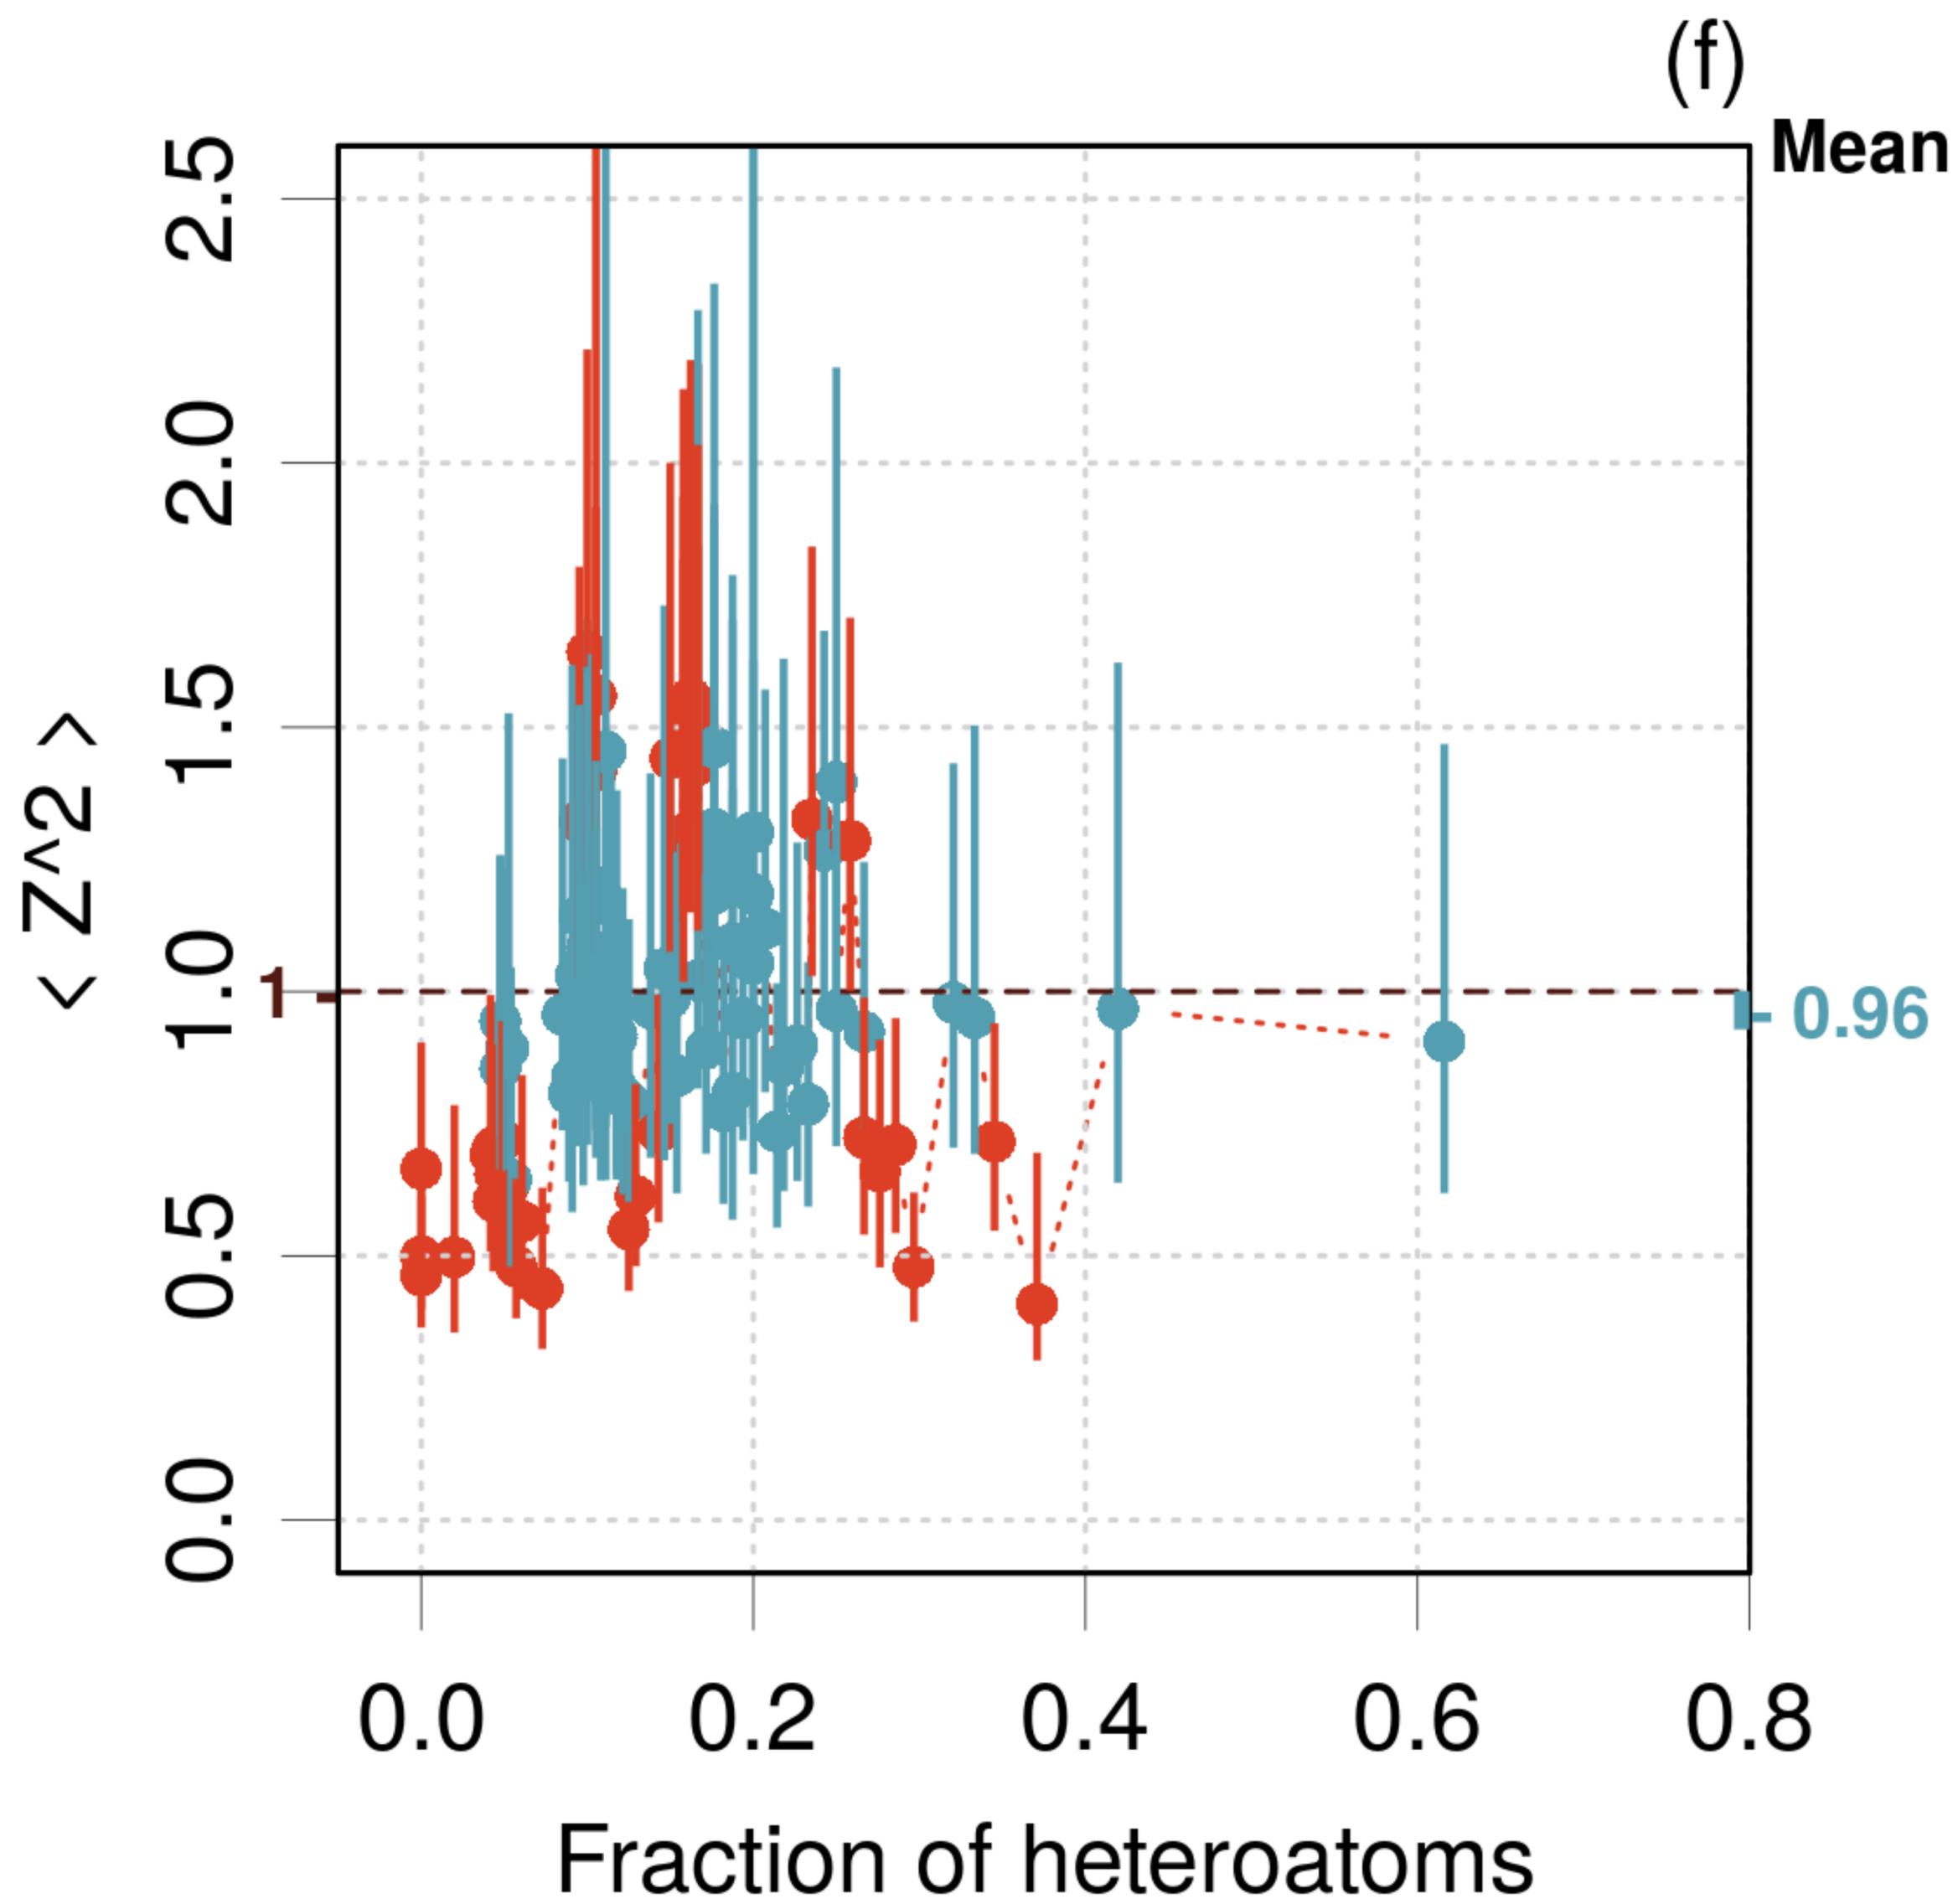




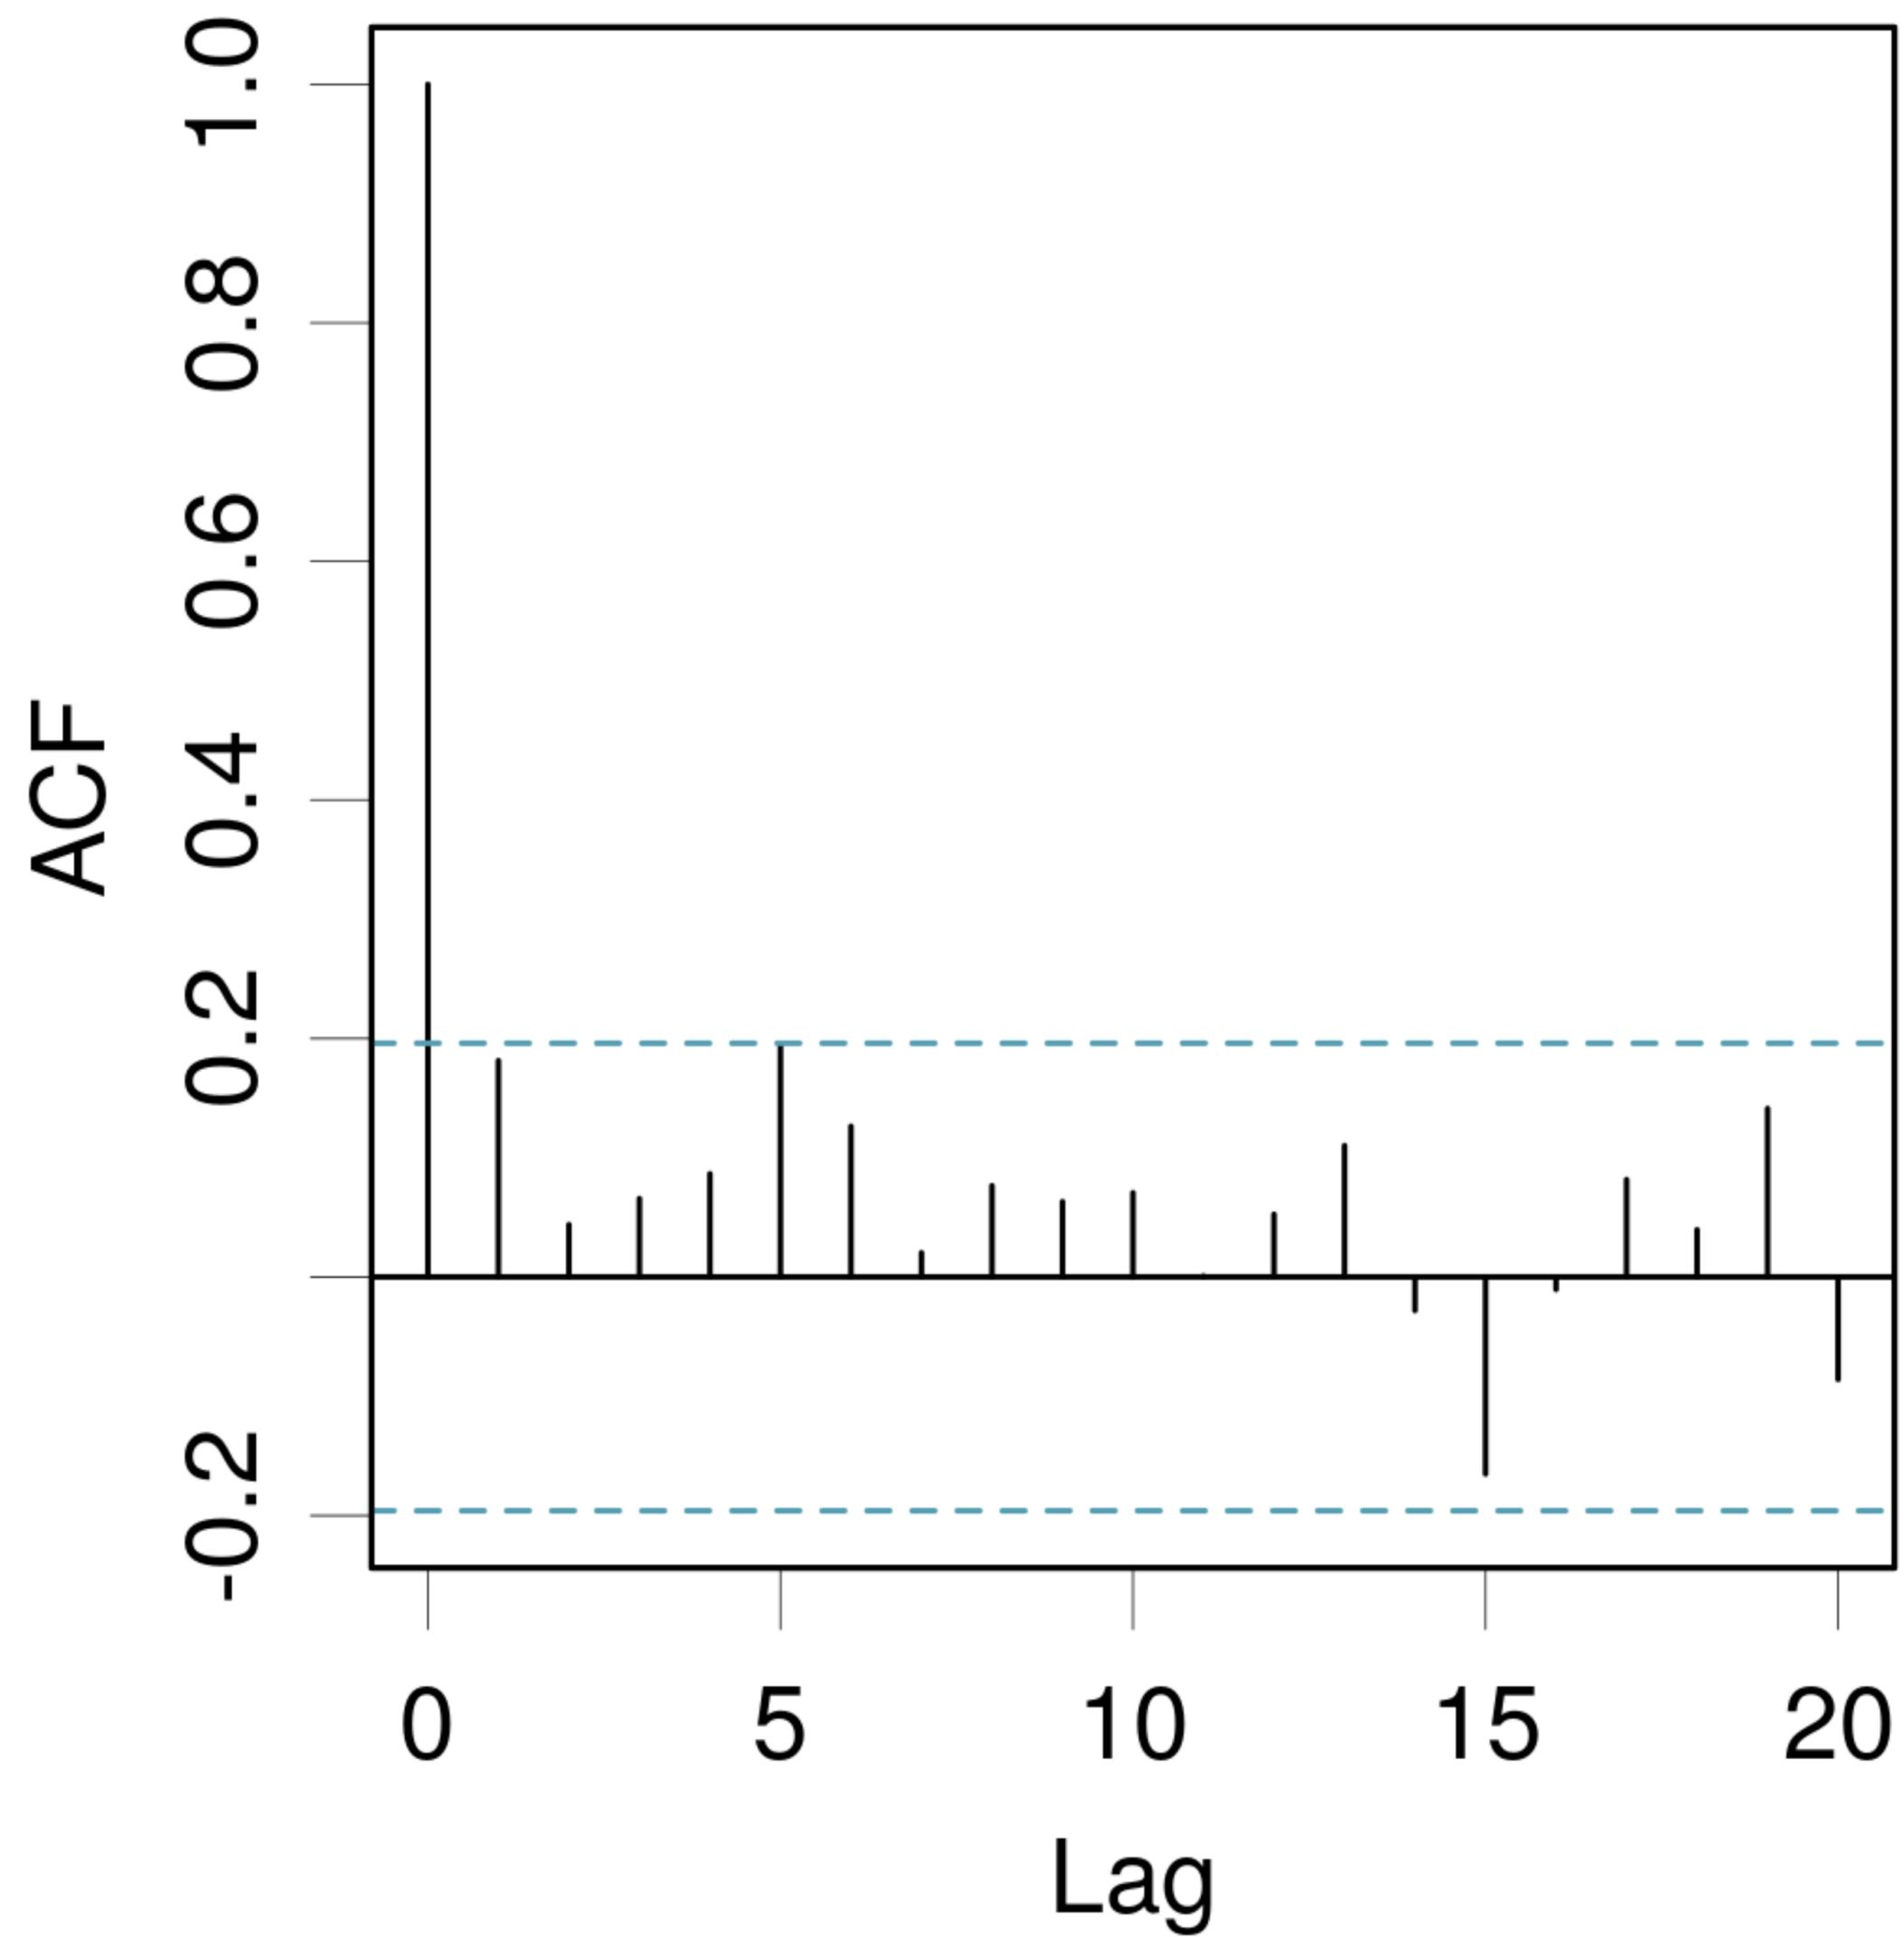




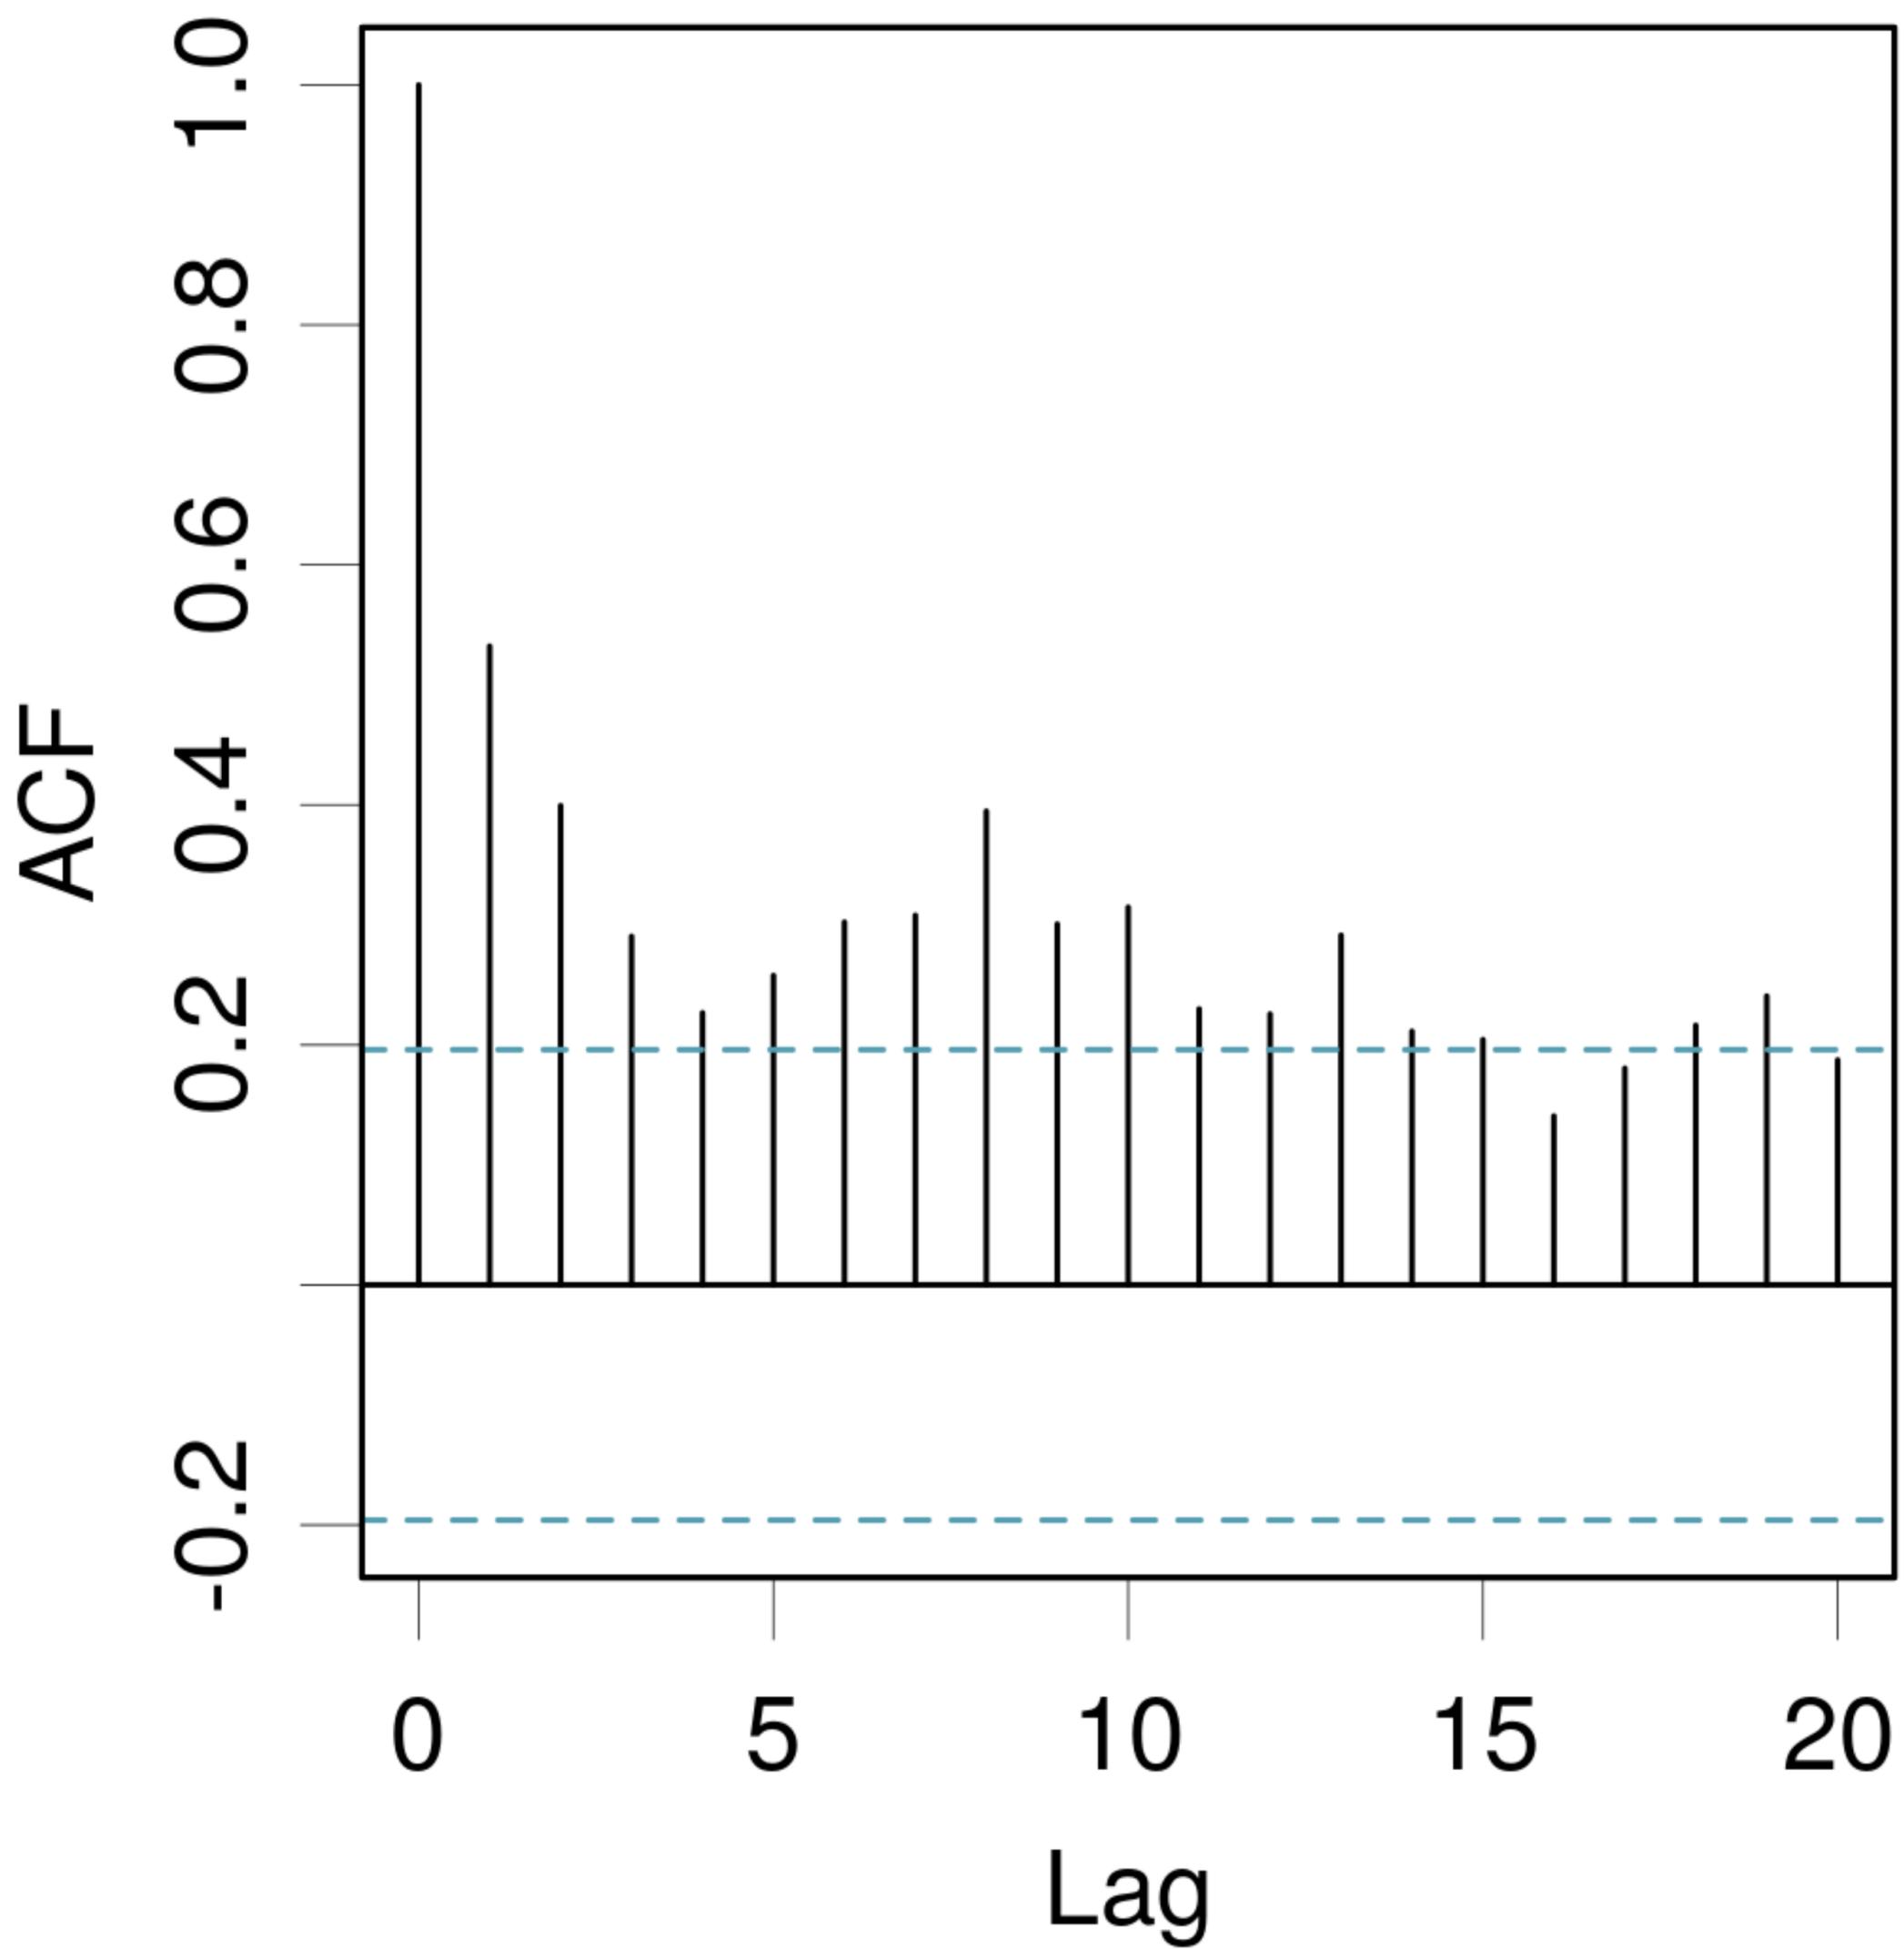




(g)



(h)



(i)

