A novel method for predicting the response variability of friction-damped gas turbine blades

T. Butlin [a,*], P. Ghaderi [a], G. Spelman [a], W. J. B. Midgley [b], R. Umehara [b]

[a] *Cambridge University Engineering Department, Trumpington Street, Cambridge, CB2 1PZ, England*

[b] *Vibration No.3 Laboratory, Vibration Research Department, Research and Innovation Center, Mitsubishi Heavy Industries, Ltd., 2-1-1 Shinhama, Arai-cho, Takasago City, Hyogo, 676-8686, Japan*

Highlights

- A new approach to predict variability of friction-damped turbine blades is presented;
- The method uses Maximum Entropy to estimate a PDF of the friction describing function;
- Computationally expensive Monte Carlo simulations of the nonlinear system are avoided;
- Comparisons with numerical and experimental tests show generally good agreement;
- The results highlight the need for accurate characterisation of uncertainty;
- The method enables rapid estimates of the response distribution.

Abstract

Predicting the response of gas turbine blades with underplatform friction dampers is challenging due to the combination of frictional nonlinearity and system uncertainty: a traditional Monte Carlo approach to predicting response distributions requires a large number of nonlinear simulations which is computationally expensive. This paper presents a new approach based on the principle of Maximum Entropy that provides an estimate of the response distribution that is approximately two orders of magnitude faster than Monte Carlo Harmonic Balance Method simulations. The premise is to include the concept of 'computational uncertainty': incorporating lack of knowledge of the solution as part of the uncertainty, on the basis that there are diminishing returns in computing precise solutions to an uncertain system. To achieve this, the method uses a describing function approximation of the friction-damped part of the system; chooses an ignorance prior probability density function

for the complex value of the describing function based on Coulombs friction law; updates the distribution using an estimate of the mean solution, the admissible domain of solutions, and the principle of Maximum Entropy; then carries out a linear Monte Carlo simulation to estimate the response distribution. The approach is validated by comparison with HBM simulations and experimental tests, using an idealised academic system consisting of a periodic array of beams (with controllable uncertainty) coupled by single-point friction dampers. Comparisons with two- and eight-blade systems show generally good agreement. Predicting the response statistics of the maximum blade amplitude reveals specific well-understood circumstances when the method is less effective. Predictions of the overall blade response statistics agree with Monte Carlo HBM extremely well across a wide range of excitation amplitudes and uncertainty levels. Critically, experimental comparisons reveal the care that is needed in accurately characterising uncertainty in order to obtain agreement of response percentiles. The new method allowed fast iteration of uncertainty parameters and correlations to achieve good agreement, which would not have been possible using traditional methods.

1 Introduction

Underplatform friction dampers are an ideal solution for reducing high amplitude vibration in the harsh operating conditions of gas turbines [1]. However, predicting the response of friction damped systems can be challenging due to nonlinearity and uncertainty: nonlinearity arises predominantly from the frictional interfaces; and there is uncertainty associated with the underlying structural and frictional properties [2].

There is a need for efficient methods that can predict the response of friction-damped structures and which take uncertainty into account, without requiring computationally demanding Monte Carlo simulations of the nonlinear system. There are two main strategies: develop computationally efficient modelling methods so that Monte Carlo studies become feasible (e.g. [3,4]); and / or develop methods for handling uncertainty that require a minimal number of nonlinear simulations (e.g. [5]).

---
* Corresponding author (tel: +44 1223 765237)
  Email addresses: tb267@cam.ac.uk (T. Butlin), pg452@cam.ac.uk (P. Ghaderi), gms41@cam.ac.uk (G. Spelman), w.j.midgley@lboro.ac.uk (W. J. B. Midgley), ryuichi_umehara@mhi.co.jp (R. Umehara).

The Harmonic Balance Method (HBM) has emerged as one of the most successful and efficient methods for predicting the nonlinear response of friction dampers [6]. The response can be predicted highly effectively using HBM with just the fundamental frequency retained [6]: in this case the nonlinearity can be characterised in terms of a describing function. It is noted that more recently there has been a shift towards using the Multi-Harmonic Balance Method (e.g. [7]). The multi-HBM approach has two main benefits: increasing solution accuracy and also allowing prediction of more complicated phenomena that are sometimes observed: for example, it has been found that single harmonic HBM is not able to accurately capture severe events such as friction-damper separation, or multiple equilibrium positions [7]. Nevertheless it remains the case that using a single harmonic is remarkably effective and the associated speed-accuracy trade-off is explicitly investigated in [8].

Despite the intrinsic efficiency of single harmonic HBM, it remains computationally expensive to use it for Monte Carlo simulations for predicting the effect of uncertainty [4]. There is a need for modelling approaches that include the effects of uncertainty explicitly or which reduce the number of nonlinear simulations that need to be run. There have been some interesting recent developments in this general area. In the context of probabilistic uncertainty, Peherstorfer et al. [5] use importance sampling together with a combination of surrogate models and high fidelity models to obtain an efficient estimate of the response statistics. Nevertheless, there is still a need for multiple simulations of a high resolution model and the approach requires specification of two models of differing complexity. Fuzzy arithmetic is an interesting approach to non-probabilistic types of uncertainty, but as described by Moens and Hanss [9] the efficiency of fuzzy arithmetic methods is still limited by the number of simulations needed to estimate response bounds for different levels of uncertainty membership. This is because the response bounds are found by optimisation, or for the upper bound 'anti-optimisation': in other words numerical optimisation is used to search the admissible set of parameters for the extreme responses.

An alternative approach to finding response bounds was presented in [10]: the approach has some similarities with fuzzy arithmetic but rather than specifying bounds in terms of the system parameters, higher level constraints are chosen such that the optimisation loop only requires linear simulations to be carried out. Analytic solutions have also been found for common examples of nonlinearities [11], and the concept of 'equivalent linear bounds' was shown to be useful for response metrics dominated by contributions from the fundamental excitation frequency: the assumptions of equivalent linear bounds are the same as for single harmonic HBM. The 'equivalent linear bounds' approach was applied to friction-damped systems in [12,13], where it was shown to be effective for predicting response bounds of a two- and eight-beam system coupled by a friction damper, with uncertainty associated with the frictional

3

interface. The bounds were shown to be somewhat conservative for typical levels of uncertainty, motivating research into probabilistic methods.

A method that has started to received significant attention is the use of Polynomial Chaos Expansion (PCE): the fundamental theoretical work was developed in [14], but it has only more recently begun to be applied in engineering applications [15]. The core approach is to describe the uncertain input parameters and response distribution in terms of a truncated set of orthogonal basis distributions, then solve the system of equations for the coefficients of the output distribution basis. This can be achieved either by Galerkin projection (referred to as an 'intrusive' method in the sense of changing the system of equations to solve), or by least squares solution using point-wise observations from the original simulation code (referred to as a 'non-intrusive' method). The intrusive methods are computationally faster but more complex to implement [16]. This class of uncertainty propagation method can be very efficient and applicable to nonlinear systems: the approach has been combined with multi-frequency HBM using both intrusive [17] and non-intrusive approaches [18], and also with a nonlinear normal mode framework [19]. The results in each study show a great deal of potential, accounting for uncertainty of strongly nonlinear systems with multi-stable frequency response curves. Nevertheless there are several underlying challenges associated with PCE methods in general: in particular, that simulation time scales poorly with the number of uncertain parameters which currently limits the complexity of system that can be tackled.

Another method that has received increasing attention in the context of structural dynamics is the principle of Maximum Entropy (or 'MaxEnt'). Given limited information about a random variable, Jaynes [20] shows that the appropriate, or least biased, probability density function (PDF) is taken to be the one that maximises its entropy subject to any constraints that incorporate known information, such as a mean value of the PDF. Soize [21] has applied this principle to systems with random matrices. In this approach the system matrices (mass, stiffness and damping) are taken to be uncertain, with their mean represented by deterministic values and probability density function assigned using Maximum Entropy. Response distributions can then be calculated efficiently for linear systems. One challenge with this approach is the difficulty of interpretation: the admissible set of system matrices is not necessarily constrained by physical arguments. For example, the ensemble of possible uncertainties can include systems where remote parts of the system are directly coupled, leading to non-physical results [22]. In itself this is not a failing of Maximum Entropy, rather that the approach does not readily lend itself to including physical constraints. In addition, this approach is mainly focused on uncertainty associated with the underlying linear system rather than nonlinear interactions.

The method presented in this paper applies the principle of Maximum Entropy, but from a different perspective. A novel method is described for estimating the response distribution of friction-damped gas turbine blades subject to a probabilistic parametric description of uncertainty, that only requires one nonlinear simulation (or equivalent computational cost). The method still uses Monte Carlo sampling to propagate uncertainty, but the central concept is to deliberately not compute the nonlinear solution for each sample and instead include 'computational uncertainty' as part of the overall uncertainty. This additional uncertainty means that the results are generally slightly conservative (but not overly so), and only the linear response of the system needs to be computed for propagating the uncertainty. This results in a dramatic reduction in computational cost, without significant loss of accuracy in predictions. The computational uncertainty is included using the principle of Maximum Entropy.

An idealised academic benchmark system has been chosen to test the method, which is described in Section 2: this consists of an experimental test rig with single-point friction couplings to represent the friction dampers (Section 2.1–2.4). Benchmark nonlinear solutions are found using the Harmonic Balance Method (Section 2.5). The framework for estimating the response distribution is described in Section 3 and comparisons with Monte Carlo experiments and simulations are presented in Section 4.

2 Academic benchmark system

An idealised test system has been chosen in order to investigate the fundamental issues associated with friction-damped gas turbine blades within a laboratory setting. The key design requirements of the system were that it should: (1) be periodic to represent the periodicity of bladed disks; (2) include frictional couplings between periodic elements; and (3) allow deliberate introduction of mistuning and other uncertainties. With this starting point, an experimental test rig was designed that consisted of a periodic array of tunable beams coupled by single-point friction dampers. The intention was to choose a system that included some of the key modelling challenges (nonlinear friction and uncertainty), while also retaining simplicity to make the analysis more tractable. This intermediate complexity rig aims to provide a benchmark for assessing new methods in a relatively controlled setting.

Experimental modal analysis was carried out to obtain a modal representation of the dynamic response of the structure. The Harmonic Balance Method was used to predict the deterministic response of the system based on the experimentally identified modal representation, in order to provide a benchmark reference for validation purposes.

A photograph of the overall experimental test rig is shown in Fig. 1(a) where the beams are an idealised representation of turbine blades and will be referred to as Beams 1–8 from left to right. A close up of the friction damper assembly is shown in Fig. 1(b). Figure 2 shows an annotated diagram for (a) a nominal beam and (b) a friction coupling arm. The beams have been water-jet cut from a single sheet of steel to form a comb-like structure of eight nominally identical beams connected at their base for straightforward alignment. The base is bolted between two heavy clamping beams each 25 mm thick to minimise coupling between the beams via the base. Each beam has height 290 mm, width 40 mm, and thickness 3 mm. This results in the dynamic response being dominated by low frequency out-of-plane bending modes. The top of each beam is folded over to provide a horizontal surface of length 31 mm for sliding contact when the beams vibrate out-of-plane. Each beam is coupled by a single-point friction damper acting between the top platforms. In order to maintain approximately constant normal force during large amplitude oscillations, a small steel wedge was secured onto the beam platform to provide the frictional contacting surface at an angle of approximately 3 degrees.

The friction couplings consist of an aluminium arm securely attached to the top of each beam with a hemispherical steel pin contact with the wedge platform on the neighbouring beam. A mass was suspended from a soft elastic bungee in order to provide an approximately constant normal force at the frictional contact. The damper arm is split in the centre, and joined by a thin (0.2 mm) sheet of spring steel, to provide an approximate pinned boundary condition about the bending axis perpendicular to the plane of the rig, while maintaining a high torsional stiffness about the other two axes.

Excitation is provided by a non-contact coil-magnet arrangement at Position (3) (see Fig. 2). The electromagnetic coils (120 m, 21 awg) were fixed to the base structure, and the neodymium cylinder magnets (diameter 10 mm, length 20 mm) were attached to each beam, positioned within the coils. Calibrated measurements of the current input into these coils provide a surrogate for the input force to each beam. The excitation pattern into the set of eight coils could be specified to allow for effective engine order (EO) excitation of the beams. Further instrumentation is provided by accelerometers at approximately the mid-point (Position 2) and near the top (Position 1) of each beam.

There is some variation in the frictional couplings between the beams due to differences in the friction properties and effective coupling stiffness between each pair of beams. Additional controllable mistuning is introduced by masses clamped to either side of the beam (in pairs of 70 g), which can be moved independently on each beam: the irregular cut-out provides 20 discrete notches
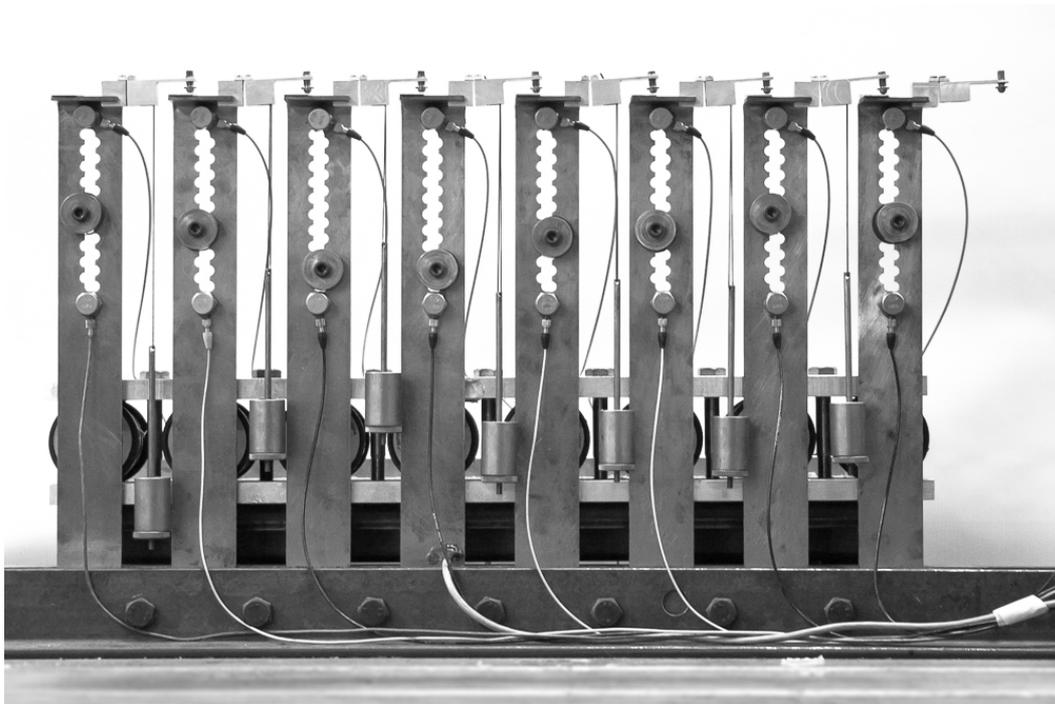
(a)
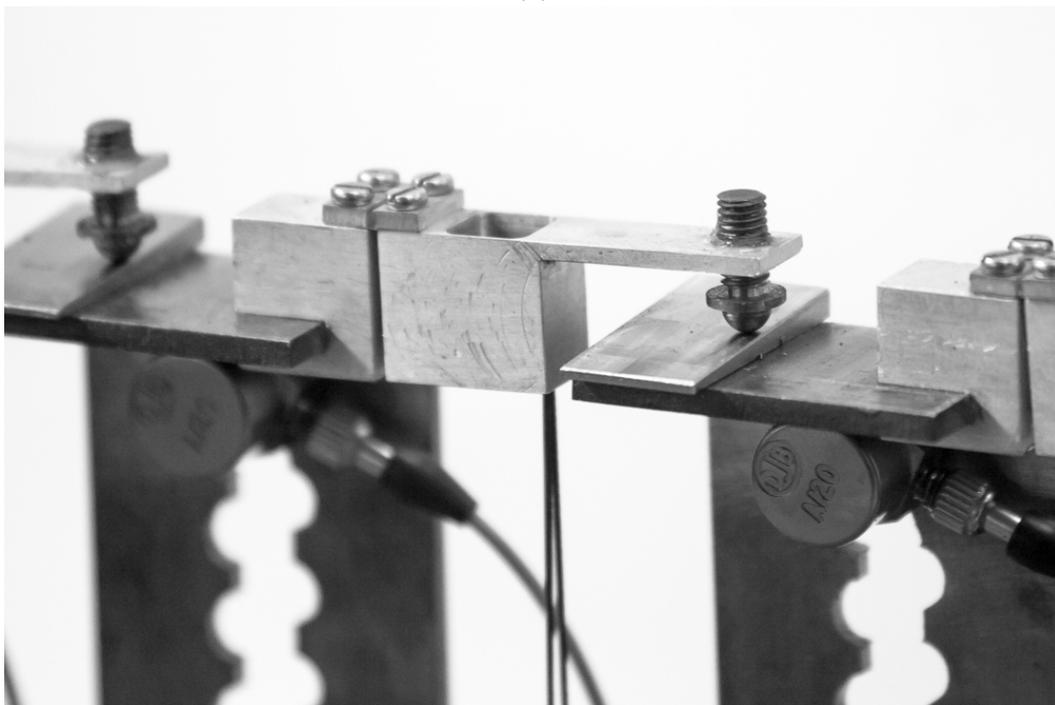
(b)

Fig. 1. Experimental test rig: (a) eight-beam test rig; (b) close up of friction damper assembly
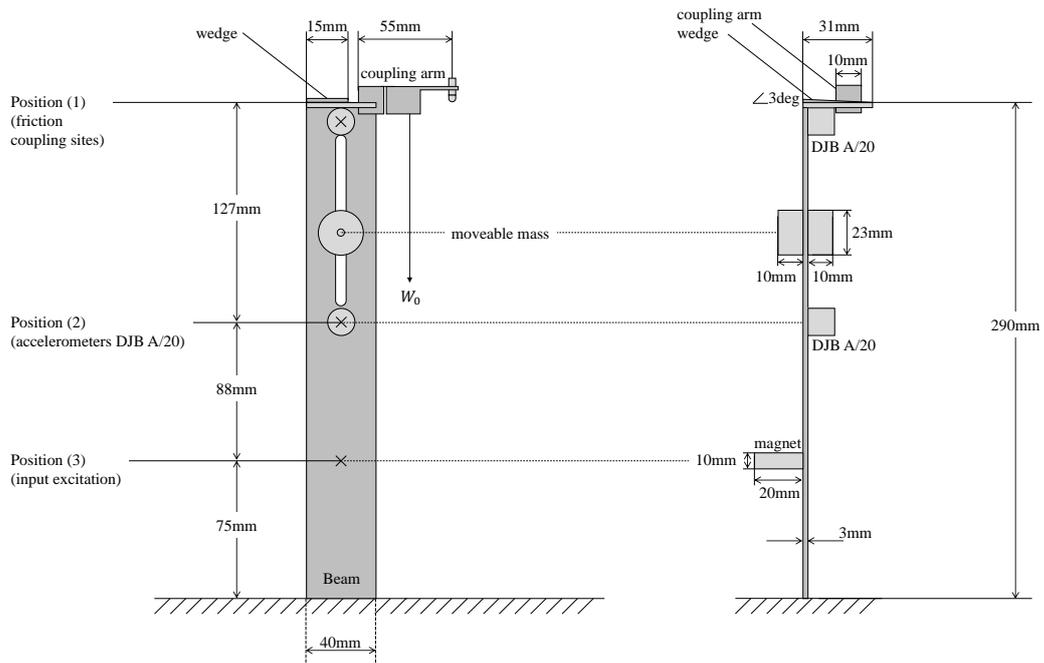
(a)
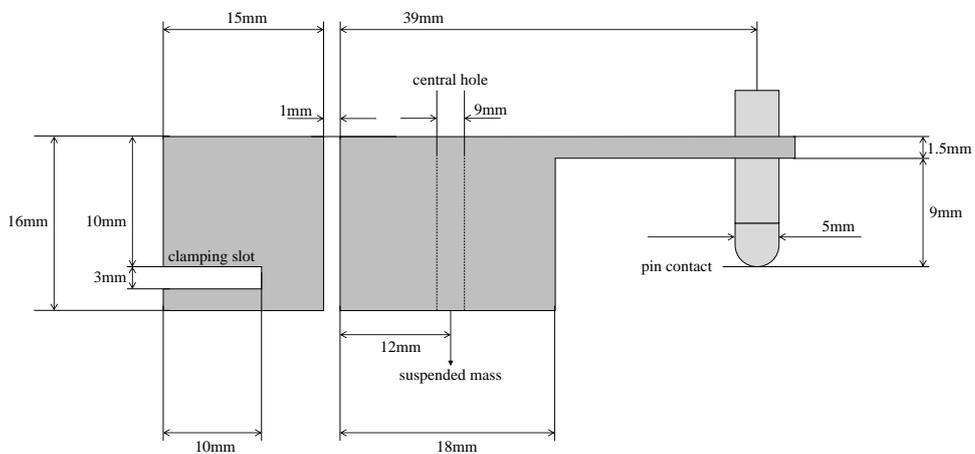
(b)

Fig. 2. Diagram showing dimensions (not to scale) of (a) a single beam within the array and (b) details of a friction coupling arm.

in which the masses can be positioned, numbered 1 to 20 from top to bottom (they will be referred to as tuning locations in this paper). These masses allowed the natural frequencies and mode shapes of each beam to be deliberately perturbed.

The full eight-beam response is rather complicated, so some tests were carried out using a two-beam subsection of the full rig. While this is a simple case, it allows for a controlled discussion of the influence of different parameters before moving onto the full eight-beam system. To achieve this, the dampers connecting the two beams to the remainder of the rig are disconnected and the unused beams are damped out using foam to prevent the dynamics of the unused beams influencing the response.

2.2   Single beam characterisation

A linear characterisation of an isolated beam with friction dampers disconnected (but not removed) was carried out using standard modal analysis techniques (e.g. [23]). For this process the tuning mass was placed in tuning location 8 on the beam (within the set of 1–20 discrete notches) which will be referred to as the 'nominal tuned location' throughout this paper. For these modal analysis measurements the input excitation was an instrumented hammer. The results were calibrated and mass-normalised modal amplitudes were identified. Tables 1 and 2 show the identified modal parameters for Beams 5 and 6 for the first five modes of each beam: frequency; damping factor; and modal amplitude (MA). The modal amplitudes are presented for three key positions on the beam:

(1)   The response of the accelerometer at the beam tip (Position 1).
(2)   The response of the accelerometer at half the beam height (Position 2);
(3)   The excitation location (Position 3);

Table 1
Modal fits for Beam 5 with mass at tuning location 8 (MA refers to the Modal Amplitude).

Mode Count	Type	Frequency (Hz)	Damping ratio	MA_1 $(kg^{-\frac{1}{2}})$	MA_2 $(kg^{-\frac{1}{2}})$	MA_3 $(kg^{-\frac{1}{2}})$
1	Bending	16.20	0.0017	2.55	1.04	0.33
2	Bending	105.8	0.0071	-2.12	1.70	1.16
3	Torsion	151.7	0.006	-0.6	0.32	0.31
4	Bending	275.5	0.017	1.36	0.19	2.82
5	Bending	550.9	0.013	-0.15	2.13	-1.58

Table 2

Modal fits for Beam 6 with mass at tuning location 8 (MA refers to the Modal Amplitude)

Mode Count	Type	Frequency (Hz)	Damping ratio	MA_1 $(kg^{-\frac{1}{2}})$	MA_2 $(kg^{-\frac{1}{2}})$	MA_3 $(kg^{-\frac{1}{2}})$
1	Bending	16.94	0.0021	2.55	0.99	0.48
2	Bending	110.5	0.0032	-1.84	2.18	1.42
3	Torsion	159.2	0.0036	0.27	0.21	-1.04
4	Bending	270.5	0.0068	0.80	0.44	2.5
5	Bending	542.4	0.02	0.06	0.72	-0.68

Modes 1, 2, 4 and 5 correspond to the first four bending modes of each beam. Mode 3 is a torsion mode of the beam at 159 Hz but it only makes a small contribution to the response over the frequency range of interest and the response measurements are made on the beam centre line, so it is not included in subsequent analysis. An example fit to the measured transfer function for a single beam using these parameters is shown in Figure 3, which shows a good fit below 150 Hz and a more approximate fit at higher frequencies. The intention was to accurately model the system for 0-150 Hz (to include the two passbands), without attempting to accurately model frequencies higher than 150 Hz. In practice it turned out that including a fit of just the first two modes does not give good response predictions in the range 0-150 Hz due to the tails of the higher frequency modes, which was corrected by including a very approximate fit of the next two bending modes.

Some peaks are visible near 240 Hz that were not fitted: these are likely to be modes of neighbouring beams in the rig making a small contribution from the weak coupling through the clamped base of the assembly. Attempts to isolate individual beams for these tests were made by damping neighbouring beams with foam inserts, however complete isolation in all cases was difficult to achieve. It was also found that the clamped boundary condition for the beams was not introducing unwanted nonlinearity: releasing the beam from rest at a large amplitude and measuring the decaying response revealed only a very weak dependence of frequency and damping on amplitude (not shown), and no significant harmonics.
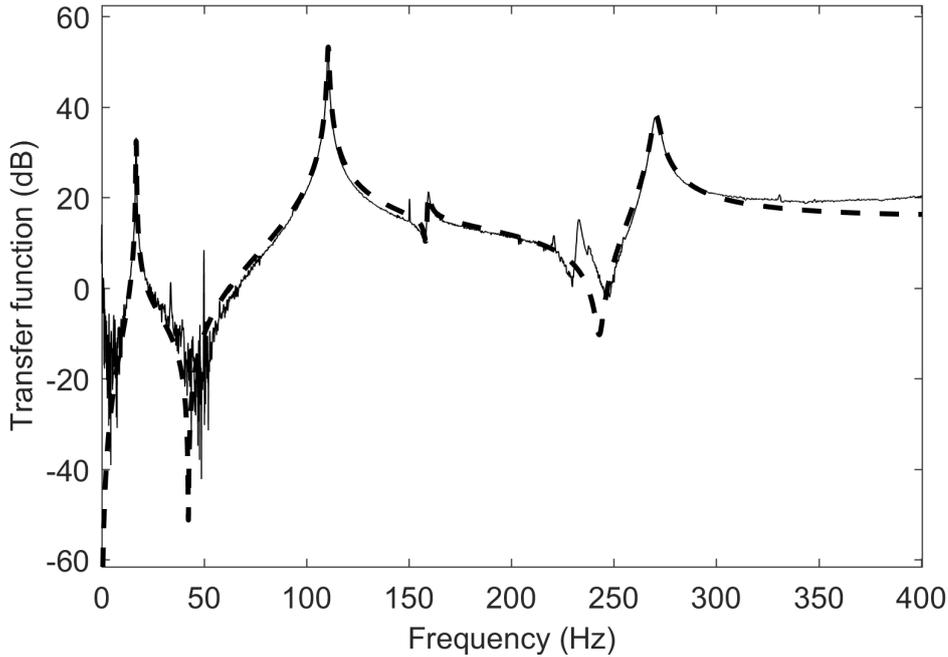
Fig. 3. Example driving point transfer function from input force to output acceleration for Beam 6 at Site 3 (shaker, near base). Solid lines are experimental results with dashed lines showing the reconstructed transfer function from modal parameters

The mistuned rig will be discussed in more detail in Section 2.3, however it is worth highlighting here the increased sensitivity of the modal parameters to any small imperfections between the beams for higher frequency modes of a given beam. In particular the modal parameters for the first and second modes on both beams are fairly similar, but by the fifth mode differ significantly. The natural frequencies show a smaller variation, while the damping shows a much more marked difference between the beams. The variation in damping is likely due to variation in the clamped lower boundary of the beams, which is very difficult to achieve consistently across the width of the rig. The uncertainty resulting from the clamped boundary condition is considered within the overall uncertainty of the beams.

2.3 Mistuning characterisation

The movable masses visible in Fig. 1 allow the natural frequencies and mode shapes for each beam to be varied. As a result there are two sources of variability in the rig:

(i) The natural variability between each beam due to manufacturing imperfections, in both the beam and friction damper properties;

11

(ii) The controllable variability in the beam properties due to the movable masses.

Initially the beam variability from the tunable masses is characterised by a set of linear modal tests, to identify modal amplitudes and natural frequencies as in Section 2.2. The variability in a single beam is considered by modal testing for each tuning location (discrete mass position) of a single beam: in this case Beam 5.

The variability in the first two modes is summarised in Table 3. When using masses to intentionally mistune the rig, the second mode shape is significantly more sensitive to the tuning location than the first mode. This proves useful as it will give two distinct levels of variability within the rig. The higher modes have not been carefully characterised in terms of uncertainty, however estimates suggest the levels of uncertainty are a similar order of magnitude to the second mode.

Table 3
Mean and variation in natural frequency and modal amplitudes (MA) due to tuning mass movement on Beam 5. Modal amplitude subscript corresponds to position number on the beam.

Parameter	Mean	%Std Dev
Mode 1 Natural Frequency (Hz)	16.2	4.7
Mode 1 MA_1 ($kg^{-\frac{1}{2}}$)	0.36	8.0
Mode 1 MA_2 ($kg^{-\frac{1}{2}}$)	1.04	3.8
Mode 1 MA_3 ($kg^{-\frac{1}{2}}$)	2.56	3.4
Mode 2 Natural Frequency (Hz)	103	5.8
Mode 2 MA_1 ($kg^{-\frac{1}{2}}$)	1.05	24
Mode 2 MA_2 ($kg^{-\frac{1}{2}}$)	1.55	25
Mode 2 MA_3 ($kg^{-\frac{1}{2}}$)	-1.98	39

Transfer functions of Beam 5 were measured for each tuning location. This ensemble is summarised in Fig. 4, which shows the percentiles of the measured transfer functions, along with the percentiles of the reconstructed transfer functions using the modal data. The percentiles of the reconstructed transfer functions are not computed directly from the individually fitted cases, but rather using natural frequencies and mode shapes from Gaussian distributions based on the measured mean and variance: this was to simulate that input uncertainty is often specified in terms of a mean and variance rather than as a known distribution. Additionally the higher modes were also randomised using the variances measured for the second beam bending modes, as the variance was not measured above the second mode because the response distribution

over the range 0–150 Hz did not require an accurate characterisation of the uncertainty for modes above 150 Hz. The results of this process provide a good fit over the range from 0–400 Hz. Randomly generated systems drawn from the parameters in Table 3, will be used as a starting point in the MaxEnt verification in Sections 4.2 and 4.4.

(a)   (b)

(c)
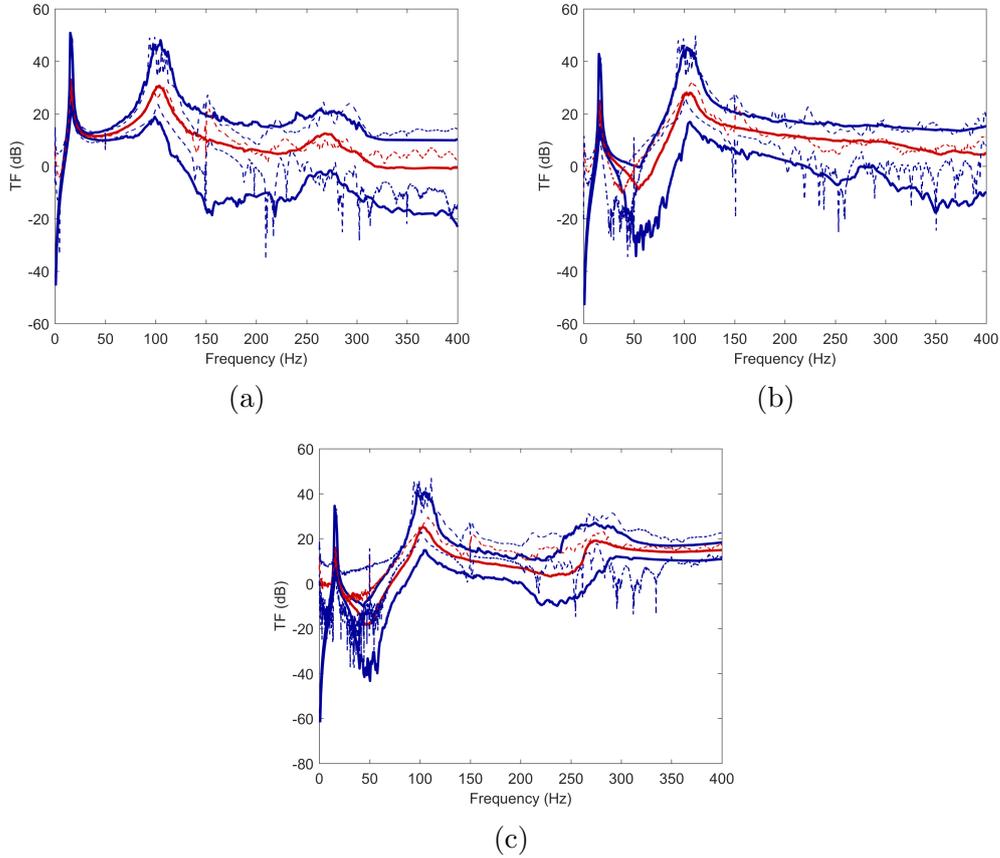
Fig. 4. Percentiles of the ensemble of uncertain transfer functions for Beam 5 at (a) Site 3 (beam tip); (b) Site 2 (mid-beam); and (c) Site 1 (shaker site). Dashed lines are experimental results with solid lines showing reconstructed transfer functions from random realisations of the modal parameters for a measured mean and variance. The three lines represent the 5th, 50th and 95th percentiles.

2.4   Coupling characterisation

To estimate the coupling properties of the friction arms, two parameters are required: the coefficient of friction and the coupling stiffness between the beams. Measurements of the coefficient of friction for steel-on-steel have been carried out on an independent pin-on-disc test rig described in [24], which provided a mean value of $\mu_0 = 0.5$.

The inter-beam tangential coupling stiffness was measured by isolating adja-

13

cent pairs of beams coupled by a single friction damper. The remainder of the dampers were disconnected and the inactive beams were also damped using foam padding. A normal preload $N_0 = 0.96\,\mathrm{N}$ was applied to the active damper with input excitation $F_0 = 0.018\,\mathrm{N}$ which was found to result in friction coupling that was approximately at the linear 'fully sticking' limit. Given that the uncoupled mode shapes of the beams and natural frequencies are known, the frequency spacing between the first and second mode in the first pass band of the coupled pair can be used to determine the coupling stiffness between the beams using a standard linear analysis.

The results of this process for every friction damper are given in Table 4. Damper 5 in Table 4, corresponds to the beam pair that will be used in the 'two-beam' experiments presented in Section 4.2, which couples Beams 5 and 6. The most apparent characteristic of these values is that they are orders of magnitude lower than would be expected of a local contact stiffness at the friction site. This is due to the lowest stiffness in the coupling arm being associated with the steel flexure, which is in series with any contact stiffness. As a result this stiffness becomes the limiting stiffness for the experimental rig. Although it means that the friction model does not include a strict contact stiffness, it fulfills the same role dynamically and does not affect the validation of the MaxEnt method that follows. The physical effect of the low coupling stiffness is that the passbands are slightly narrower than would be otherwise expected.

There is significant variation in the coupling stiffness between the beams, and more than would be expected in a gas turbine. The variation is because each coupling arm is assembled as a collection of parts, and these assembly details are not tightly controlled: the arm is glued onto the folded section of the beams, the spring steel flexure is bolted to the two sections of the arm, and the pin itself is threaded and glued into the end of the arm. However, these factors were simply considered as part of the variability of the test rig and did not limit the validity of the later comparisons.

Table 4
Coupling stiffness parameters between rig beams

Damper Number	1	2	3	4	5	6	7
Coupling Stiffness ($\mathrm{Nm^{-1}}$)	7200	4500	6500	7200	4800	8000	7000

2.5 Harmonic Balance Method

The Harmonic Balance Method (HBM) is commonly used for predicting the response of friction damped turbine beams, e.g. [25]. The solution is approximated as a truncated series of harmonic terms with fundamental frequency

usually chosen to be the input forcing frequency, and the numerical precision of predictions increases with the number of terms included in the expansion [26]. This is a robust and well-understood method, and it has been found that satisfactory (albeit more approximate) predictions can be achieved even when only the fundamental frequency is retained in the expansion [6]. This apparently severe assumption still gives useful results because for friction damped systems it is usually the case that the output displacement response is dominated by the fundamental excitation frequency. The approach represents a form of linearisation and the nonlinear system is characterised by an amplitude and frequency dependent 'describing function'. This first-order approach is used in the present study as the intention is to focus on the effects of uncertainty rather than to achieve high accuracy deterministic predictions, and there are diminishing returns in accurately predicting a specific response in the presence of uncertainty. Nevertheless it is recognised that some kinds of dynamic response require a multi-HBM approach, but these cases fall beyond the scope of the present work.

The frequency-domain system representation is summarised in Fig. 5. The total force $\mathbf{F}$ acting on the beams is the sum of external forces $\mathbf{F}_{\mathrm{ext}}$ and internal nonlinear friction forces $\mathbf{F}_{\mathrm{nl}}$. The linear structural dynamics can be characterised by the Frequency Response Function (FRF) matrix $\mathbf{D}(\omega)$ such that the output response is given by $\mathbf{Y} = \mathbf{DF}$. A subset of these output states $\mathbf{Y}_{\mathrm{nl}}$ are associated with the nonlinear friction dampers, and these states provide the input to the friction describing function $\mathbf{K}(\omega, \mathbf{Y}_{\mathrm{nl}})$. An output metric is defined by a mapping M, and is simply an output quantity of interest chosen by the user.
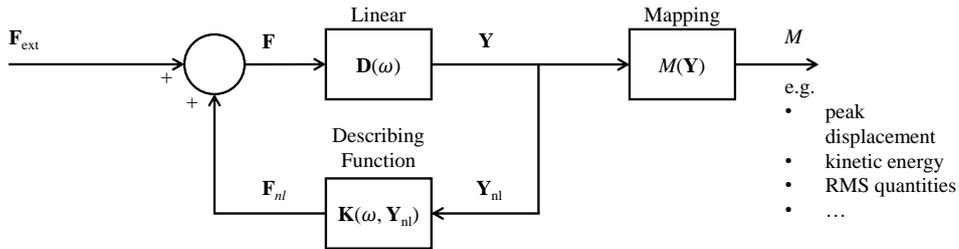

Fig. 5. Summary of friction-damped system representation using describing function $\mathbf{K}$ to characterise the behaviour of friction dampers.

The friction model used is Coulomb's law with a coefficient of friction μ_0 and normal pre-load N_0 in series with a tangential contact stiffness k_c as per the macro slip model of [6]. A schematic of the model is shown in Fig. 6. This is a very simple model, and there are two main reasons that we do not

15

include other friction models in this study: (1) the focus is on the approach to uncertainty modelling, which is demonstrated using Coulomb friction with contact stiffness, but is not intrinsically limited to this choice; and (2) it is a commonly used model in this context, where it often serves as the underlying law within distributed-contact Finite Element models.

It is convenient to define a non-dimensional parameter s:

$$s \equiv \frac{\mu_0 N_0}{k_c A}. \tag{1}$$

where A is the relative displacement amplitude between a given pair of contacting surfaces, i.e. $y_{n+1} - y_n = A \cos \omega t$.
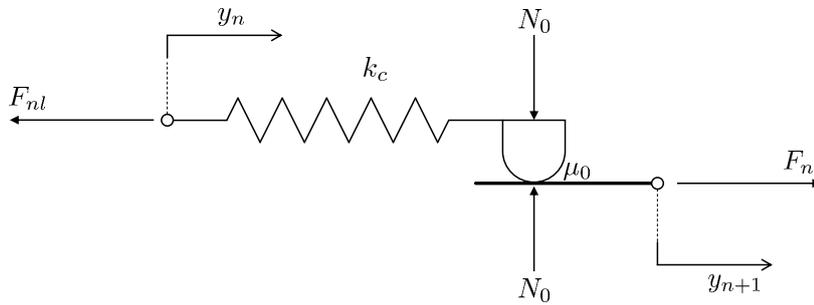

Fig. 6. Schematic of the friction model used for a single damper, consisting of a contact stiffness in series with a Coulomb friction sliding contact.

The describing function K for a single damper can be written as a function of s:

$$K = K_r(s) + \mathrm{i} K_i(s), \tag{2}$$

where

$$K_r = \frac{k_c}{\pi} \left(\arccos\left(1 - 2s\right) - 2(1 - 2s)\sqrt{s(1-s)} \right) \tag{3}$$

$$K_i = \frac{4k_c}{\pi} s\left(1 - s\right), \tag{4}$$

which is valid for $0 < s \le 1$. When $s \to 0$ then $K \to 0$ which corresponds to large relative amplitude (the fully slipping limit), and when $s = 1$ then $K = k_c$ which corresponds to small relative amplitude (the fully sticking limit). These results were derived in [6], and the expressions presented here use notation consistent with the present study.

Figure 7 shows two projections of the describing function K using parameters based on the experimental test rig described in Section 2.4: $k_c = 4800 \, \mathrm{Nm}^{-1}$,

16

$\mu_0 = 0.5$ and $N_0 = 0.48\,\text{N}$. Four points have been identified by a '$\times$' indicating relative beam amplitudes of 10^{-5}, 10^{-4}, 10^{-3}, and $10^{-2}\,\text{m}$. When a friction damper between a pair of beams is fully sticking, the relative displacement between beams is small and the coupling is given by a purely real value equal to the contact stiffness k_c. As the relative amplitude increases, the friction dampers begin to slip, leading to a complex value of K which follows the arc shown. At high amplitudes, the friction coupling becomes insignificant and $K \to 0 + 0\text{i}$. Although friction is still present, its relative effect on the system dynamics starts to become insignificant.


(a)

(b)
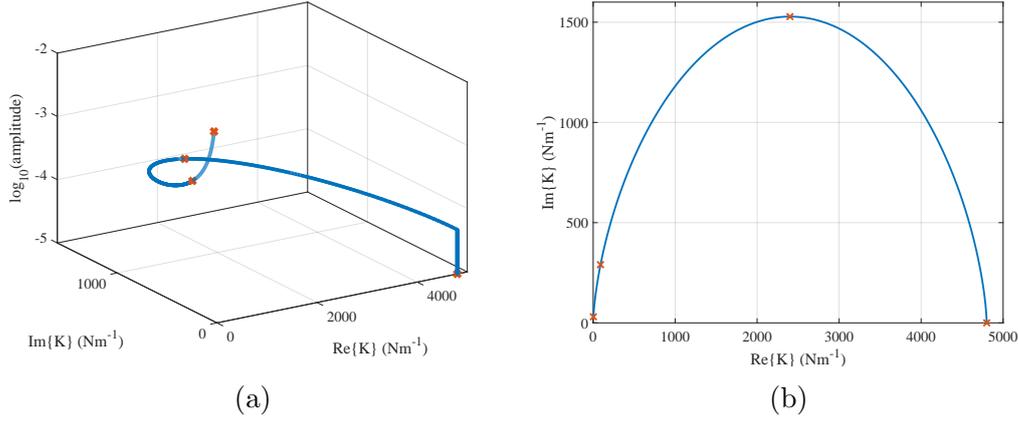
Fig. 7. Complex value of describing function K: (a) real and imaginary parts of K_i as a function of relative amplitude between beam pairs; and (b) projection of real and imaginary parts of K.

Eq. (2) represents the describing function for a single damper, so the response of N dampers is parameterised by the vector $\mathbf{s} = [s_1 \; s_2 \; \cdots \; s_N]^T$ corresponding to the diagonal matrix of describing functions $\mathbf{K}$ (see Fig. 5). The solution procedure adopted was to:

(1) guess a vector of values $\mathbf{s}_{\text{trial}}$
(2) compute the corresponding diagonal describing function matrix $\mathbf{K}$ using Eq. (2)-(4):

$$\mathbf{K} = \begin{bmatrix} K_1 & 0 & \cdots & 0 \\ 0 & K_2 & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & K_N \end{bmatrix} \quad (5)$$

(3) find the response $\mathbf{Y} = (\mathbf{I} - \mathbf{DK})^{-1}\mathbf{DF}_{\text{ext}}$
(4) compute $\mathbf{s}_{\text{out}}$, with $s_{\text{out},j} = \mu_{0,j}N_{0,j}/k_{c,j}A_j$, and A_j being the relative displacement amplitude across the jth friction damper computed from $\mathbf{Y}$
(5) iterate until the residual difference is zero: $\mathbf{R} \equiv \mathbf{s}_{\text{trial}} - \mathbf{s}_{\text{out}} = 0$.

The numerical solution was found using Matlab's `fsolve` function.

The HBM model is used as a benchmark reference model to valide the MaxEnt method, described in the next section. For any given comparison, the benchmark response percentiles were computed using an ensemble size of $N = 1000$ simulations. This was empirically found to give adequately converged estimates for the 5th and 95th percentiles of the response.

2.6 Validation of HBM with experimental results

The linear structural dynamics are well understood and accurately characterised over the bandwidth of interest of 0–150 Hz. Introducing the frictional coupling makes predictions much more challenging due to both nonlinearity and uncertainty. As a starting point, two beams (5 and 6) were effectively isolated from the rest of the assembly, rather than starting with the full eight-beam friction-coupled test rig. To isolate this pair of beams, all the friction dampers were disconnected except for the one connecting Beams 5 and 6. This effectively decoupled the other beams because the base was tightly clamped by a thick steel beam. However, some small residual coupling was observed (due to finite stiffness of the clamping structure), so damping was added using foam inserts to further reduce the effect of the other beams. This was found empirically to be sufficient to observe the expected behaviour for a two-beam system.

A sinusoidal input force excitation was applied to Beam 5 at Position (3) and the output acceleration response was measured at Position (2). A slow continuous frequency sweep was used for the input spanning 10–150 Hz. Convergence checks on the rate of change of frequency were carried out to ensure that unwanted transient effects were not significant (not shown). Figure 8 shows a comparison between HBM predictions and experimental results for (a) input excitation amplitude $F_0 = 0.18$ N and (b) $F_0 = 0.72$ N. The dashed line shows the experimental results and the solid line is the HBM prediction, showing the response for the driven beam. The coefficient of friction was independently measured to be $\mu_0 = 0.5$. It can be seen that overall there is reasonable agreement, with very good agreement in the range 0–100 Hz. This range includes the first two peaks at 17 Hz (beams in-phase) and 40 Hz (beams out-of-phase) corresponding to the first bending mode of the beams, i.e. the first passband of the coupled system. As expected the friction damper does not significantly affect the in-phase mode (17 Hz) where there is theoretically no relative motion between the beams, but it has an increasing affect on the out-of-phase mode (40 Hz) where there is significant relative motion between the beams.

The level of agreement is more approximate in the second passband in the range 100-150 Hz: the two peaks again correspond to the in- and out-of-phase modes, in this case for the second bending mode of the beams. The reason for the larger discrepancy here is likely to be because the second modes of the beams are not tuned as accurately as the first modes of the beams (due to manufacturing and assembly details), while the HBM model assumes identical beams for these tests. These are details that could all be characterised in more detail, but the emphasis of this study is on predicting the response distribution arising from uncertainty rather than high fidelity modelling of the deterministic components, so correction of these details has not been carried out in this study.
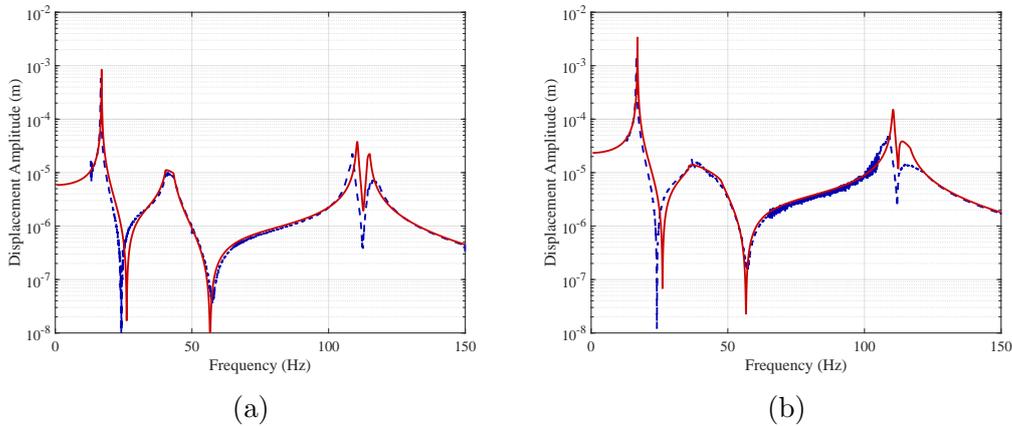

Fig. 8. Two-beam sub-assembly: comparison of HBM predictions (solid) with experimental swept-sine tests (dashed): (a) input force amplitude $F_0 = 0.18$ N; and (b) $F_0 = 0.72$ N.

A similar comparison was carried out for the full eight-beam assembly. In this case the excitation was applied to all eight-beams (rather than just one) in a pattern corresponding to the experimentally identified third passband mode of the assembly: this approximately corresponds to Engine Order 2 excitation, which has two nodal diameters and theoretically should only excite the corresponding passband mode for a perfectly tuned assembly.

Figure 9 shows the eight-beam comparison between HBM predictions and experimental results. In order to simplify the figure the maximum response across the eight beams is shown at any given frequency. It can be seen that the results are qualitatively in broad agreement: peaks are at approximately the correct amplitudes, and the passband modes are more lightly damped for the lower excitation in (a) than in (b). However, there are significant differences in the details: this is wholly representative of how difficult it can be to obtain good agreement for assembled structures with frictional interfaces, and motivates the need for an approach that incorporates uncertainty intrinsically.
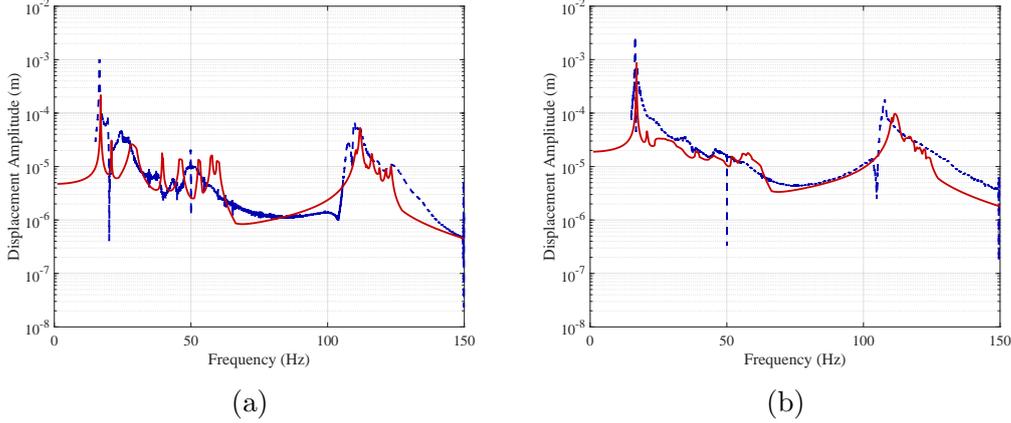
Fig. 9. Eight-beam full assembly: comparison of HBM predictions (solid lines) with experimental swept-sine tests (dashed): (a) input force amplitude $F_0 = 0.18\,\mathrm{N}$; and (b) $F_0 = 0.72\,\mathrm{N}$.

3 Theoretical framework of MaxEnt approach

This section presents the MaxEnt method that has been developed, together with the supporting theoretical results that underpin the approach.

3.1 Overview of Method

The new approach is based on the assertion that the computational cost of finding the exact solution for each sample of a Monte Carlo simulation is too great, and that expending computational effort on finding exact solutions for a system with uncertain parameters leads to diminishing returns, i.e. there is little point finding accurate solutions to a system that is imprecisely defined. Therefore, the choice is made to not attempt to solve the full nonlinear equations for each Monte Carlo sample and instead include computational uncertainty alongside structural and frictional uncertainty by sampling the describing function $\mathbf{K}(\mathbf{s})$ using a suitable probability density function (PDF). The phrase 'computational uncertainty' is defined here to be the uncertainty introduced by not knowing (or choosing not to compute) the HBM solution to the system shown in Fig. 5. Specifically this is represented as a PDF in the non-dimensional parameter $\mathbf{s}$ defined in Eq. 1 which determines the operating point of each friction damper. The core idea is that the HBM solution is itself unknown and is therefore uncertain, because it is infeasible to compute (for many samples in a Monte Carlo simulation).

The PDF is estimated using the principle of Maximum Entropy: this is a two-step process that requires identifying an 'ignorance prior' that represents complete lack of knowledge of the solution, then updating the prior subject to

20

additional information (the mean solution) using MaxEnt. The ignorance prior is a subjectively chosen PDF that aims to characterise no prior knowledge. For example, the ignorance prior distribution for throwing a dice is usually taken to be a uniform distribution across the six possible outcomes; if many experiments were carried out it may be discovered that the mean solution was not consistent with this, and it would be possible to apply MaxEnt to update the PDF with this additional information.

The initial choice of the PDF for the ignorance prior of complex values of the describing function $\mathbf{K}(\mathbf{s})$ falls to engineering judgement. In [12] a uniform distribution was applied to the real and imaginary parts of each K_i over an admissible space which allowed the uncertainty to include friction-model uncertainty. But it has been found that Coulomb's friction law appears to be a useful approximation of friction-dampers, which specifies a different constraint that each $K_i(s_i)$ must lie on a defined curve in the complex $\mathbf{K}$-space. An example of the describing function was shown in Fig. 7 using parameters based on the experimental test rig. The describing function is conveniently parameterised by the non-dimensional variable s which takes the range $0 < s \leq 1$ describing fully slipping through to fully sticking.

The 'ignorance prior' PDF $m(s_i)$ for a given damper is chosen based on the 'Principle of Indifference' (see Jayne's [20]), which for continuous distributions leads to a uniform distribution. Therefore the ignorance prior $m(s_i)$ is chosen to be a uniform distribution over the normalised operating point s_i of each damper such that $0 < s \leq 1$ where $s_i = 0$ and $s_i = 1$ correspond to slipping and sticking limits respectively.

Additional information can be used to update the PDF of the normalised operating point s_i of the describing function for each damper: a force saturation limit that discards any solutions that exceed the Coulomb friction force; and an estimate of the mean solution which can be found in several ways, e.g. by using the solution from one HBM simulation based on a nominal or 'mean' set of model parameters.

The force constraint effectively defines the admissible domain of the PDF of s. The 'mean' estimate provides a set of $s_{i,\mathrm{mean}}$ values that can be used to update the ignorance prior using Maximum Entropy. The underlying principle of MaxEnt is that it provides a method for updating an ignorance prior using new information in a way that introduces no additional bias, in other words it includes additional information in a way that makes no extra assumptions. When the mean is known then the effect of applying MaxEnt is to re-weight the ignorance prior by an exponential function over the admissible domain, as shown in the next section. This is readily implemented by assigning weightings to the original samples drawn from the ignorance prior.

The proposed method is summarised in Fig. 10 and is as follows: generate an uncorrelated ensemble of describing function values specified by the normalised operating point s_i for each damper; generate a parallel ensemble of uncertain parameters from a set of input distributions that describe the system uncertainty; combine the samples pairwise and discard those that exceed the Coulomb friction force limit; estimate the mean solution $s_{i,\mathrm{mean}}$, e.g. using an HBM simulation of the 'mean' model; re-weight the samples using an exponential distribution to achieve a mean set of $s_{i,\mathrm{mean}}$ as given by the mean estimate; and propagate each sample to an output metric of interest.
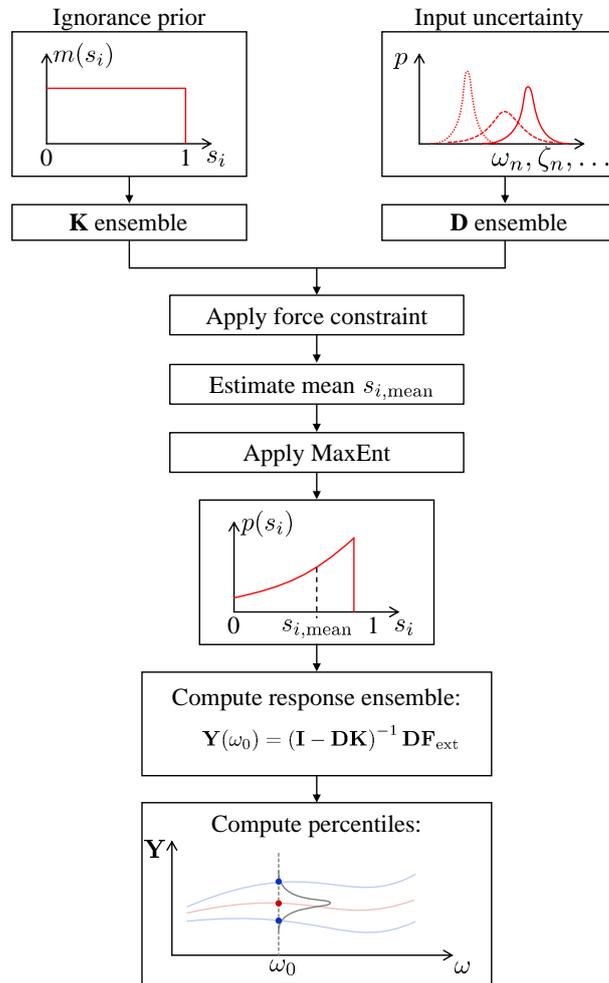

Fig. 10. Summary of MaxEnt approach.

3.2 Maximum Entropy

The principle of Maximum Entropy for a continuous random variable is based on maximising the relative 'information entropy' H:

$$H(p, m) = -\int_{\mathbf{D}} p(\mathbf{s}) \log \frac{p(\mathbf{s})}{m(\mathbf{s})} d\mathbf{s} \tag{6}$$

where $p(\mathbf{s})$ is the multivariate probability density function that is sought, $\mathbf{D}$ is the domain of admissible values of $\mathbf{s}$, and $m(\mathbf{s})$ is the 'ignorance prior' that is required to satisfy invariance of H under a change of variables $\mathbf{s} \to f(\mathbf{s})$.

If the mean normalised operating point $s_{i,\mathrm{mean}}$ of the describing function of each damper is specified then there is a constraint function for each friction damper:

$$\int_{\mathbf{D}} s_i p(\mathbf{s}) d\mathbf{s} = s_{i,\mathrm{mean}} \tag{7}$$

and an overall normalisation constraint:

$$\int_{\mathbf{D}} p(\mathbf{s}) d\mathbf{s} = 1 \tag{8}$$

Equation 7 defines the marginal expectation $s_{i,\mathrm{mean}} \equiv E[s_i]$, the integration is taken to be over the region $\mathbf{D}$ (valid values of each s_i parameter) and $d\mathbf{s} \equiv ds_1 ds_2 \ldots ds_i \ldots ds_{N-1} ds_N$.

The problem of maximising Eq. 6 subject to Equations 7 and 8 can be expressed as a variational problem using Lagrange multipliers λ_n to enforce the constraints. This is achieved by defining a functional U as follows:

$$U = H(p, m) + \lambda_0 \left[\int_{\mathbf{D}} p(\mathbf{s}) d\mathbf{s} - 1 \right] + \sum_{i=1}^{N} \lambda_i \left[\int_{\mathbf{D}} s_i p(\mathbf{s}) d\mathbf{s} - s_{i,\mathrm{mean}} \right] \tag{9}$$

$$= \int_{\mathbf{D}} F(p, \mathbf{s}) d\mathbf{s} - \lambda_0 - \lambda_i s_{i,\mathrm{mean}} \tag{10}$$

where N is the number of friction dampers and $F(p, \mathbf{s})$ is defined to be:

$$F = -p(\mathbf{s}) \log \frac{p(\mathbf{s})}{m(\mathbf{s})} + \lambda_0 p(\mathbf{s}) + \sum_{i=1}^{N} \lambda_i s_i p(\mathbf{s}) \tag{11}$$

The functional U does not depend on the derivative of p, so the maximum occurs when:

$$\frac{\partial F}{\partial p} = 0 \tag{12}$$

$$\frac{\partial U}{\partial \lambda_0} = 0 \tag{13}$$

$$\frac{\partial U}{\partial \lambda_i} = 0 \tag{14}$$

These three equations give:

$$-\log \frac{p(\mathbf{s})}{m(\mathbf{s})} - 1 + \lambda_0 + \sum_{i=1}^{N} \lambda_i s_{i,\text{mean}} = 0 \tag{15}$$

$$\int_{\mathbf{D}} p(\mathbf{s}) \mathrm{d}\mathbf{s} - 1 = 0 \tag{16}$$

$$\int_{\mathbf{D}} s_i p(\mathbf{s}) \mathrm{d}\mathbf{s} - s_{i,\text{mean}} = 0 \tag{17}$$

The Maximum Entropy PDF $p(\mathbf{s})$ is therefore:

$$p(\mathbf{s}) = m(\mathbf{s}) \mathrm{e}^{\lambda_0 - 1} \prod_{i=1}^{N} \mathrm{e}^{\lambda_i s_i} \tag{18}$$

The factor $\mathrm{e}^{\lambda_0 - 1}$ is simply a normalising constant such that the total integral is unity, and the Lagrange multipliers λ_i are exponential weightings such that the marginal expectations of each describing function parameter s_i are equal to the HBM mean solution.

3.3 Methods for estimating the mean solution

Conceptually the simplest method for obtaining an estimate of the mean solutions s_i is to assume that it is given by the HBM solution of the 'mean model', i.e. the system with nominal parameter values. However, it could be argued that this is a poor choice because it may represent an extreme case such as a perfectly 'tuned' set of turbine beams. This will not result in a good estimate of the mean solution of an ensemble of mistuned systems. A better estimate could be obtained by taking a subset of HBM simulations and computing the mean from that: but this comes with a computational cost.

By way of illustration Fig. 11(a) shows a set of mean s-values calculated using $N = 1$, 10 and 1000 HBM simulations: this is for a representative two-beam test system with an intermediate level of forcing $F_0 = 0.3\,\text{N}$. The frequencies where $s = 1$ correspond to the damper fully sticking, while the frequencies where $s < 1$ correspond to the friction dampers being active. This shows that the dampers are slightly active near $17\,\text{Hz}$ (s=0.96) and more active near $40\,\text{Hz}$

(s=0.91): these correspond to the in-phase [1 1] and out-of-phase [1 −1] modes of the first passband. The second passband shows a wider range of frequencies where the dampers are active (90-140 Hz): this is because the uncertainty in this frequency range is sufficiently large that the two passband modes are not always distinct. It can be seen that using a single HBM simulation gives a poor estimate of the ensemble mean (dotted line). The y-axis does not show the full range of this case, with $s = 0.27$ and $s = 0.66$ at the two off-axis minima. Using 100 HBM simulations produces the dot-dashed estimate which is noisy but in the right vicinity of the $N = 1000$ reference (solid).


<div style="text-align:center">(a)   (b)</div>
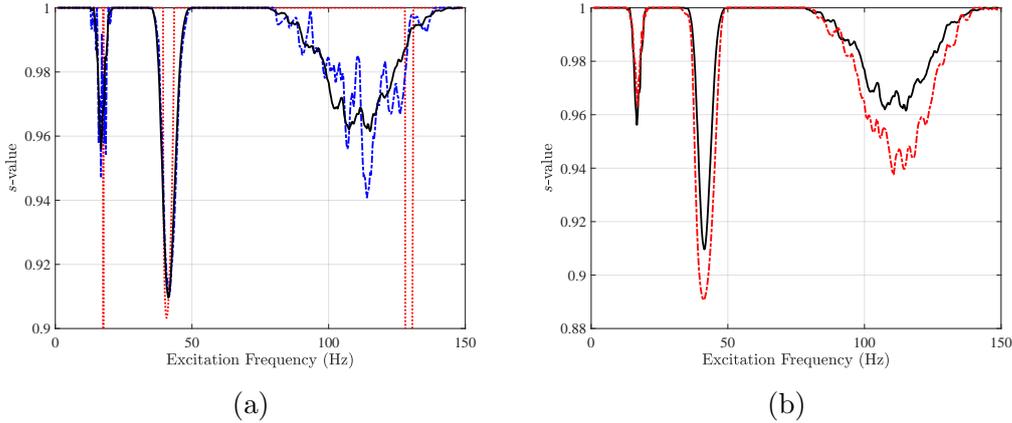
Fig. 11. Estimating s-values using (a) a subset of N HBM simulations with $N = 1$ (dotted), $N = 100$ (dot-dashed), $N = 1000$ (solid); and (b) using the self-consistency method (dot-dashed) compared with the $N = 1000$ (solid) reference.

An alternative approach is to use a self-consistency argument. Rather than solving the HBM equations for each system sample, the solution $s_{i,\mathrm{mean}} = \mathrm{E}[s_i]$ is sought that results in the expected value of the residue function being zero:

$$\mathrm{E}[R_i] = \mathrm{E}[s_i] - \mathrm{E}[s_{2,i}] \tag{19}$$

Figure 11(b) shows a comparison between the estimated mean compared with the reference mean from $N = 1000$ HBM simulations. Although the estimate is imperfect, with an apparently consistent bias, it has been found that output distribution is insensitive to the exact value for the mean, and this method has been found to be effective and efficient. The self-consistency approach can be made faster by using a subset of ensemble systems, with a similar trade-off as for HBM in (a). For subsequent simulations $N = 100$ has been found to be effective. Note that this does not carry the same computational cost as 100 HBM simulations: this ensemble only affects the computation of the residual in Eq. 19 (which is fast), not the iterative solution (which is the limiting factor in HBM).

4 Results and comparisons

The results are presented in four parts: Section 4.1 shows the two-beam comparison between MaxEnt predictions and Monte Carlo HBM simulations; Section 4.2 shows the comparison with two-beam experimental results; Section 4.3 presents the eight-beam comparison between MaxEnt and Monte Carlo HBM; and Section 4.4 shows the eight-beam comparison with experimental results.

4.1 Two-beam comparison: MaxEnt and HBM

The nominal system parameters identified from the two-beam experiment that have been used for HBM simulations are summarised in Table 5. The input excitation for the two-beam case was just applied to Beam 5 using the shaker near the base (Site 3).

Table 5
Nominal properties of experimental test rig.

Coefficient of friction, μ_0	0.5
Damper coupling stiffness, k_c (Nm^{-1})	4800
Beam natural frequencies, f_n (Hz)	[16.94, 110.5, 270.5, 542.4]
Beam damping factor, ζ_n	[0.0020, 0.0032, 0.0068, 0.020]
Normal force (N)	0.48

To illustrate the behaviour of the nominal tuned system, Fig. 12 shows a set of frequency response curves of the driven beam for different amplitudes of input force, with $F_0 \in [0.01 \quad 0.03 \quad 0.1 \quad 0.3 \quad 1 \quad 3 \quad 10]$ N. The response has been normalised by the input force to compress the dynamic range of the y-axis. It can be seen that over this range of input forces the response transitions from a linear 'fully-sticking' limit at $F_0 = 0.01$ N (thin line) to a linear 'fully-slipping' limit at $F_0 = 10$ N (bold line). At the low-amplitude limit the first two peaks correspond to the first passband (first bending mode of the beams) and the next two peaks correspond to the second passband (second bending mode of the beams). The two peaks within each passband are the $[1 \quad 1]$ and $[1 \quad -1]$ modes. At intermediate levels of input force the friction damper becomes active, and this is the regime of operation for which they are designed. As the input force amplitude increases, the damper suppresses the second peak of each passband until the beams become effectively uncoupled.

In order to test the MaxEnt approach, uncertainty is applied simultaneously to the coefficient of friction μ_0, the damper contact stiffness k_c, the beam natural frequencies f_n, and the beam damping factors ζ_n. For each variable
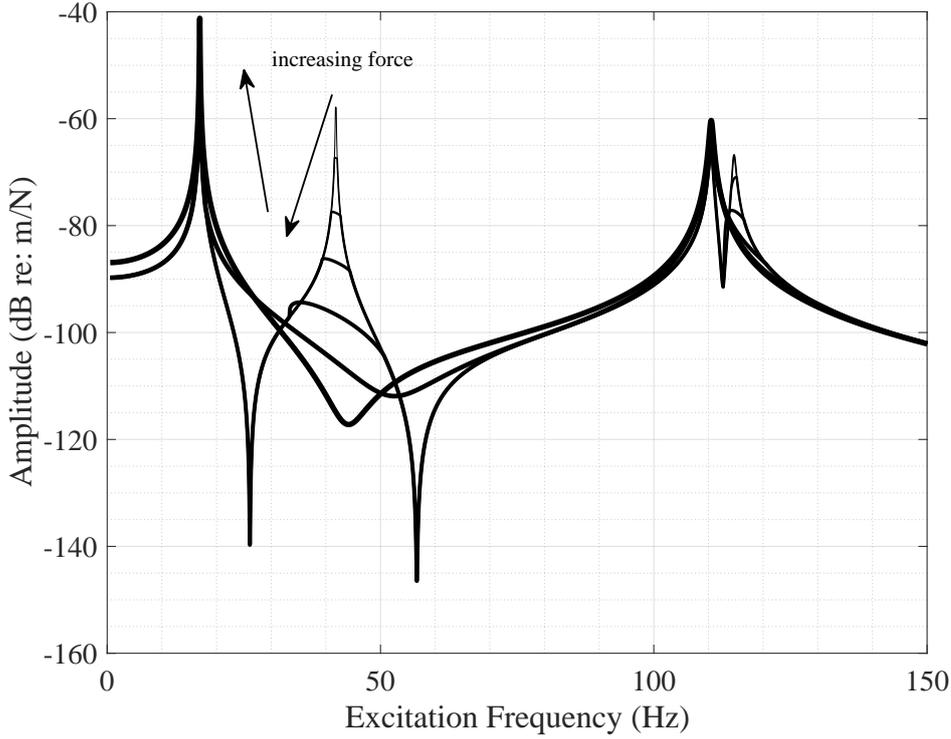
Fig. 12. Sequence of frequency responses for two-beam system for a range of input force amplitudes $F_0 \in [0.01 \quad 0.03 \quad 0.1 \quad 0.3 \quad 1 \quad 3 \quad 10]$ N: lines weighted according to force input level, with increasing line weight corresponding to increasing force amplitude.

the uncertainty is chosen to be independently normally distributed with mean given by the nominal parameters in Table 5 and standard deviation σ_p defined as a percentage of the mean. A Monte Carlo set of 1,000 HBM simulations was carried out for comparison with the MaxEnt predictions. Figure 12 shows an example set of distributions for these parameters for the case $\sigma_p = 10\%$.

Figure 14 shows the response percentiles using an intermediate input force $F_0 = 0.3$ N (in the middle of the range of Fig. 12) for two levels of uncertainty: (a) $\sigma_p = 1\%$; and (b) $\sigma_p = 10\%$. The lines correspond to the 5, 50 and 95 percentiles across all beams, inferred by their sequence from smallest to largest amplitude.

The Monte Carlo HBM results are based on 1,000 simulations which took approximately 2,200 seconds to compute: this was the observed time taken for the calculation (not the CPU time), utilising parallel processing on an eight-core 3.2GHz computer. This compares with 20 seconds for the MaxEnt simulations using the self-consistency method for estimating the mean, using the same computer, i.e. two orders of magnitude faster. It is recognised that comparisons of computational speed are affected by many complicating fac-

27
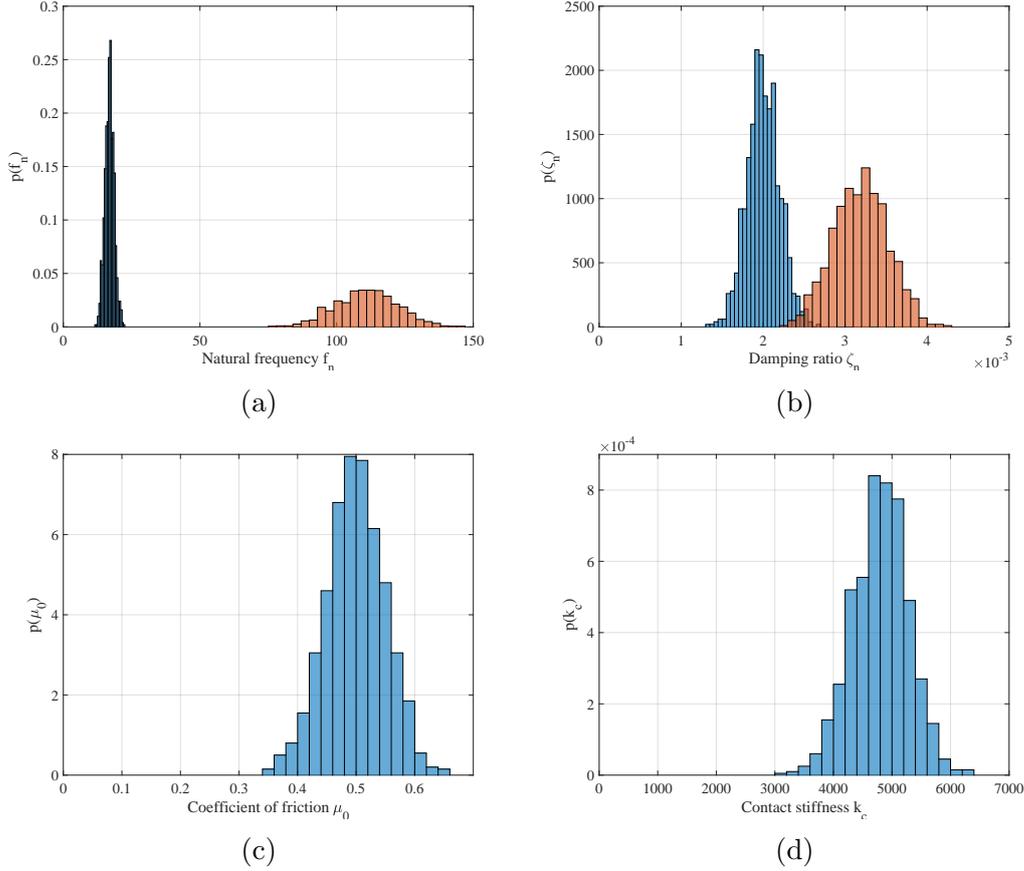
Fig. 13. Summary of input probability density functions used for the uncertain parameters of each beam or friction damper for the case $\sigma_p = 10\%$: (a) natural frequencies; (b) damping ratios; (c) coefficient of friction; and (d) coupling stiffness.

tors, nevertheless this represents a significant improvement. The increase in speed is because the self-consistency method for estimating the mean carries only a slightly greater computational cost compared to a single HBM simulation, and the rest of the uncertainty propagation involves the prediction of a linear response.

The choice to use an ensemble size of $N = 1,000$ was empirically found to give converged estimates of the 5th and 95th percentiles: one indicator of this convergence is the relatively low noise levels within the predicted frequency response percentiles.

It is clear from Fig. 14 that the MaxEnt percentile estimates are generally very good in both cases. In (a) the small uncertainty gives very tightly defined frequency response distributions, with negligible response uncertainty for 0–40 Hz and very small uncertainty at higher frequencies. For most frequencies there is no visible difference between MaxEnt and HBM percentiles. Some differences are apparent on close inspection just below 50 Hz and 120 Hz. In
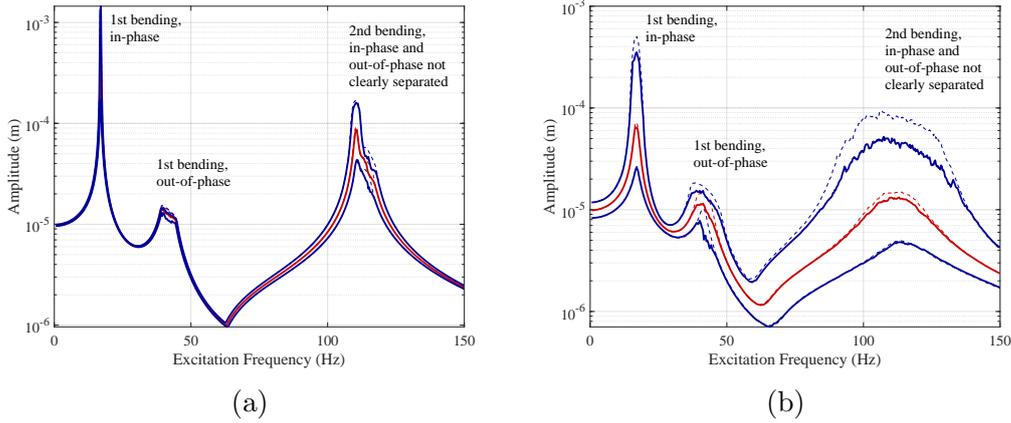
Fig. 14. Comparison of MaxEnt estimates (solid) with HBM Monte Carlo results (dashed), showing response amplitude percentiles (5, 50 and 95% inferred by sequence of increasing amplitude) using a relative standard deviation of: (a) $\sigma_p = 1\%$ and (b) $\sigma_p = 10\%$. These results are for a two-beam system with an intermediate level of input excitation force $F_0 = 0.3\,\mathrm{N}$ applied to one beam, using 1,000 samples for MaxEnt and HBM.

(b) the response uncertainty is much larger and is accurately estimated by the MaxEnt estimate across most frequencies. Some moderate differences can be observed at the 17 and 40 Hz peaks and in the range 90–130 Hz. These differences have been carefully investigated and some of the reasons are presented in the following discussion.

A zoom of the friction-damped 40 Hz peaks are shown in Fig. 15: it can be seen that the percentiles agree well then suddenly deviate for a range of frequencies before agreeing again. This is because the 'mean' solution is generally in the 'sticking' regime with $s = 1$, resulting in a MaxEnt PDF being a delta function at $s = 1$. In other words when the mean model is sticking the MaxEnt method does not introduce any additional uncertainty, hence there is extremely close agreement between MaxEnt and HBM. When the mean model gives $s < 1$, then MaxEnt introduces additional uncertainty in the distribution of s. It might be expected that this distribution would be conservative with a response distribution wider than the true distribution, but this effect is balanced by the force constraint which reduces the admissible domain of s-values. This is made clear by applying MaxEnt without the force constraint, which results in Fig. 16: the predicted distribution becomes much wider than the HBM distribution when the friction dampers are active. Overall the force constraint has been found to be beneficial: without it the predictions are can be overly conservative.

The results shown in Fig. 14 are representative of several persistent characteristics of these MaxEnt predictions: low amplitude excitation gives accurate predictions of uncertainty because the friction dampers are in the sticking regime; small uncertainty results are accurately predicted because the com-
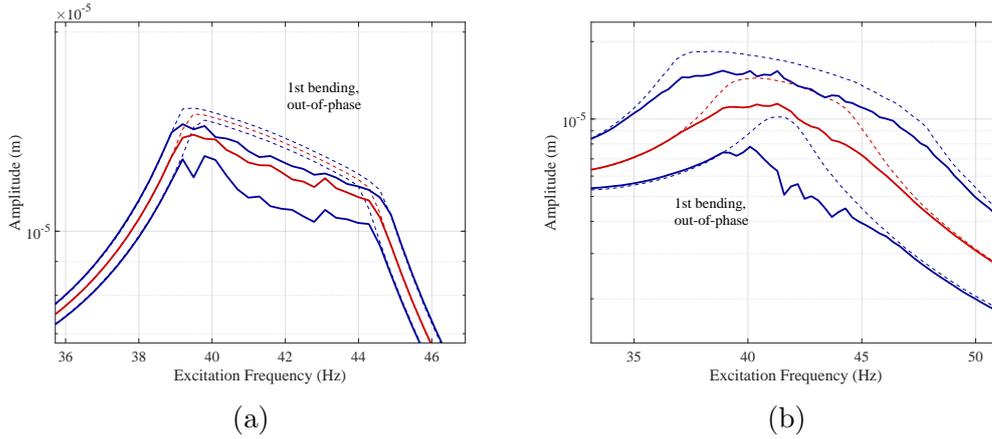
Fig. 15. Expanded view of friction-damped peak near 40 Hz from Fig. 14.
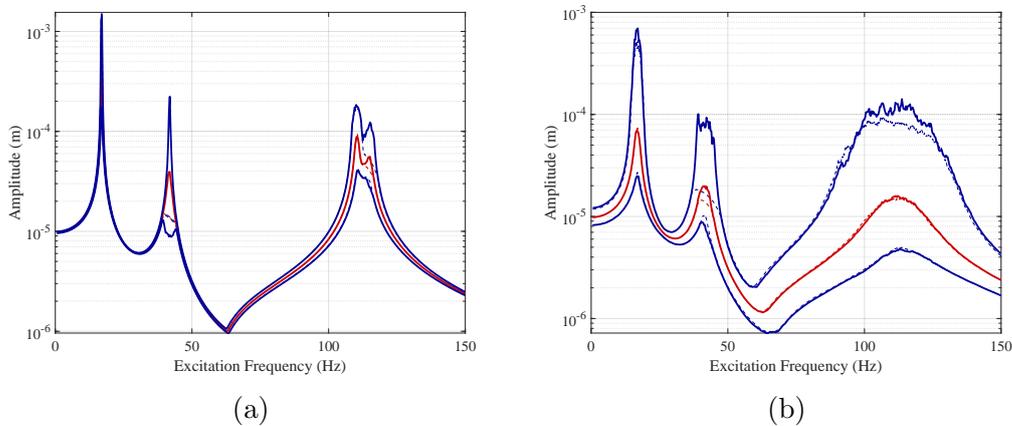
Fig. 16. Effect of applying MaxEnt without a force constraint: test case is otherwise the same as Fig. 14.

bination of the force constraint and MaxEnt weighting to the mean solution gives a tightly defined PDF for s; high amplitude excitation combined with significant uncertainty generally results in greater discrepancies particularly for the tails of the response distribution.

These cases with more significant discrepancies may not be particularly problematic: gas turbine blades are manufactured to extremely tight tolerances, so uncertainty associated with structural properties is generally considered to be small; but other properties cannot be so tightly controlled such as the inter-blade contact stiffness which is a function of the wear-profile and start-up conditions for the individual dampers. It is also important to be clear as to what precision is expected of the tails of a predicted distribution: percentiles from Monte Carlo reference simulations are sensitive to the precise details of uncertainty characterisation, and as will be seen there is a practical limit as to the precision of this input uncertainty characterisation. This means that Monte Carlo HBM will also be an approximation to the actual response

30

statistics of the blades because the input uncertainty distributions may not be known accurately, so care is needed before discounting MaxEnt as a fast estimate of response distributions when there are relatively small differences compared with HBM.

Nevertheless, the reasons for the high-amplitude and high-uncertainty discrepancies have been investigated in detail. In summary there are two primary reasons for discrepancies that have opposite effects on the response distribution predictions: (1) computational uncertainty; and (2) nonlinearity-induced localisation. As already discussed, the additional 'computational uncertainty' introduced by the method tends to give a conservative prediction of distribution tails and a reasonably good estimate of the mean. The effect of localisation is that for some members of an ensemble of systems, high response amplitudes are associated with a small number of beams with corresponding s values close to 0, despite a mean s much closer to 1. This means that very few samples are generated within the MaxEnt algorithm that have the specific conditions needed for localisation to be represented in the estimated distribution, resulting in an underprediction of the response. In other words smooth exponential distributions of s are not a good representation of the response when localisation occurs, and the discrepancy is specific to when the friction dampers are active as the sticking limit is well represented by MaxEnt even if localisation is present.

Figure 17 shows the full Monte Carlo HBM solutions associated with the two-beam example of Fig. 14(b), taking an excitation-frequency snapshot at (a) $f = 40\,\text{Hz}$ and (b) $f = 120\,\text{Hz}$. The plots show the response amplitudes of each beam against their corresponding s-values (calculated by HBM), with the driven beam shown as crosses and the undriven beam as circles. It can be seen in (a) at $40\,\text{Hz}$ that the solutions are reasonably tightly distributed around $s = 0.9$. In (b) at $120\,\text{Hz}$ the solutions are more broadly scattered in the range $0.5 < s \leq 1$ and for smaller values of s the directly driven beam response is much larger than the second beam. This is an example of localisation and the sampling strategy adopted here for MaxEnt does not reliably capture the tails of this s-distribution that correspond to high amplitude localised response. Some effort has been made to explore different sampling strategies in order to more reliably capture the tails of the s-distribution and it is expected that further improvements are possible, but this remains a challenge for future work.

4.2   Two-beam comparison: MaxEnt and Experiment

Initial experimental verification of the MaxEnt method is performed on a two-beam subset of the full experimental rig. This is achieved by disconnecting all
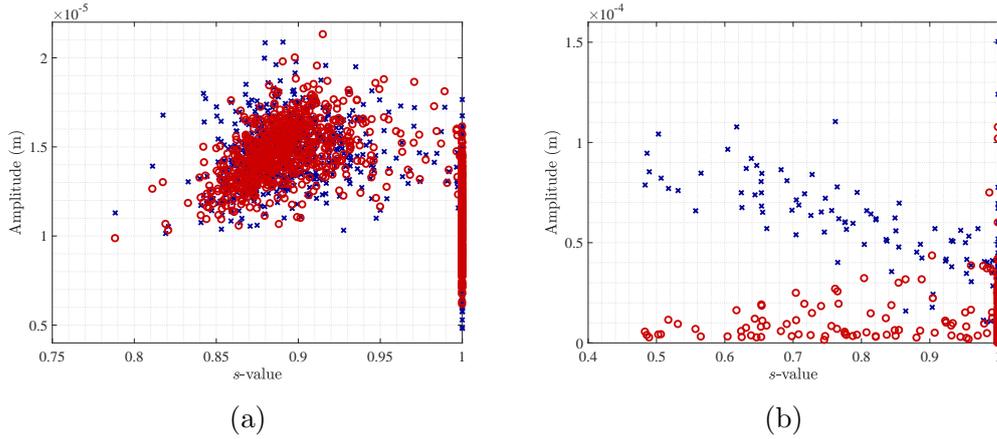
Fig. 17. Frequency snapshots of Monte Carlo HBM solutions for (a) $f = 40\,\text{Hz}$; (b) $f = 120\,\text{Hz}$. Test case corresponds to Fig. 14; crosses show the directly driven beam solutions and circles show the second-beam solutions.

the friction dampers except one and damping all other beams in the rig using foam inserts. The beams chosen for the two-beam experiment are Beams 5 and 6. The two-beam system illustrates many of the important physical characteristics and provides a highly controlled case to highlight the importance of having a solid model of the uncertainty in the system.

Experimental data was collected for excitation into a single beam (Beam 5) at three amplitudes, using a single normal force. An ensemble of 50 random systems was generated by moving the masses on both Beams 5 and 6, and all the data presented in this section will be shown as percentiles of the response across the ensemble. Percentiles will be presented at the 10, 50 and 90th levels: slightly tighter than the numerical results due to the smaller size of the experimental ensemble compared with HBM simulations. As identified in figure captions, percentiles may be for the full set of data across both beams, or only the maximum of the response at a particular beam location.

The full set of parameters used in the two-beam experiment not already introduced elsewhere is given in Table 6.

Table 6
Experimental parameters for two-beam tests

Parameter	Value
Excitation Force (N)	[0.18, 0.36, 0.72]
Normal Force (N)	0.48
Coupling Stiffness Nm^{-1}	4800

32

4.2.1 Fitting uncertainty

The initial uncertainty model used in the MaxEnt comparisons is based on the results from Section 2.3, and Table 3. The uncertain parameters used as inputs to the MaxEnt method were taken to be the beam natural frequencies and mode shapes. Uncertainty in the coupling stiffness parameters and coefficient of friction were small and did not significantly change the prediction percentiles. Two response metrics are presented in the figures that follow: (i) percentiles based on the maximum response across the beams at Site 2 (mid-beam) for each case and (ii) percentiles based on all beam responses at Site 2.

As a starting point, a Gaussian uncertainty model approximately based on Table 3 was used as the input for MaxEnt: the standard deviation for the natural frequencies was taken to be $\sigma_p = 5\%$, and $\sigma_p = 5\%$ for the first mode amplitudes and $\sigma_p = 30\%$ for the remaining modes. Figure 18(a,b) shows a comparison of the MaxEnt predicted percentiles (solid) with the experimental results (dashed), with (a) showing the response distribution across all beams; and (b) showing the response distribution of the maximum beam response. It can be seen that the overall level of agreement is good for both metrics: the main features and corresponding uncertainty are broadly captured by the MaxEnt predictions. In particular the two distinct modes at 17 and 40 Hz are predicted, including the higher level of friction-damping for the out-of-phase mode at 40 Hz. The broader 'peak' corresponding to the second passband is also present and the MaxEnt approach captures how the two modes can no longer be distinguished, which is also true for the experimental results.

However some details differ, for example the 10th percentile in the 60–100 Hz region (a) and (b), and the high frequency 'tail' of the second passband above 120 Hz is not accurately captured. Another apparent discrepancy is the difference between the 50th percentiles in (a) at 25 Hz where the MaxEnt method completely misses the antiresonance. This is in fact a post-processing artefact particular to the two-beam system: the metric in (a) includes the response of both beams but only one of the beams has an antiresonance, so the 50th percentile algorithm can arbitrarily switch between the selected beams. Also of note in (b) is a peak in the experimental 50th percentile just below 120 Hz that is not present in the MaxEnt prediction. This feature is not present in (a) when considering the metric of all the beams, which highlights how the metrics involving the maximum of the response tend to highlight extreme features that are less well captured by the MaxEnt method.

While these comparisons are generally good, careful consideration of the uncertainty model allows for improvements. The initial guess for an uncertainty model assumed a Gaussian distribution of the parameters. However the mistuning characterisation of the linear rig shows that a uniform distribution for

the natural frequencies is more representative as the tuning locations of the masses were randomised uniformly. More significantly it is apparent from the linear characterisation that the natural frequencies and mode shapes are correlated. As the mass is moved from tip to base of the beam the first natural frequency increases while the second natural frequency decreases. Further, the magnitude of the second modal amplitude at the tip of the beam increases as the natural frequency decreases, while the modal amplitudes of the second mode shape at the excitation location and mid beam decrease before increasing again. Investigation has shown the most important of these effects is the correlation of the second mode shape and second natural frequency, so the effect of the MaxEnt fits as a result of this additional correlation is shown in Figure 18(c,d): (c) shows the response distribution across all beams and (d) shows the maximum response distribution.


(a)

(b)

(c)

(d)
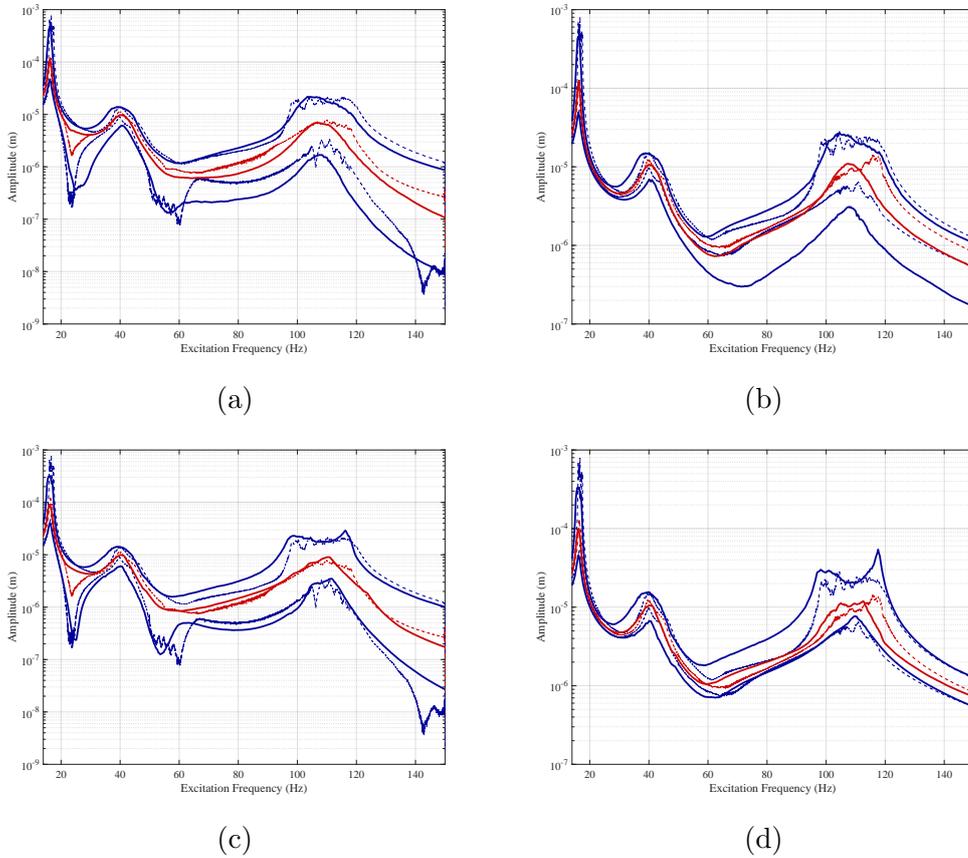
Fig. 18. Comparison of MaxEnt predicted percentiles (solid) with experimental results (dashed) for the two-beam system with an intermediate input force amplitude of $F_0 = 0.36$ N: (a,b) show MaxEnt predictions using a Gaussian input uncertainty; (c,d) uses a uniform distribution for the natural frequencies correlated with modal amplitude uncertainty.

There are two significant changes caused by the change to a uniform distribution that can be seen in Fig. 18(c,d). Firstly the edges of the second passband particularly at the 90th percentile level are noticeably sharper and this is

also apparent in the 50th percentile, as the uniform distribution tightens up the range of natural frequencies across the numerical ensemble. Perhaps most striking is that the 10th percentile line is much less conservative than in (a,b). The 90th percentile is still somewhat conservative for the maximum of the response, but when considering the percentiles of all the beams (c) this is far less apparent, and the fit with experiment is generally very good. It can also be seen that the 50th percentile prediction is improved for both metrics (c) and (d).

4.2.2 Other two-beam comparisons

For completeness Fig. 19 shows a comparison of MaxEnt predictions with the experimental results for input excitation amplitudes of $F_0 = 0.18\,\mathrm{N}$ and $F_0 = 0.72\,\mathrm{N}$ (the case $F_0 = 0.36\,\mathrm{N}$ was presented above). The uncertainty model is chosen as the final model involving correlated mode shapes corresponding to Fig. 18(c,d). It is encouraging that the quality of the fits are broadly similar to the previous case.

In summary these results suggest that a very careful model of the uncertainty in the system is required to recover close agreement between the experimental and MaxEnt results. This highlights how uncertainty propagation tools can only be as good as the input uncertainty characterisation. In practice, some sources of uncertainty are epistemic in nature and are not necessarily associated with an ensemble of instances. Therefore care is needed when comparing results, and when interpreting predictions.

4.3 Eight-beam comparisons: MaxEnt and HBM

The nominal system parameters identified from the eight-beam experiment were used for HBM simulations and are summarised in Table 7. The difference from the two-beam case is the damper contact stiffnesses, which have been individually estimated. The MaxEnt method does not rely on knowing values for these individual coupling stiffnesses, but this detail was included as the data was readily available. It is clear that there is considerable variability across the dampers such that even if the beams were perfectly tuned there is still mistuning due to the coupling variability. This is likely to be a factor in gas turbines, where the blades themselves are manufactured to a very tight tolerance so the variability in contact configuration of the dampers may be the dominant cause of mistuning.

The input force excitation pattern across the beams for the following tests is taken to be the ideal mode shape associated with the first mode of a passband that involves relative motion between the beams (i.e. Mode 2 of the passband
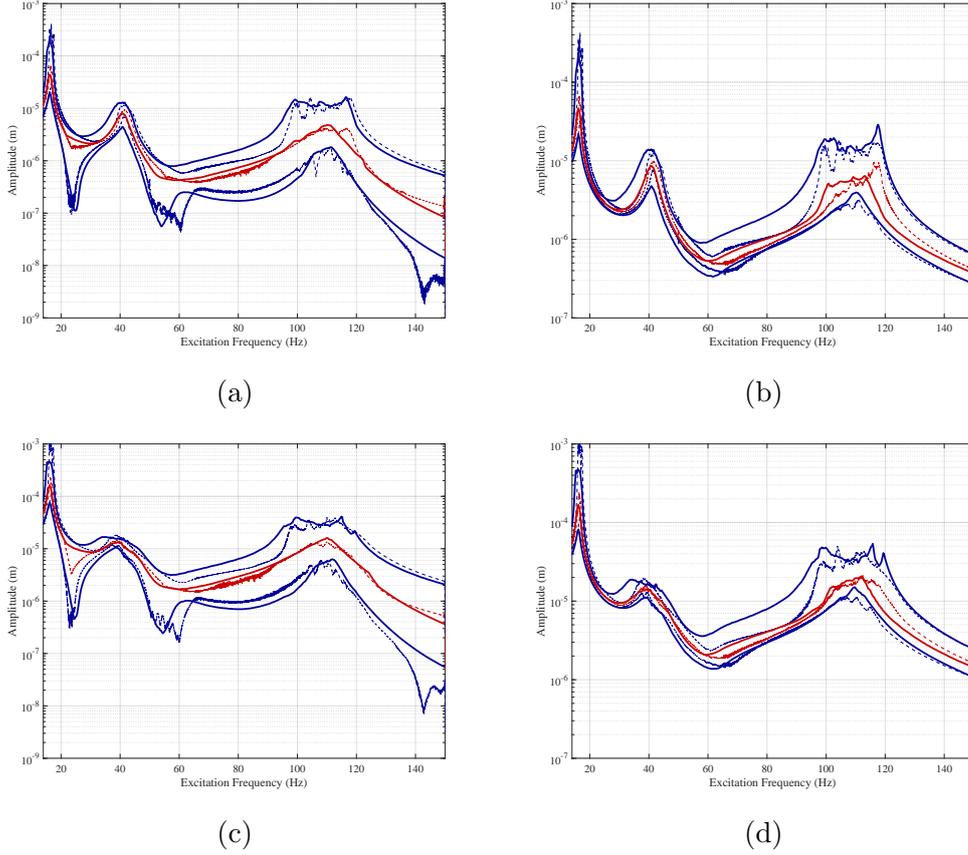
Fig. 19. Comparison of MaxEnt predicted percentiles (solid) with experimental results (dashed) for the two-beam system with input force amplitude of (a,b) $F_0 = 0.18\,\text{N}$ and (c,d) $F_0 = 0.72\,\text{N}$. The MaxEnt input uncertainty model uses a uniform distribution for the natural frequencies correlated with modal amplitude uncertainty.

Table 7
Nominal properties of experimental test rig.

Coefficient of friction, μ_0	0.5
Damper Contact stiffness, k_c (N/m)	[7200, 4500, 6500, 7200, 5000, 8000, 7000]
Beam natural frequencies, f_n (Hz)	[16.94, 110.5, 270.5, 542.4]
Beam damping factor, ζ_n	[0.0020, 0.0032, 0.0068, 0.020]
Normal force (N)	0.48

for a tuned system which has a single nodal line between Beams 4 and 5). In the linear sticking-limit for a perfectly tuned system this has the effect of only exciting the corresponding mode from each passband, but in the presence of uncertainty and mistuning then this forcing pattern excites all modes.

Figure 20 shows two sets of frequency response curves of the maximum re-

sponse across all beams for different amplitudes of input force, with $F_0 \in [0.01 \; 0.03 \; 0.1 \; 0.3 \; 1 \; 3 \; 10]$ N. The response has been normalised by the input force to compress the dynamic range of the y-axis. The first case shown in (a) is for the idealised perfectly tuned system with all contact stiffnesses equal $k_c = 4800 \, \mathrm{Nm}^{-1}$; case (b) shows the results using the parameters in Table 7. It can be seen in both cases that the friction damper has most visible effect on the first passband response (15–60 Hz). At low input amplitudes, (a) shows that the Mode 2 excitation pattern only excites the corresponding mode as expected. As the input force increases and the friction dampers become active, then damping increases and eventually the frequency response transitions to an effectively uncoupled response with natural frequency at $f_n = 16.94$ Hz. For the mistuned system in (b) the response is more complicated: at low input force eight modes are visible in the frequency response function with the dominant mode corresponding to the excitation pattern, as the dampers become active at higher amplitudes some of these are affected more strongly than others, until at high amplitudes the response tends to nearly a single peak at 16.94 Hz. For the subsequent comparisons with MaxEnt the mistuned case (b) will be presented.


(a) (b)
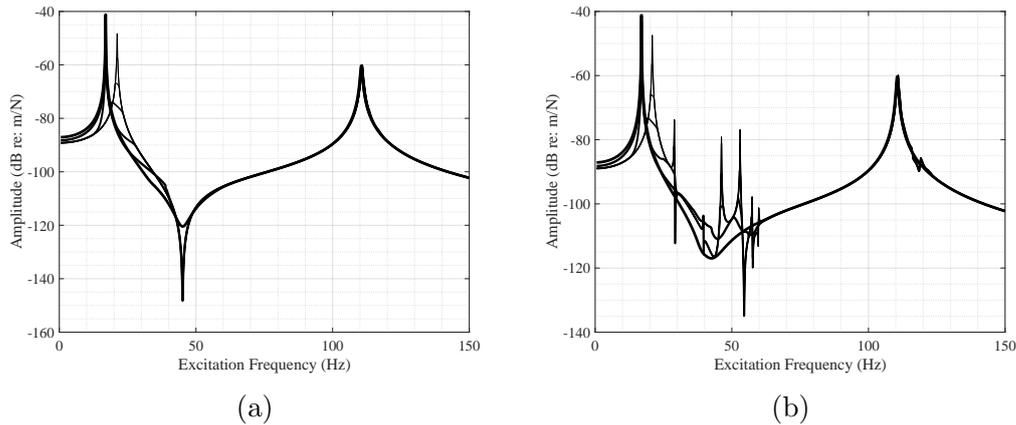
Fig. 20. Sequence of simulated frequency responses using HBM for eight-blade system using a range of input force amplitudes $F_0 \in [0.01 \; 0.03 \; 0.1 \; 0.3 \; 1 \; 3 \; 10]$ N: (a) tuned system; (b) mistuned system. Lines weighted according to force input level, with increasing line weight corresponding to increasing force amplitude.

Figure 21 shows a comparison between MaxEnt and HBM of the response distribution percentiles (5th, 50th and 95th) for a sequence of cases: the plots on the left (a,c,e) show $\sigma_p = 1\%$ and on the right (b,d,f) show $\sigma_p = 10\%$; with pairs (a,b), (c,d) and (e,f) corresponding to increasing input force $F_0 = 0.1$, 0.3, and 1 N. These excitation amplitudes span a range of representative cases for which the friction dampers are active. It can be seen that in general agreement is good: at low amplitudes (a,b) when the beams are in the 'sticking' regime then agreement is exact as before; at intermediate amplitudes (c,d) then the MaxEnt method gives slightly conservative predictions; while for high amplitudes then the localisation effect dominates and (e,f) show small

under-prediction. For all excitation levels there is good agreement when uncertainty is small: this is closer to realistic levels of uncertainty in precision gas turbine blades, suggesting that this method could be extremely valuable for fast estimates of response distributions during the design phase of a turbine. The relative contribution of each uncertain parameter to the response variability has not been calculated here: this falls beyond the scope of the present study and is left for future work.

The metric chosen for comparisons so far has been the maximum response amplitude across all beams. Predicting extreme-response statistics is very challenging, so it is also interesting to compare the HBM and MaxEnt distributions across all beams. Figure 22 shows a set of comparisons using this metric for a range of excitation amplitudes and uncertainty levels: left shows $\sigma_p = 1\%$; right shows $\sigma = 10\%$, and force amplitudes $F_0 = 0.1$, 0.3, and 1 N correspond to pairs (a,b), (c,d), and (e,f) respectively. It can be seen that the MaxEnt predictions for the overall response statistics are extremely good, providing accurate distribution percentiles for all cases and where small discrepancies occur the MaxEnt percentiles are slightly conservative.

Comparing Fig. 22 with Fig. 14 it may be thought that the results for the eight-beam system are better than for the two-beam. Some care is needed here: this comparison is not like-for-like as Fig. 22 is based on the metric across all beams, not the extreme response as in Fig. 14. The more reliable comparison is between Fig. 21 with Fig. 14, and the results have a similar level of discrepancy. This highlights that MaxEnt predictions are better for estimating distribution percentiles for the overall response than for the extreme response.
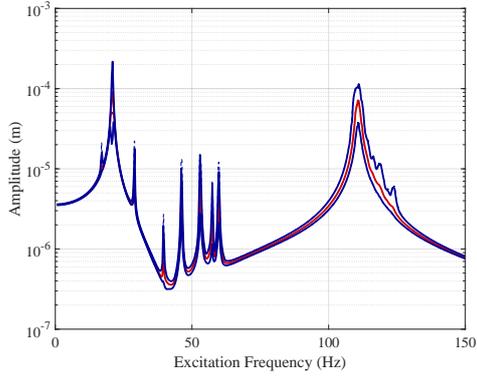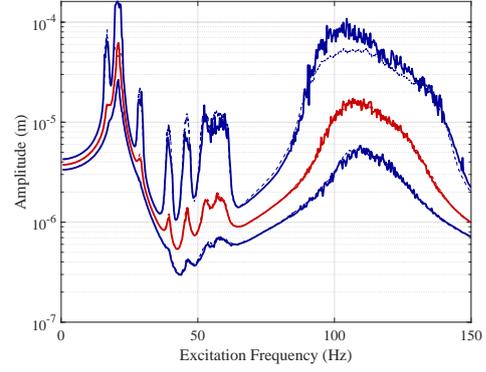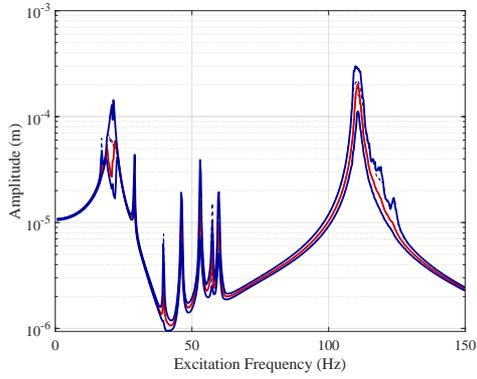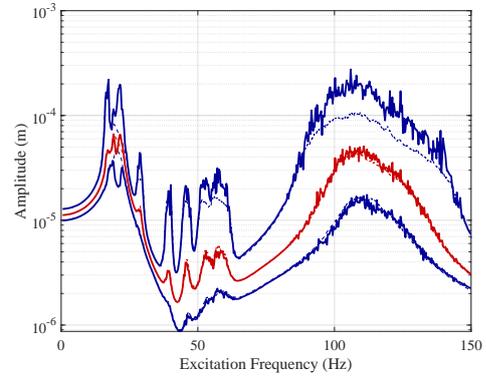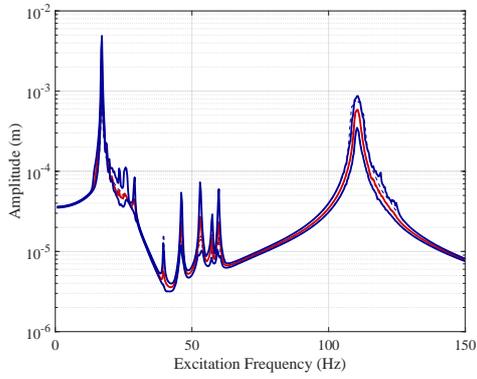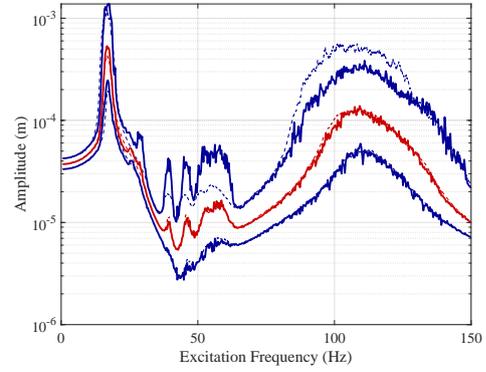
Fig. 21. Comparison between MaxEnt and HBM Monte Carlo for a range of input excitation amplitudes and uncertainty levels for the eight-beam system: (a,c,e) $\sigma_p = 1\%$; (b,d,f) $\sigma = 10\%$; with $F_0 = 0.1$, 0.3, and 1 N associated with pairs (a,b), (c,d), and (e,f) respectively. The percentiles represent the response distribution for the maximum beam response, using 1,000 samples for MaxEnt and HBM.
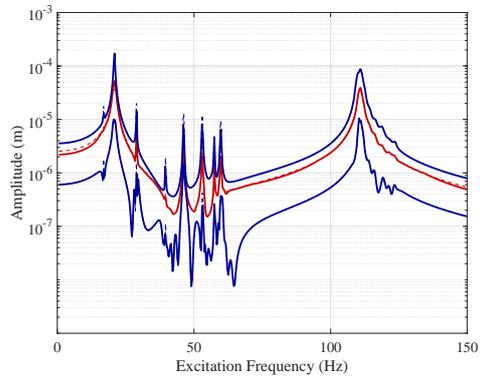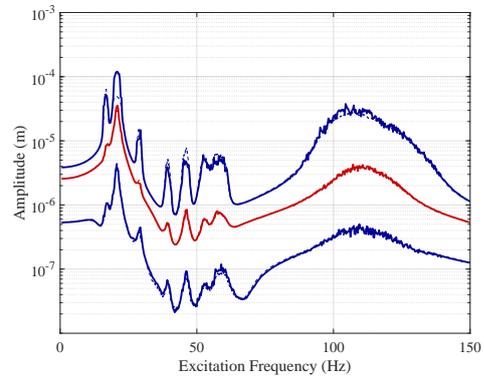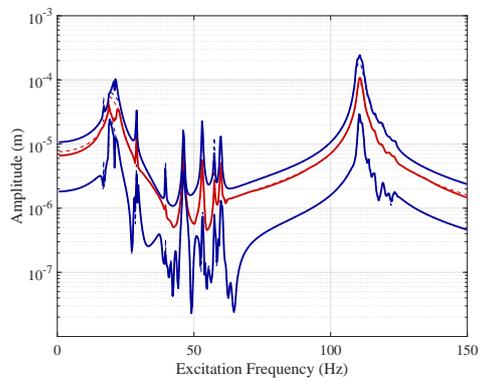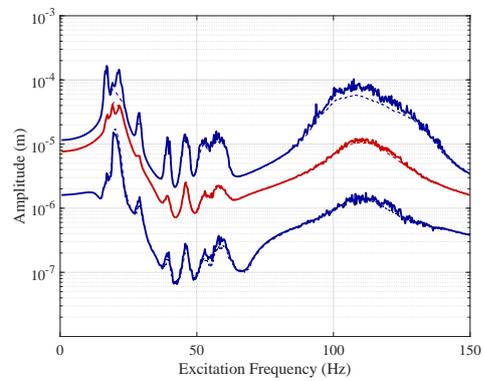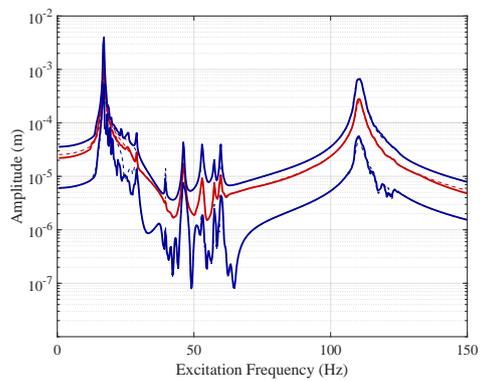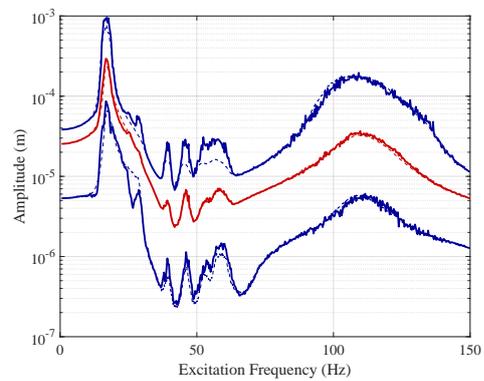
Fig. 22. Comparison between MaxEnt and HBM Monte Carlo for the response distribution across all beams. The sequence is otherwise equivalent to Fig. 21.

40

The final characterisation of the coupled rig required for the eight-beam tests are the engine order mode shapes across the rig. The excitation pattern for the eight-beam tests was chosen to have similar dynamic properties to engine order excitation in the true turbine, i.e. to attempt to select individual modes from each passband. A key difference for the non-circular beam array is that repeated frequency doublet modes are not present, so travelling wave excitation (with shifted phase polarisations of engine order excitation) is not necessary. To achieve the mode-selection property the excitation pattern was chosen to be as similar as possible to one of the global mode shapes of the assembly.

Measuring the engine order mode shapes is challenging as the coupled linear rig in a fully sticking configuration is difficult to achieve. However, estimates of these mode shapes can be made using a small amplitude excitation ($F_0 = 0.018\,\text{N}$) and a relatively high normal force ($N_0 = 0.96\,\text{N}$) applied to the friction dampers. Excitation patterns were then chosen to best fit these measured mode shapes. The excitation pattern for Engine Order 2 is given in Table 8 (normalised to the peak amplitude), which has two nodal diameters consistent with the expected mode shape. This force pattern was then used as the deterministic input amplitude to each beam of the assembly.

Table 8
Normalised excitation force pattern (N)

Beam	1	2	3	4	5	6	7	8
EO 2	-0.29	-0.59	0.59	1.0	1.0	0.59	-0.59	-0.29

The exploration of the two-beam rig provided an understanding of the relative importance of different rig parameters for fitting an accurate response. Although precise characterisation was needed for accurate agreement, it is reassuring that if the uncertainty model is too conservative and does not include important correlations between the modal parameters then it tends to provide a conservative estimate of the response.

To complete the experimental work, MaxEnt predictions were compared with experimental results for the eight-beam rig. An ensemble of 50 random systems was generated by moving the masses on all beams, and the data presented in this section will be shown as percentiles of the response across the ensemble. Percentiles will be presented at the 10, 50 and 90th levels. The uncertainty model was based on the two-beam case using the correlations between parameters required to give the fits given in Fig. 18(c,d). The mean modal parameters for all beams were chosen to be the average parameters across the two beams given in Table 6. Engine order excitation was implemented using EO2 given

in Table 8. Excitation amplitudes were at the nominal levels of 0.18, 0.36 and 0.72 N corresponding to the maximum force input.

The results of this test for the three nominal excitation amplitudes are shown in Fig. 23: (a,b) shows the case for $F_0 = 0.18\,\text{N}$; (c,d) $F_0 = 0.36\,\text{N}$; and (e,f) $F_0 = 0.72\,\text{N}$. The response distribution across all beams is shown in the left plots (a,c,e) and the maximum response distributions are shown in the right plots (b,d,f). As for the two-beam case the fits for the percentiles based on all beam responses (a,c,e) are generally better than those when considering the maximum response of the beams as a metric (b,d,f). Overall, the agreement is remarkably good: the main features in each figure are captured by the MaxEnt predictions. Although the fit is not as good as for the two-beam case, the greater variability in the rig between beams and coupling links means this is not surprising.

One of the interesting characteristics of the eight-beam tests is the different ranges of modal and statistical overlap that are apparent. Two distinct ranges can be identified: high modal overlap across an entire passband (e.g. 90–130 Hz); and low to intermediate modal overlap where some modal structure can be identified (e.g. 14–60 Hz).

In regions of high modal overlap the MaxEnt prediction performs best and with a slight bias towards overestimating the response, particularly at the higher frequency edge of the pass-band. This illustrates that regions of the response with high modal overlap seem to be less sensitive to accurate characterisation of the details of the uncertainty model, as long as the broad level of uncertainty is correct.

For regions of intermediate modal overlap, the quality of the fit remains good particularly at the 50th percentile level. However, some of the peaks in the 90th percentile are overly conservative in the MaxEnt prediction. This may indicate that some of the details of the uncertainty model or the operating point do not exactly match the experimental results, and as a result extreme predictions (and particularly the maximum response), do not correlate as well between MaxEnt and experimental results. The good agreement between the MaxEnt predictions and benchmark HBM results (Figs. 21 and 22) suggest that the differences are likely to be predominantly due to inaccurate characterisation of the input uncertainty rather than being an approximation inherent to the MaxEnt method.

In summary it is clear that it is extremely challenging to obtain good agreement between the experimental and predicted ensembles of results. The good agreement between MaxEnt and the Monte Carlo HBM simulations show that it is not a problem with MaxEnt, but rather with the long-standing challenge of reliably characterising uncertainty in nonlinear systems. The results pre-

sented here are just a representative sample of cases and many more were investigated. This highlights one of the key successes of the new MaxEnt approach, providing the ability to carry out fast predictions using different modelling details.
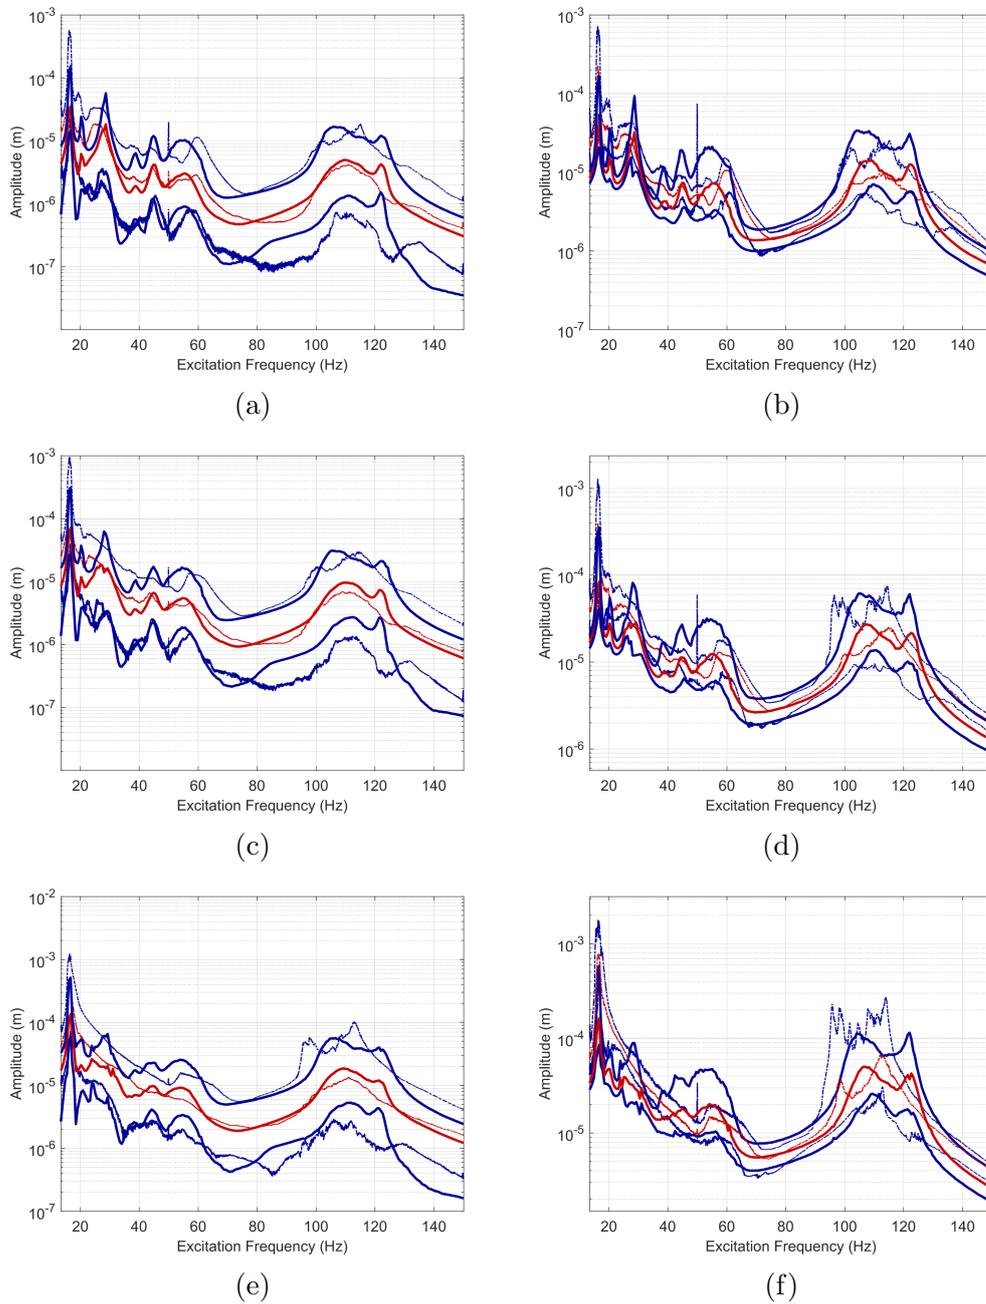

Fig. 23. Max Ent fit to experimental data, left column contains percentiles for all beam data, and right column shows only the maximum response across the beams. Row 1 is for an amplitude of 0.16 N, row 2 is for an amplitude of 0.36 N and row 3 is 0.72 N.

5 Conclusions

There is a need for efficient methods for predicting the response distribution of friction-damped gas turbine blades, that include the effects of uncertainty without requiring computationally expensive Monte Carlo simulations of the full nonlinear system. This paper has presented a novel method that provides a fast estimate of the response distribution using the principle of Maximum Entropy. The new method is approximately two orders of magnitude faster than Monte Carlo Harmonic Balance Method (HBM) simulations, and gives a reliable approximation of the response statistics across a wide range of excitation amplitudes and for different levels of uncertainty.

The method is based on using a describing function approximation of the nonlinear friction behaviour, but assigns a probability density function to the solution rather than computing individual solutions for all members of a randomised system. An initial ignorance prior distribution is updated by the principal of Maximum Entropy subject to an estimate of the mean solution and the admissible domain of the solution set.

The method has been validated by comparison with benchmark Monte Carlo simulations using the Harmonic Balance Method, and also by comparison with experimental results from a carefully designed test rig. The numerical results reveal generally good agreement between MaxEnt and Monte Carlo HBM, with some deviations in the maximum response percentiles when there is both high amplitude input force combined with high levels of uncertainty: the moderate discrepancies are caused by localisation which is not accurately represented by the MaxEnt distribution. This is not a failing of MaxEnt, and instead reveals that more information would be needed to be able to include this effect. When MaxEnt is used to predict the overall response statistics across all beams rather than just the peak response statistics, then the predictions agree extremely well with Monte Carlo HBM.

The comparisons with experimental results were particularly revealing. In order to obtain good agreement a great deal of care is needed to accurately characterise the uncertainty itself, including relevant correlations between uncertain parameter distributions. This challenge is common to all parametric uncertainty propagation methods: the accuracy of the predicted response distribution depends on the quality of the input uncertainty data. The MaxEnt approach here enabled iterating the input uncertainty characterisation over several attempts in order to obtain good agreement between MaxEnt and experimental distributions. This demonstrates the effectiveness of the method as this iterative process would not have been feasible using traditional approaches.

Acknowledgements

References

[1] J. H. Griffin. A Review of Friction Damping of Turbine Blade Vibration. *International Journal of Turbo and Jet Engines*, 7(3-4):297–308, 1990.

[2] K. Y. Sanliturk, D. J. Ewins, and A. B. Stanbridge. Underplatform Dampers for Turbine Blades: Theoretical Modeling, Analysis, and Comparison With Experimental Data. *Journal of Engineering for Gas Turbines and Power*, 123(4):919, 2001.

[3] K. Y. Sanliturk and D. J. Ewins. Modelling Two-Dimensional Friction Contact and Its Application Using Harmonic Balance Method. *Journal of Sound and Vibration*, 193(2):511–523, jun 1996.

[4] M. Krack, S. Tatzko, L. Panning-von Scheidt, and J. Wallaschek. Reliability optimization of friction-damped systems using nonlinear modes. *Journal of Sound and Vibration*, mar 2014.

[5] B. Peherstorfer, T. Cui, Y. Marzouk, and K. Willcox. Multifidelity importance sampling. *Computer Methods in Applied Mechanics and Engineering*, 300:490–509, 2016.

[6] K. Y. Sanliturk, M. Imregun, and D. J. Ewins. Harmonic Balance Vibration Analysis of Turbine Blades With Friction Dampers. *Journal of Vibration and Acoustics*, 119(1):96, 1997.

[7] L. Pesaresi, L. Salles, A. Jones, J. S. Green, and C. W. Schwingshackl. Modelling the nonlinear behaviour of an underplatform damper test rig for turbine applications. *Mechanical Systems and Signal Processing*, 85:662–679, 2017.

[8] C. Gastaldi and T. M. Berruti. A method to solve the efficiency-accuracy trade-off of multi-harmonic balance calculation of structures with friction contacts. *International Journal of Non-Linear Mechanics*, 92:25–40, jun 2017.

[9] D. Moens and M. Hanss. Non-probabilistic finite element analysis for parametric uncertainty treatment in applied mechanics: Recent advances. *Finite Elements in Analysis and Design*, 47(1):4–16, 2011.

[10] T. Butlin. Anti-optimisation for modelling the vibration of locally nonlinear structures: An exploratory study. *Journal of Sound and Vibration*, 332:7099–7122, jul 2013.

[11] T. Butlin. Response bounds for complex systems with a localised and uncertain nonlinearity. *Journal of Sound and Vibration*, 384:227–252, 2016.

[12] T. Butlin, P. Ghaderi, and G. Spelman. Estimating response bounds for uncertain friction-damped structures. In *Proceedings of ISMA 2016 International Conference on Noise and Vibration Engineering*, Leuven, 2016.

[13] T. Butlin, G. Spelman, P. Ghaderi, W.J.B. Midgley, and R. Umehara. Predicting response bounds for friction-damped gas turbine blades with uncertain friction coupling. *Journal of Sound and Vibration*.

[14] N. Wiener. The homogeneous chaos. *Amer. J. Math*, 60897(4):936, 1938.

[15] G. Stefanou. The stochastic finite element method: Past, present and future. *Computer Methods in Applied Mechanics and Engineering*, 198(9-12):1031–1051, 2009.

[16] M. Hadigol and A. Doostan. Least squares polynomial chaos expansion: A review of sampling strategies. *Computer Methods in Applied Mechanics and Engineering*, 332:382–407, 2018.

[17] J. Didier, J. J. Sinou, and B. Faverjon. Nonlinear vibrations of a mechanical system with non-regular nonlinearities and uncertainties. *Communications in Nonlinear Science and Numerical Simulation*, 18(11):3250–3270, 2013.

[18] T. Roncen, J. J. Sinou, and J. P. Lambelin. Experiments and numerical simulations of nonlinear vibrations of a beam with non-ideal boundary conditions and uncertainties. *Procedia Engineering*, 199:808–813, 2017.

[19] A.M. Panunzio, L. Salles, C. W. Schwingshackl, and M. Gola. Asymptotic numerical method and polynomial chaos expansion for the study of stochastic non-linear normal modes. In *Proceedings of the ASME Turbo Expo*, volume 7B, pages 1–8, 2015.

[20] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

[21] C. Soize. Stochastic modeling of uncertainties in computational structural dynamicsRecent theoretical advances. *Journal of Sound and Vibration*, 332(10):2379–2395, may 2013.

[22] J. Legault, R.S. Langley, and J. Woodhouse. Physical consequences of a nonparametric uncertainty model in structural dynamics. *Journal of Sound and Vibration*, 331(25):5469–5487, dec 2012.

[23] E Skudrzyk. *Simple and complex vibratory systems*. Pennsylvania State University Press, Pennsylvania, 1968.

[24] A. Cabboi, T. Putelat, and J. Woodhouse. The frequency response of dynamic friction: Enhanced rate-and-state models. *Journal of the Mechanics and Physics of Solids*, 92:210–236, 2016.

[25] E. P. Petrov and D. J. Ewins. Analytical Formulation of Friction Interface Elements for Analysis of Nonlinear Multi-Harmonic Vibrations of Bladed Disks. *Journal of Turbomachinery*, 125(2):364, 2003.

[26] C. Gastaldi and T. M. Berruti. A method to solve the efficiency-accuracy trade-off of multi-harmonic balance calculation of structures with friction contacts. *International Journal of Non-Linear Mechanics*, 92(May 2016):25–40, 2017.