JASA ARTICLE

In-situ inference of the thermoacoustic properties of an industrial combustion system[a)]

Matthew Yoko,[1] Claus Lahiri,[2] André Fischer,[2] and Matthew P. Juniper[1,b)]
[1]*Department of Engineering, University of Cambridge, Cambridge, United Kingdom*
[2]*Rolls-Royce Deutschland, Dahlewitz, Germany*

ABSTRACT:
A Bayesian framework is proposed for building and calibrating physics-based models of industrial thermoacoustic systems, using the Rolls-Royce SCARLET test rig as a case study. Several candidate models are constructed, and their uncertain parameters are inferred directly from experimental data. Bayesian model comparison is then used to identify the most probable candidate model, balancing data fit and model complexity. The selected physics-based model reproduces both non-reacting and reacting measurements with high precision. This model is then used to confirm the findings of recent work, which demonstrates an inconsistency in a commonly used method for measuring the flame response in complex combustion chambers. This paper goes further to provide an improved method for identifying the flame response from data. Moreover, because this process learns the parameters of a flame model rather than just processing the experimental data, the model can interpolate and extrapolate, and provide deeper insight into the underlying physics.
© *2026 Acoustical Society of America*. https://doi.org/10.1121/10.0043969
(Received 19 December 2025; revised 9 April 2026; accepted 17 May 2026; published online 8 June 2026)
[Editor: Ning Xiang] Pages: 5150–5169

I. INTRODUCTION

Thermoacoustic oscillations arise in combustion systems when an unsteady heat release rate interacts constructively with acoustic waves in the combustion chamber, leading to self-sustained pressure fluctuations. In practical combustors, such as gas turbines or rocket engines, this feedback can amplify small perturbations, posing risks of excessive noise, vibration, and even structural failure. From the perspective of acoustics, thermoacoustics extends the familiar modelling framework to include a nonlinear source term associated with fluctuating heat release rate. This "flame response" source term describes how the heat release rate of the flame fluctuates in response to acoustic velocity perturbations. To represent this coupling in a tractable way, it is common to approximate the flame response using a linear transfer function, known as the flame transfer function (FTF).

Physics-based models of thermoacoustic systems are typically qualitatively accurate, but not quantitatively accurate. This is often because thermoacoustic oscillations are extremely sensitive to the values of unknown model parameters, chiefly the flame response (Juniper and Sujith, 2018). Small errors in these parameters can lead to large errors in the model predictions, but we rarely know the values of these parameters to sufficient precision *a priori*.

In most cases, the flame response is obtained through direct measurement. This requires (i) a measurement of the acoustic velocity, either using a hot wire (Gatti *et al.*, 2018; Kornilov *et al.*, 2007; Mejia *et al.*, 2016), optical methods (Birbaud *et al.*, 2006; Cuquel *et al.*, 2013; Ducruix *et al.*, 2000), or the multi-microphone method (AEsøy *et al.*, 2022; Nygård and Worth, 2021; Seybert and Ross, 1977), and (ii) a measurement of the perturbed heat release rate, typically using optical methods (Birbaud *et al.*, 2006; Cuquel *et al.*, 2013; Ducruix *et al.*, 2000; Gatti *et al.*, 2018; Kornilov *et al.*, 2007; Mejia *et al.*, 2016). Most of these methods are limited to measuring the flame response under laboratory conditions and on simplified geometries because (i) industrially relevant geometries rarely permit optical access, (ii) the operating conditions in industrially relevant combustion chambers are too harsh for the delicate instruments that are required to perform these measurements, and (iii) current optical methods for measuring heat release rate perturbations are not accurate for industrially relevant flames, such as technically premixed spray flames (Peterleithner *et al.*, 2016; Schuermans *et al.*, 2010; Shreekrishna and Lieuwen, 2013).

In order to measure the flame response in industrially relevant conditions, it is common to use indirect methods that only require microphone measurements. The most widely used of these is the two-source method (Munjal and Doige, 1990; Paschereit *et al.*, 1999; Schuermans *et al.*, 1999). By forcing the system from the upstream and downstream boundaries and measuring the acoustic response using microphones, the transfer matrix of the system can be computed. From the transfer matrices measured under non-reacting and reacting conditions, the FTF can be extracted.

Rolls-Royce developed the SCARLET (SCaled Acoustic Rig for Low Emission Technology) test rig to

[a)]This paper is part of a special issue on Bayesian Methods in Acoustical Science and Engineering.
[b)]Email: mpj1001@cam.ac.uk

 © 2026 Acoustical Society of America

https://doi.org/10.1121/10.0043969

directly measure flame responses of industrially relevant burners under realistic operating conditions. The rig, which is described in detail in Sec. II, uses the two-source method to measure the FTF. In recent work, however, Eder *et al.* (2025) demonstrated that the two-source method, called the burner-flame transfer method (BFTM) in Eder *et al.* (2025), introduces artefacts in the measured FTF when applied to complex combustor geometries, such as the one in SCARLET. The artefacts arise because the method assumes that the acoustic behaviour of the chamber is the same under non-reacting and reacting conditions. This is not true for industrially relevant chambers, which feature area changes and perforated liners between the flame and the measurement section, the acoustic response of which changes with temperature.

Eder *et al.* (2025) proposed a model-based inference approach to correct this error, in which a low-order model of the rig is used to account for the variation in acoustic behaviour of the chamber between the non-reacting and reacting cases. The experimental data is then used to close the model, yielding an improved estimate of the flame response, but only at the tested frequencies. The authors noted a residual error at high frequencies, which they ascribed to errors in the physics-based model.

In this paper, we propose an alternative approach for obtaining the flame response from rigs such as SCARLET. Like Eder *et al.* (2025), we construct a detailed low-order, physics-based model of the rig. Unlike Eder *et al.* (2025), who use the data to close their model directly, we identify the model parameters that we do not know to sufficient precision *a priori*, and use a rigorous Bayesian framework to infer these from the experimental data. This means that our model, once trained on the data, produces a noise-free, fully interpretable reconstruction of the data, and can interpolate and extrapolate to unseen frequencies and operating conditions. It is often the case that the model parameters are a function of the operating condition. In this case we must also learn a higher-level function that maps operating condition to model parameters, which is the subject of a companion paper (Yoko *et al.*, 2026).

In previous work, we proposed and demonstrated a Bayesian framework for assimilating experimental data into models of thermoacoustic systems. We developed the framework by applying it to several systems with incremental steps in complexity: an electrically heated Rijke tube (Juniper and Yoko, 2022), a laminar conical flame in a duct (Yoko and Juniper, 2024a), and a weakly turbulent conical flame in a duct (Yoko and Juniper, 2024b). In the last of these papers, we compared the inferred flame response to direct measurements from chemiluminescence and found them to be almost identical. Using these simple systems, we have therefore demonstrated this approach as a viable method for inferring the flame response with the flame *in situ* using only microphone measurements, which we now apply to an industrial test rig.

II. EXPERIMENTAL CONFIGURATION

The SCARLET test rig, illustrated in Fig. 1, consists of four key components: the combustion chamber, measurement sections, sirens, and dampers. The combustion chamber is a modular, single can combustor with an easily replaceable fuel spray nozzle. It includes the major features that would be found in a modern aero-engine, such as a bypass annulus and an effusion-cooled liner. On either end of the combustion chamber is a measurement section, which is a circular duct fitted with microphones at five axial locations. At the end of each measurement section, two sirens are coupled to the rig to provide acoustic forcing. Finally, a damper is fitted on either end of the sirens to minimize the reflection of acoustic waves back into the rig. At the downstream end of the rig there is a back-pressure valve, which provides control over the pressure in the rig.

The rig operates at realistic engine conditions, with inlet pressures up to 30 bar, and inlet temperatures up to 950 K. The maximum inlet mass flow rate is 4 kg/s, and the combustor pressure drop can be varied between 2% and 5% using the back-pressure valve. At any given operating condition, the rig is forced sequentially in the upstream and downstream directions using the sirens. The sirens produce sinusoidal excitation over a range of amplitudes for frequencies up to 1.2 kHz, which is the theoretical limit for plane wave propagation in the measurement sections at ambient conditions. The excitation amplitude can be varied by shifting the phase relation between the two sirens of each pair. Following a forcing test, the recorded microphone measurements are Fourier-decomposed to obtain the complex pressure at each microphone position, which we assimilate into our physics-based model of the rig.

III. PHYSICS-BASED MODEL OF SCARLET

We model the SCARLET rig using an acoustic network modelling code that was written to carry out this work. This study required a code that (i) could handle the geometric complexity of the rig, including branch splits and junctions, swirl injectors, and effusion-cooled panels, (ii) included the

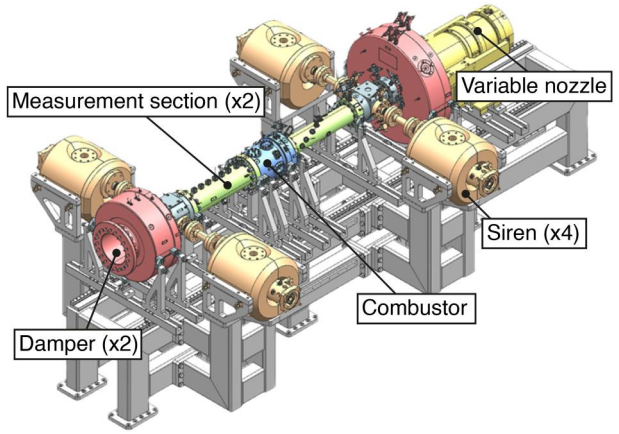

FIG. 1. Diagram of the SCARLET rig highlighting the major components.

effects of an appreciable mean flow, and (iii) was differentiable with respect to all model parameters.

The code follows the typical network modelling approach, where a complex geometry is decomposed into a series of ducts that are connected by jump conditions. At present, the code contains jump conditions to describe: area increases and decreases; branch splits and junctions; perforated liners; swirl injectors, and various models for the flame. The code is modular, so new jump conditions can be added easily. The equations solved at each jump and boundary condition are summarized in Appendix A.

Once the network is assembled, the jump conditions and boundary conditions are solved simultaneously to calculate either (i) the steady flow in each duct, (ii) the forced acoustic response of the system, or (iii) the unforced acoustic response of the system. Importantly for our inference framework, it is possible to cheaply obtain the derivative of the steady and acoustic flow variables with respect to any model parameter. The derivatives are calculated using a combination of a custom automatic differentiation class and the adjoint method. Adjoints are used in parts of the code that the user is unlikely to modify, while automatic differentiation is used in the user-modifiable parts of the code. This allows new users to extend the code without having to derive adjoints for their modifications.

An example network for SCARLET is sketched in Fig. 2. The majority of the geometric features are modelled according to standard network modelling practice, as outlined by Dowling and Stow (2003). There is, however, a lack of consensus on how to model swirl injectors and perforated liners, and there are many plausible models for the flame response. To resolve this, we propose several candidate models for each of these features, summarized in the following, and rank them against each other using Bayesian model comparison to select the best model for each feature, given the experimental data.

For the swirl injector, the baseline model is a plain orifice with an unknown effective area and no pressure recovery at the sudden expansion (Dowling and Stow, 2003). We then consider models that additionally include any combination of (i) partial or full pressure recovery at the sudden expansion, (ii) momentum and/or energy loss due to drag on the vanes, and (iii) inertia of the fluid slug in the injector. For the perforated liner, we compare three candidate models: (i) lumping patches of holes together and treating them as orifices Eder *et al.* (2025), (ii) Howe's Rayleigh conductivity model (Howe, 1979), and (iii) Howe's model with thickness correction (Jing and Sun, 1999). We note that some studies have additionally shown liner damping to be sensitive to the grazing flow, which the previously mentioned models neglect (Bauer, 1977; Howe *et al.*, 1996; Rao and Munjal, 1986). However, this effect scales with grazing flow Mach number, which is $\mathcal{O}\left(10^{-2}\right)$ considered for all cases in this paper. The effect of grazing flow is therefore negligible, and the data contains no information to infer the parameters of a grazing flow correction, or to discriminate between models with and without this effect. For the flame,

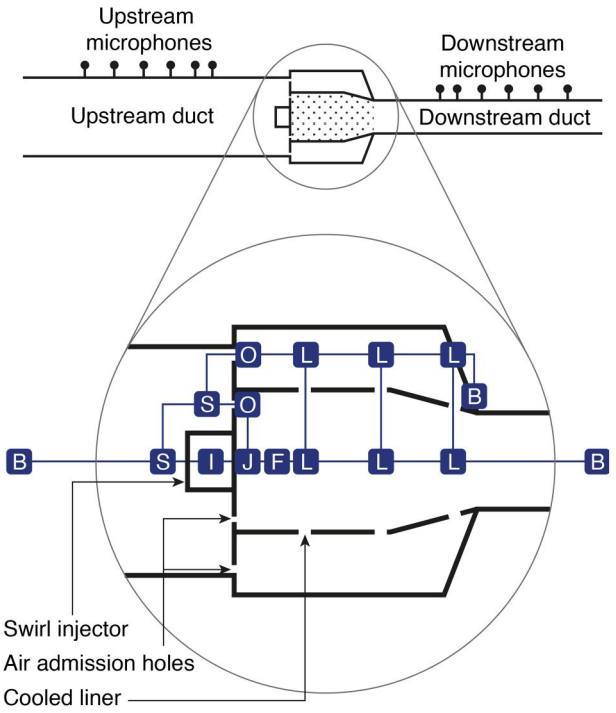

FIG. 2. Sketch of an acoustic network configuration for modelling SCARLET. The network consists of boundary conditions (B), and jump conditions including branch splits (S) and junctions (J), orifices (O), which are just area decreases followed by increases, and special elements for the injector (I), liner (L) and flame (F).

we compare discrete distributed time delays with one to four delays (Polifke, 2020) and Gaussian distributed time delays with one to four delays (Alemela *et al.*, 2010).

We supply boundary conditions for the acoustic waves in the form of reflection coefficients at the unforced end and an incoming wave of unit strength at the forced end. When forcing upstream, we additionally supply a downstream boundary condition to describe the conversion of entropy waves to acoustic waves. We assume zero incoming entropy fluctuations at the upstream boundary.

In previous work on laboratory rigs in unforced conditions, we calculated the acoustic reflection coefficients from suitable analytical expressions (Juniper and Yoko, 2022; Yoko and Juniper, 2024a). The end terminations of the SCARLET rig, however, are too complex to model analytically or numerically. We therefore calculate the reflection coefficients and their uncertainties, directly from the microphone data as it becomes available. As in the two-source method, this uses a portion of the available data to characterize the boundary conditions, but forcing from both ends provides sufficient degrees of freedom to determine both the reflection coefficients and the model parameters.

If we want to evaluate the model at a frequency or test condition for which data is not available, we must estimate the reflection coefficients based on previously studied conditions. For this, we use a Gaussian process model, which is trained on several years' worth of test data on the SCARLET rig, covering a wide range of test conditions. We note, however, that the most valuable output of the model is

https://doi.org/10.1121/10.0043969

the scattering or transfer matrix of the burner, which is independent of the reflection coefficients, so the fidelity of the Gaussian process model for this interpolation is not critical. Whether the reflection coefficients are estimated from the data or from the Gaussian process, their uncertainty is propagated to the model predictions.

IV. BAYESIAN DATA ASSIMILATION

The Bayesian data assimilation framework that we use has been described in detail in previous papers (Juniper and Yoko, 2022; Yoko and Juniper, 2024a). Here, we provide a brief overview of the framework.

We propose several plausible models of the rig, which we call $\mathcal{H}_i$, each of which contains a set of parameters, $\boldsymbol{a}$, about which we have incomplete knowledge. We encode our lack of knowledge about these parameters by describing them with probability distributions rather than point estimates. We then conduct data assimilation on two levels: parameter inference and model selection. At the parameter inference level, we assume that each candidate model is structurally correct and attribute any model-data mismatch to uncertainty in the parameters. We then use experimental data to perform a Bayesian update on the parameter probability distributions, thereby learning an improved estimate of the parameter values from the data. At the model comparison level, we relax the assumption that each candidate model is correct and calculate the likelihood of each model, given the data. We then select the model with the highest likelihood as the best model for this case. We now briefly describe the details of each level.

A. Parameter inference

For a given candidate model, $\mathcal{H}_i$, we wish to infer the most probable values of the model parameters $\boldsymbol{a}$, given some experimental data, D. We begin by declaring our state of knowledge about the parameters before seeing the data, which we call the prior distribution, $P(\boldsymbol{a}|\mathcal{H}_i)$. In this work we use a multivariate normal distribution for the prior,

$$P(\boldsymbol{a}|\mathcal{H}_i) = \mathcal{N}(\boldsymbol{a}_p, \boldsymbol{C}_a), \tag{1}$$

where $\boldsymbol{a}_p$ is the prior mean, and $\boldsymbol{C}_a$ is the prior covariance matrix.

We propose a pragmatic approach to defining the prior distribution, capitalizing on the engineer's ability to estimate and bound physical quantities. For each parameter, we define the values that we are 99% confident the parameter cannot be (i) larger than and (ii) smaller than. These two values are easy to estimate and fully define a Gaussian distribution. We use the mean of the values as our prior mean, and the range defines the six-sigma range for the 99% confidence interval of a Gaussian, allowing us to compute the diagonal entries of the covariance matrix. We have no prior information about the correlations between parameters, so we set the off-diagonal terms of the covariance matrix to zero.

Next, we define a noise model for the data, which allows us to calculate the likelihood of observing the data, given the model parameters, $P(D|\boldsymbol{a}, \mathcal{H}_i)$. In this work we assume that the data is corrupted by additive Gaussian noise with zero mean and covariance $\boldsymbol{C}_d$, which is often the best model for experimental measurement noise (Jaynes, 2003, Sec. 19.2.1). This gives a likelihood function, which is a Gaussian distribution centered around the model output,

$$P(D|\boldsymbol{a}, \mathcal{H}_i) = \mathcal{N}(s(\boldsymbol{x}, \boldsymbol{a}), \boldsymbol{C}_d), \tag{2}$$

where $s(\boldsymbol{x}, \boldsymbol{a})$ is the output of the model for inputs $\boldsymbol{x}$ and parameters $\boldsymbol{a}$, and $\boldsymbol{C}_d$ is the covariance matrix of the data noise. In this work, the model inputs are the forcing and operating conditions, while the model outputs are the Fourier-decomposed pressures at each microphone position.

We estimate the data noise covariance, $\boldsymbol{C}_d$, from the residual variance between the measured time-series and the time-series reconstructed from the Fourier-decomposed signal at each forcing frequency. This assumes that the uncertainty in the Fourier-decomposed microphone measurements is dominated by incoherent broadband noise, which is probably the case given the highly turbulent mean flow. The resulting noise level is approximately frequency-independent.

Having defined our prior and likelihood functions, we use Bayes' rule to calculate the posterior distribution of the parameters, given the data,

$$P(\boldsymbol{a}|D, \mathcal{H}_i) = \frac{P(D|\boldsymbol{a}, \mathcal{H}_i)P(\boldsymbol{a}|\mathcal{H}_i)}{P(D|\mathcal{H}_i)}. \tag{3}$$

The most probable parameters are those which maximize the left-hand side. When identifying the most probable parameters, we can neglect the denominator of Eq. (3) because it does not depend on the parameters, and we maximize the numerator. We cheaply find the peak of the posterior by reframing the inference problem as an optimization problem. We define a cost function, $\mathcal{J}$, to be the negative log of the numerator from Eq. (3), which we minimize using a gradient-based minimization algorithm, with gradients calculated automatically by the network code. The resulting parameter estimate is called the maximum *a posteriori* (MAP) estimate.

Once we have found the peak of the posterior distribution, we estimate the posterior uncertainty in the parameter values using Laplace's method (Jeffreys, 1973; MacKay, 2003). This method approximates the posterior distribution by performing a second-order Taylor expansion around the peak of the true posterior, yielding a Gaussian approximation centred at the MAP point, with posterior covariance

$$\boldsymbol{C}_{a|D}^{-1} = \boldsymbol{C}_a^{-1} + \boldsymbol{J}^\top \boldsymbol{C}_d^{-1} \boldsymbol{J}, \tag{4}$$

where $\boldsymbol{J}$ is the Jacobian matrix of the model outputs with respect to the parameters, evaluated at the MAP point.

B. Model selection

Once the posterior distribution of each candidate model's parameters has been found, we calculate the posterior probability of each model itself. This allows us to quantitatively rank the candidate models against each other. We once again use Bayes' rule to calculate the posterior probability of each model,

$$P(\mathcal{H}_i|D) \propto P(D|\mathcal{H}_i)P(\mathcal{H}_i). \tag{5}$$

The first factor on the right-hand side of Eq. (5) is the denominator of Eq. (3), which is called the marginal likelihood, or the evidence. The second factor is the prior probability that we assign to each model. If we have no reason to prefer one model over another *a priori*, we assign equal prior probabilities to all models and rank them according to their marginal likelihood. Recalling that we have approximated the posterior distribution with a Gaussian, we can cheaply compute the marginal likelihood,

$$P(D|\mathcal{H}_i) \approx P(D|\boldsymbol{a}_{MAP}, \mathcal{H}_i) \times P(\boldsymbol{a}_{MAP}|\mathcal{H}_i)|2\pi\boldsymbol{C}_{a|D}|^{1/2}, \tag{6}$$

from which we see that the marginal likelihood can be decomposed into two factors. The first factor, called the best-fit likelihood (BFL), measures how well the model fits the data, and rewards the model for a good fit. The second factor, called the Occam factor (OF), penalizes the model for its complexity. This process, therefore, naturally identifies the simplest model that is capable of describing the data. Further details can be found in MacKay (2003).

It is worth expanding on the OF to investigate how it penalizes model complexity. Taking the logarithm of the OF and substituting in the Gaussian function for the prior, evaluated at the MAP point, we have,

$$\log \mathrm{OF} = -\frac{1}{2}(\boldsymbol{a}_{MAP} - \boldsymbol{a}_p)^\top \boldsymbol{C}_a^{-1}(\boldsymbol{a}_{MAP} - \boldsymbol{a}_p) - \log \frac{|\boldsymbol{C}_a|^{1/2}}{|\boldsymbol{C}_{a|D}|^{1/2}}. \tag{7}$$

The first term penalizes models whose MAP parameters are far from the prior mean, i.e., we prefer models for which we have precise prior information. To understand the second term, we note that the square root of the determinant of a covariance matrix is proportional to the volume of the uncertainty ellipsoid defined by the covariance. The second term, therefore, penalizes models where the posterior accessible volume is much smaller than the prior accessible volume, i.e., models for which the parameters must be delicately tuned to fit the data.

We can gain further insight into the second term by substituting in Eq. (4) for the posterior covariance,

$$-\log \frac{|\boldsymbol{C}_a|^{1/2}}{|\boldsymbol{C}_{a|D}|^{1/2}} = -\log |\boldsymbol{C}_a|^{1/2}|\boldsymbol{C}_a^{-1} + \boldsymbol{J}^\top \boldsymbol{C}_d^{-1}\boldsymbol{J}|^{1/2}. \tag{8}$$

By the matrix determinant lemma, we can rewrite this as

$$-\log |\boldsymbol{C}_a|^{1/2}|\boldsymbol{C}_d^{-1} + \boldsymbol{J}\boldsymbol{C}_a\boldsymbol{J}^\top|^{1/2}|\boldsymbol{C}_d^{-1}|^{1/2}|\boldsymbol{C}_a^{-1}|^{1/2} = -\log \frac{|\boldsymbol{C}_d + \boldsymbol{C}_s|^{1/2}}{|\boldsymbol{C}_d|^{1/2}}, \tag{9}$$

where $\boldsymbol{C}_s = \boldsymbol{J}\boldsymbol{C}_a\boldsymbol{J}^\top$ is the covariance of the model outputs induced by the prior uncertainty in the parameters. The denominator, $|\boldsymbol{C}_d|^{1/2}$, is the volume of the solution space that the data occupies, while the numerator, $|\boldsymbol{C}_d + \boldsymbol{C}_s|^{1/2}$, is the volume of the solution space that the model can access, given the prior uncertainty in the parameters. The OF therefore penalizes elaborate, flexible models that can access a large volume of solution space, relative to the volume occupied by the data. These models could describe many different datasets, so they are penalized for their lack of specificity.

V. RESULTS

We infer the unknown parameters in two steps. In the first step, we use non-reacting experiments to infer the acoustic parameters of the rig. We then fix these parameters and use data from reacting experiments to infer the parameters of the flame response model.

As mentioned before, we infer the parameters by assimilating the Fourier-decomposed pressure measurements at each microphone into our model. For brevity, in the body of this paper, we show the more concise scattering matrix representation of the data fit, while the fit of the Fourier-decomposed pressure at each microphone is shown in Appendix B. We additionally show the transfer matrix representation of the data fit in Appendix C.

When computing derived quantities such as scattering or transfer matrices from the microphone measurements, we propagate the measurement uncertainty through the matrix inversion. The resulting uncertainty is frequency-dependent because the condition number of the inversion depends on, among other things, the ratio of microphone spacing to acoustic wavelength. For the transfer matrix in particular, this condition number is typically poor, amplifying the microphone measurement uncertainties, as seen in Appendix C. This motivates our preference for scattering matrices throughout this paper.

A. Application to non-reacting data

Our inference framework requires that we first infer the most probable parameters for each candidate model, then rank the models against each other. We will report on this in the reverse order, first showing the outcome of model ranking so that we can report on parameter inference for only the most likely model.

1. Outcome of model ranking

We compare ten candidate models for the swirl injector. As described previously, the injector models include any combination of the effects of (i) sudden area changes, (ii) partial or full pressure recovery, (iii) drag as a source in the momentum and/or energy equations, and (iv) the inertia of

https://doi.org/10.1121/10.0043969

the slug of fluid in the injector orifice. The ten combinations of these effects that we study are summarized in Table I. Figure 3 shows the model ranking metrics for the ten candidate models, measured relative to the most likely model (model 4).

We are particularly interested in the log of the marginal-likelihood ratio between each model and the most likely model, which is shown in dark blue in Fig. 3. This quantifies how much less likely each model is than the most likely model, given the data. Jeffreys (1983, p. 432) and Kass and Raftery (1995, Sec. 3.2) provide some guidance on how to interpret these values: for a natural logarithm, log-ratios between 0 and −1 are barely worth mentioning; log-ratios between −1 and –3 indicate moderate evidence against the weaker model; log-ratios between –3 and −5 indicate strong evidence against the weaker model, and log-ratios more negative than −5 indicate very strong evidence against the weaker model.

From Fig. 3, we see that model 4 (sudden contraction and expansion with full pressure recovery, and drag in both the momentum and energy equations) has the most evidence in the data. We note, however, that model 9, which additionally includes the inertia of the fluid in the injector, is only slightly less likely than model 4, with a log-marginal likelihood ratio of −1.3 (only just outside the threshold of "barely worth mentioning"). Models 4 and 9 are adaptations of the popular $L-\zeta$ model (Schuermans *et al.*, 1999), which has been used in previous work to model the injector in the SCARLET rig, particularly with L set to zero (similar to model 4) (Eder *et al.*, 2025).

It is also worth mentioning that model 7 fits the data the best (largest BFL), and only has two unknown parameters. This would make it seem like a good candidate, but we find it is much less likely than model 4 (large, negative log-marginal likelihood). This is because in order to fit the data, the parameters for this model must be driven to values that we believed *a priori* were extremely unlikely (an effective area less than 10% of the geometric area, and an effective length more than twice as long as the geometric length). The OF penalizes this heavily, preventing this model from being selected as the most likely model.

Finally, we note that models that assume full pressure recovery ($R=1$) consistently outperform those that assume no pressure recovery ($R=0$), even though the latter is commonly used to model swirl injectors. This is because assuming no pressure recovery asserts that the flow exits the injector as a free jet, which is not representative of the flow field in a strongly swirling flow. These flows typically feature recirculation zones and regions of decelerated flow that enable at least partial pressure recovery. We therefore find that the model with the most evidence, given this data, assumes total pressure recovery and implicitly lumps any losses at the expansion with the drag on the swirl vanes.

The same model ranking process is applied to selecting the best model for the perforated liner, where the candidate models are (i) lumping patches of holes together and treating them as orifices (Eder *et al.*, 2025), (ii) Howe's original model, which assumes an infinitely thin liner (Howe, 1979), and (iii) Howe's model with thickness correction (Jing and Sun, 1999). In this case, we find that Howe's model with thickness correction is the most likely model. Relative to this, the log-marginal likelihood ratio is –3.6 for Howe's original model (strong evidence against), and –406 for the orifice model (very strong evidence against). The poor performance of the orifice model is therefore a likely reason for the model discrepancies noted by Eder *et al.* (2025) in their attempts to model SCARLET.

As discussed in Sec. IV B, the model-ranking procedure depends on the prior through the OF. Because the prior encodes our current state of knowledge, which is subjective, we assess the sensitivity of the model ranking by increasing and decreasing the prior covariance by 50%. This is most relevant for the injector comparison, where the primary discriminant between models 4 and 9 is the OF. Increasing the prior covariance (weaker prior information) does not change the ranking, whereas decreasing it (stronger prior information) slightly favours model 9 over model 4. We nevertheless retain model 4 as the preferred injector model, since we have no reason to adopt a tighter prior. For the liner comparison, the data provide strong evidence in favour of Howe's model with thickness correction, and the ranking is unchanged under either scaling.

TABLE I. Summary of the ten swirl injector models. Each model considers the effective area of the injector (A_{eff}) as an unknown parameter (a). In addition to this, models can include pressure recovery at the sudden expansion (R), which can either be inferred from the data or set to zero (no pressure recovery) or one (full pressure recovery). The drag of the swirl vanes can be included as a source in the momentum (D_M) and/or energy (D_E) equations. Finally, the inertia of the fluid slug in the injector (I) can be set to zero or inferred from the data.

Model	Params	A_{eff}	R	D_M	D_E	I
1	1	a_1	0	0	0	0
2	2	a_1	a_2	0	0	0
3	2	a_1	0	a_2	a_2	0
4	2	a_1	1	a_2	a_2	0
5	2	a_1	1	a_2	0	0
6	2	a_1	0	0	0	a_2
7	2	a_1	1	0	0	a_2
8	3	a_1	0	a_2	a_2	a_3
9	3	a_1	1	a_2	a_2	a_3
10	4	a_1	a_2	a_3	a_3	a_4

2. Inferring the parameters of the best model

There are five acoustic parameters that we learn from the non-reacting data, which are (i) the effective area of the injector, (ii) the drag coefficient of the swirl vanes, (iii) and (iv) the effective areas of two sets of air admission holes (called the heat shield cooling holes and annulus metering holes), and (v) the effective area of the liner cooling holes.

To develop confidence in our approach, we initially learned these parameters sequentially by performing a series of tests with the rig in the following conditions: (1a) the

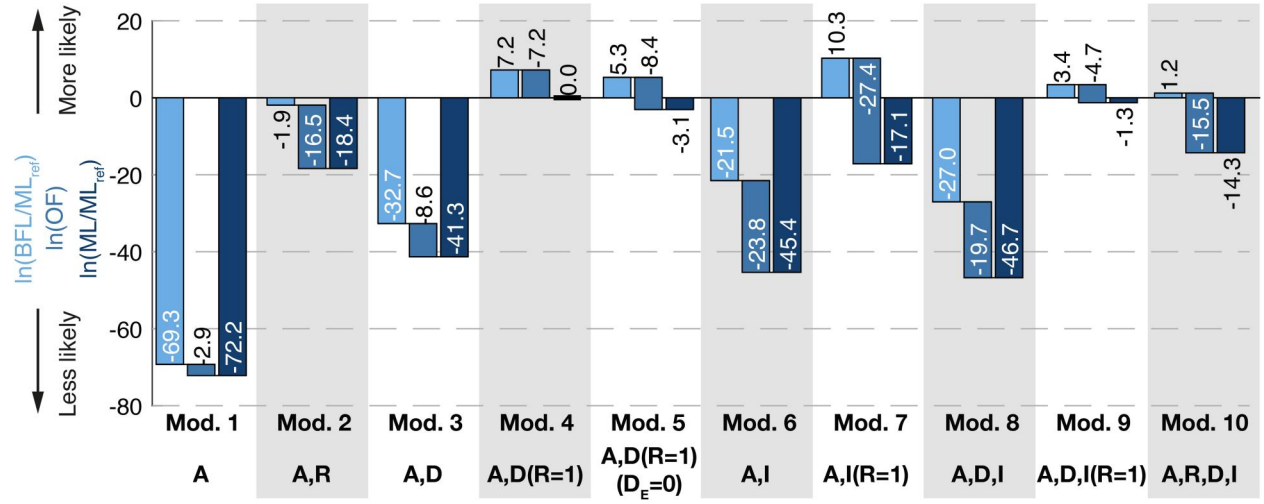

FIG. 3. Model ranking metrics for ten candidate models for the acoustic response of the swirl injector. The ten models are summarized in Table I, and the letters below each model indicate which parameters are inferred. If a parameter is fixed to a non-standard value, this is shown in brackets. The model ranking metrics are decomposed into three quantities: the log-BFL [ln(BFL)] rewards a model for a good data fit, the log-OF [ln(OF)] penalizes a model for its complexity, and the log-marginal likelihood [ln(ML)] gives the overall likelihood for the model. We rank the models based on their relative ln(ML) and choose the reference ln(ML) to be that of the most likely model. The ln(ML) bar for each model, therefore, shows how much less likely each model is than the most likely model. For example, model 9 is $e^{-1.3} = 0.27$ times as likely as model 4.

swirl injector was replaced with a plain orifice, and the liner, heat shield cooling holes, and annulus metering holes were taped shut, allowing us to learn the orifice effective area with a simple physical model, (1b) the tape blocking the heat shield cooling holes and annulus metering holes was removed, allowing us to learn the effective areas of these features, (1c) all tape was removed, allowing us to learn the effective area of the liner holes. Experiments (2a-c) repeat (1a-c) with the swirl injector in place, allowing us to learn the parameters of the swirl injector model. While this approach was invaluable for developing and validating the modelling framework, we find that it is possible to learn unique parameters for all features of the rig from a single experiment, which is a more practical result, so we report only this.

To define the prior distribution for the five parameters, we use the approach described in Sec. IV A. We specify a plausible range for the contraction coefficients of each of the orifices to be between 0.3 and 0.9, and the drag coefficient of the vanes to be between 0 and 10.

We show the prior and posterior model fit in Fig. 4. The prior model prediction (gray) already fits the data well, with only relatively minor discrepancies. This demonstrates that, because the pure acoustic problem is not highly sensitive to the model parameters, a good model produces reasonable predictions with relatively poor parameter estimates. We see from the prior prediction uncertainty bounds, however, that the uncertainty in the parameters drives large uncertainty in some of the matrix elements, particularly the reflection of the forward travelling wave (r+).

The posterior prediction is noticeably improved, with the predicted scattering matrix falling within the error bars of the vast majority of the data points. While it might be tempting to accept the small deviation of the prior model, we will see later that this small deviation can cause a significant error when inferring the flame response.

3. Sensitivity of the parameters to operating conditions

When we move to reacting to data, the conditions inside the rig change drastically. It is therefore important to understand how sensitive the model parameters are to the conditions inside the rig. To do so, we assess data from 20 different non-reacting operating conditions. We use data from a baseline condition to learn the parameters of the non-reacting model. We then compare the predictions of the model against the other 19 conditions. The baseline and range of the tested operating conditions are summarized in Table II.

Figure 5 shows the scattering matrices for the 20 datasets and the model. When plotted against the Helmholtz number, the experimental datasets collapse onto one another, indicating that the dominant acoustic difference between the cases arises from the change in sound speed due to the temperature difference. This suggests that the remaining acoustic properties of the rig are largely insensitive to the operating condition.

We confirm this by overlaying the model trained on the baseline dataset, which fits all datasets equally well. Notably, the trained model extrapolates in Helmholtz number by almost a factor of two, and remains reasonably accurate. This result confirms that even substantial variations in internal conditions do not significantly affect the rig's drag or discharge coefficients.

B. Application to reacting data

Having characterized the non-reacting rig, we now use the same approach to learn the flame response from reacting data. Motivated by the insensitivity of the acoustic parameters to the conditions in the rig, we fix the acoustic

https://doi.org/10.1121/10.0043969

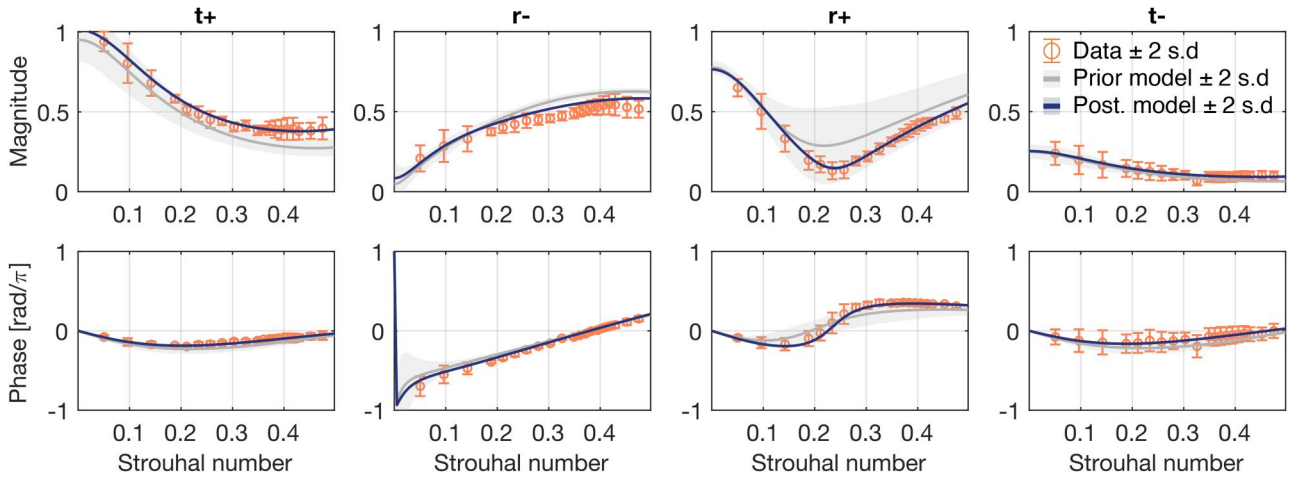

FIG. 4. Scattering matrix of the non-reacting SCARLET rig, showing the (orange) measured, (gray) prior model predictions and (blue) posterior model predictions, each of which is presented with a 95% confidence interval. The top row shows the magnitude of the matrix entries, while the bottom row shows the phase, each plotted against Strouhal number. The columns show (t+) the transmission of the forward travelling wave, (r−) the reflection of the backward travelling wave, (r+) the reflection of the forward travelling wave, (t−) and the transmission of the backward travelling wave.

parameters that we learned in Sec. V A 2 , i.e., we assume that the four effective areas and the swirl vane drag coefficient do not change between the reacting and non-reacting case. The results in Fig. 5 strongly support this assumption, but they do not test whether these parameters are affected by the temperature gradient introduced by the flame. We note, however, that the parameters in question, discharge coefficients and a drag coefficient, are geometric properties of the orifices and swirl vanes, and are not expected to depend significantly on the local temperature gradient. The dominant acoustic effect of the temperature gradient is the impedance mismatch at the interface between the hot and cold regions, which is captured explicitly by the network model and does not rely on the values of these parameters. We therefore proceed with this assumption.

1. Outcome of flame model ranking

As mentioned in the introduction, we model the flame's response to acoustic forcing as a transfer function, $\mathcal{F}$, between velocity perturbations and heat release rate oscillations

$$\mathcal{F} = \frac{q'/\overline{q}}{u'/\overline{u}}, \tag{10}$$

where q is the heat release rate, u is the velocity, $(\cdot)'$ denotes perturbations, and $\overline{(\cdot)}$ denotes mean quantities.

TABLE II. Summary of the 20 non-reacting operating conditions that we use to evaluate the sensitivity of the model parameters to the operating condition.

Parameter	Units	Baseline	Min	Max
$\dot{m}$	[kg/s]	3.3	0.3	3.9
T	[K]	825	282	825
P	[bar]	23	1.3	25
ΔP	[%]	4.0	3.1	4.8

We compare eight candidate models for the FTF, grouped into two classes of distributed time delays: (i) discrete distributed time delays (DTD) with one to four delays, $\mathcal{F} = \sum_{i=1}^{N} n_i e^{(-s\tau_i)}$, and (ii) Gaussian distributed time delays (GTD) with one to four delays, $\mathcal{F} = \sum_{i=1}^{N} n_i e^{(-s\tau_i + 1/2s^2\sigma_i^2)}$. For each term in a sum of discrete time delays, we must infer two parameters: the interaction index, n_i, and the time delay, τ_i. For each term in a sum of Gaussian time delays, we must infer the same two parameters, plus a third: the distribution factor, σ_i.

The model ranking metrics for the eight candidate models are shown in Fig. 6, ranked relative to the most likely model, which is a discrete time delay model with three time delays (3 DTD). While a Gaussian time delay model consistently fits the data better than a discrete time delay model with the same number of delays (larger BFL), the Gaussian delay models are heavily penalized by the parametric flexibility offered by the distribution factor (more negative OF).

We note that this is not a general result for two reasons. First, it is specific to the burner studied here, which is characterized by a high injection velocity and a relatively compact flame, meaning perturbations traverse the flame in a much shorter time scale than the acoustic period. As will be seen later, this leads to a relatively simple FTF which is easily described by a sum of discrete delays. Second, the model ranking result is influenced by our prior knowledge, which is particularly poor for the spread factor of the Gaussian delay model because it is difficult to estimate *a priori*. Nevertheless, increasing or decreasing each model's prior parameter covariance by 50% does not change the ranking. For the remainder of this paper, we select the discrete time delay model with three delays (3 DTD) as the best model for this burner.

2. Inferring the flame response

The prior and posterior fit of the most likely flame model is shown in Fig. 7. Unlike the non-reacting case, we see that the prior model predictions are poor, and our prior

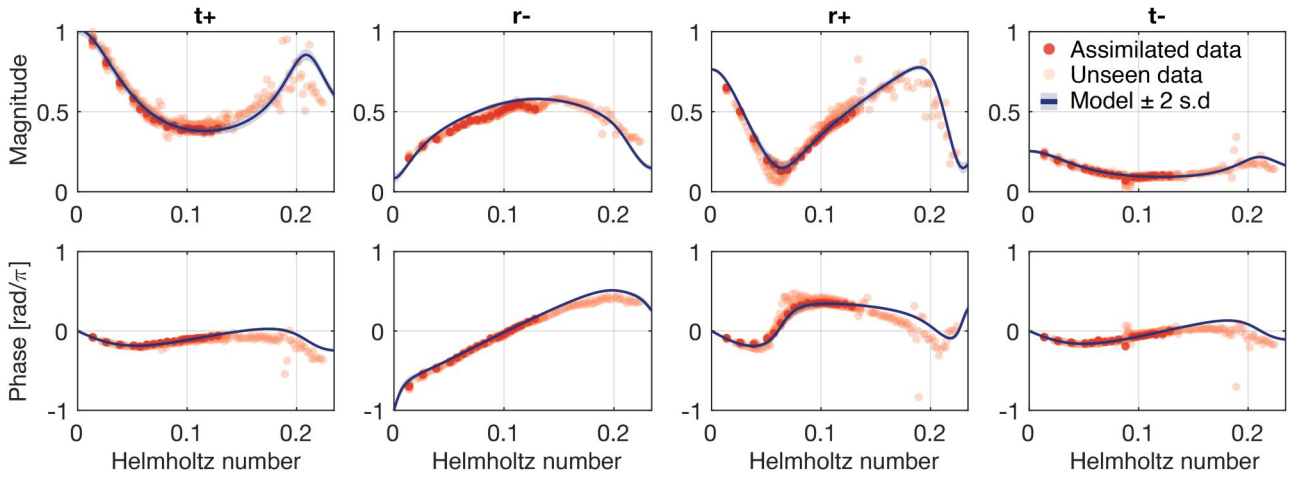

FIG. 5. Same as Fig. 4, but showing data collected from 20 operating conditions plotted against Helmholtz number. The model, trained on the data shown with darker shading, is overlaid.

uncertainty is large. This is because (i) the thermoacoustic model predictions are extremely sensitive to the values we choose for the flame model parameters, and (ii) we have poorer prior information about the flame model parameters than we did about the acoustic parameters. Nonetheless, after updating our knowledge about the flame model parameters using the experimental data, we find that we can fit the data almost exactly over the full frequency range. Moreover, the posterior uncertainty is negligible because the data contains ample information about the model parameters, allowing us to precisely infer the parameter values.

Having inferred the flame response from the microphone measurements, we now seek to validate this estimate. We do so in three ways: (i) by comparing the inferred FTF to direct measurements, (ii) by evaluating the time delays extracted from the data to ensure they are physically reasonable, and (iii) by examining the low-frequency limit of the inferred FTF.

In previous work on simpler rigs, we validated the inference framework by comparing the inferred flame response to direct chemiluminescence measurements, and found them to be almost identical (Yoko and Juniper, 2024b). No such independent measurement is available for the SCARLET rig, so we must validate our inferred FTFs indirectly. We note that, in the absence of independent flame diagnostics, the validity of the inferred FTF is conditional on the correctness of the acoustic model. We have taken care to validate the acoustic model thoroughly in the preceding sections, but this limitation should be borne in mind when interpreting the results shown in the following.

The only experimentally available FTF is obtained by applying the two-source method to the microphone measurements. However, as shown by Eder *et al.* (2025), the two-source method introduces systematic artefacts when applied to complex combustor geometries such as SCARLET. Our inferred FTF does not contain these artefacts, so it cannot be compared directly to the experimental FTF.

To bridge this gap, we use the calibrated model as follows. With the inferred FTF fixed in the model, we compute

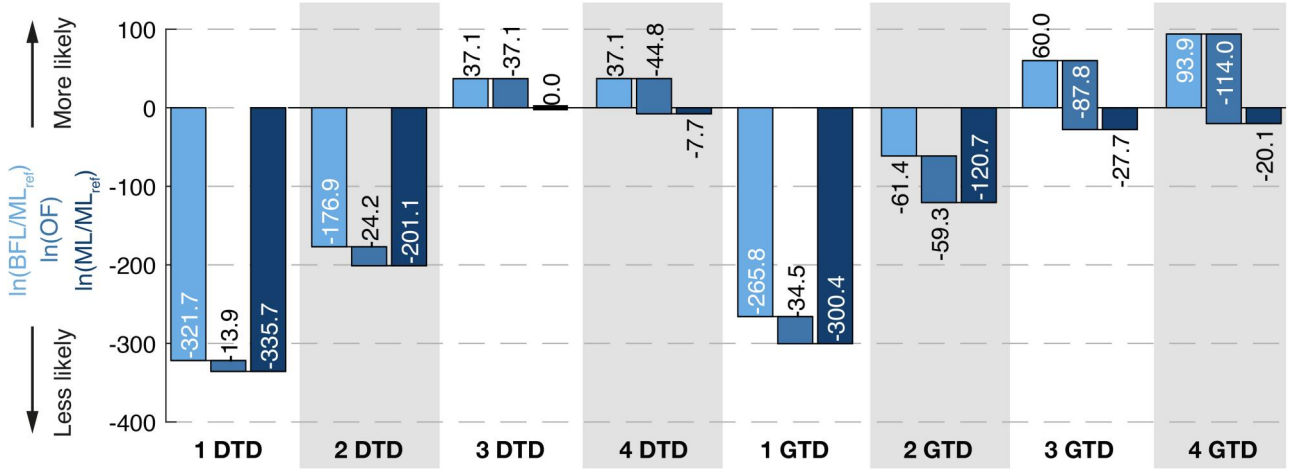

FIG. 6. Same as Fig. 3 but for ranking the eight flame models: DTD with 1–4 delays and GTD with 1–4 delays. We see that a flame model with three discrete time delays has the most evidence, given the data.

https://doi.org/10.1121/10.0043969

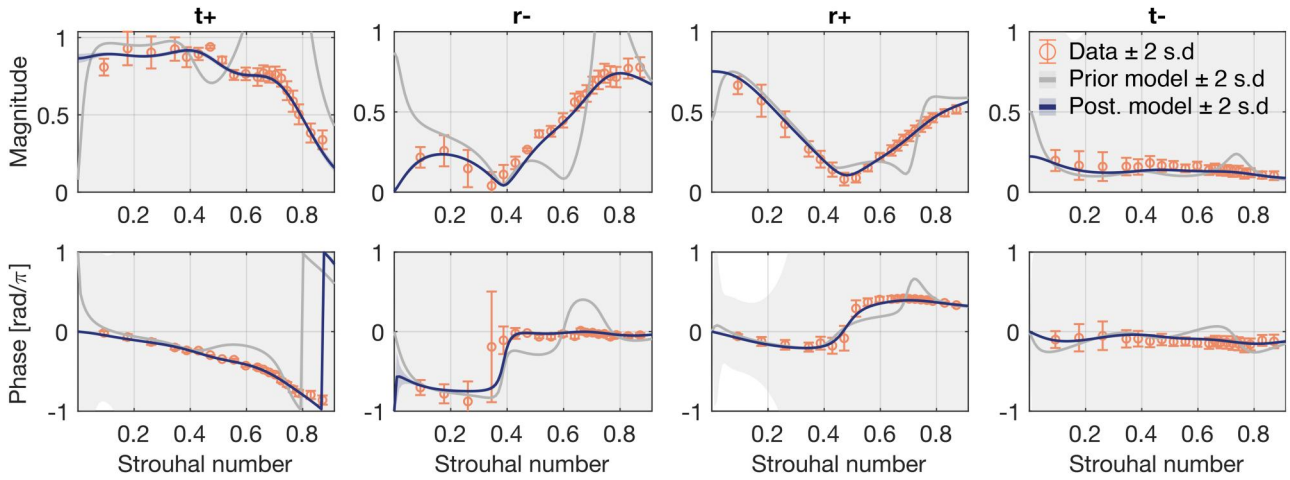

FIG. 7. The same as Fig. 4, but for the reacting results.

the acoustic pressures at the microphone locations and then apply the two-source method (2SM) to these synthetic measurements. This produces a "model 2SM" FTF that is distorted by the same artefacts as the experimental FTF, and can therefore be compared to it on equal footing. The comparison is shown in Fig. 8, where three FTFs are plotted: (i) our inferred FTF, (ii) the model 2SM FTF, and (iii) the experimental 2SM FTF. The close agreement between the two 2SM FTFs confirms that our model reproduces the experimental data precisely. Moreover, because the two-source method distorts both the model and experimental FTFs in the same way, the agreement implies that the true FTF in the experiment is close to our inferred FTF, provided that the acoustic model is a faithful representation of the rig (a condition supported by the non-reacting validation in Figs. 4 and 5).

As the second validation step, we assess the three $n-\tau$ pairs that we learned from the data, shown in Table III. The time delays have been nondimensionalized by the injector bulk velocity and the distance from the injector inlet to the centre of heat release of the flame.

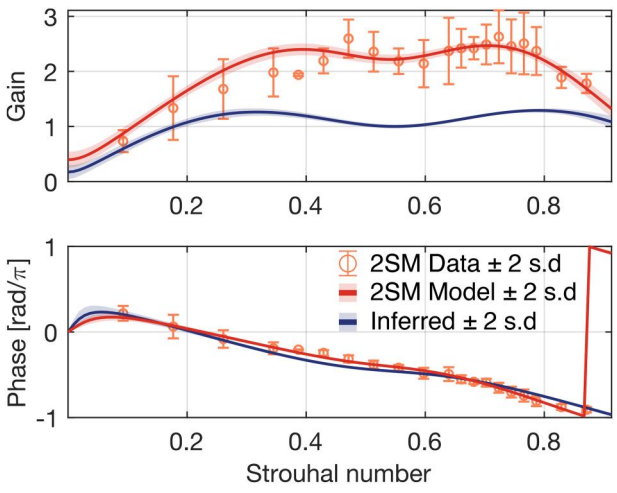

FIG. 8. FTFs obtained by applying the two-source method (2SM) to the data, shown as orange error bars, and to the posterior model, shown as a red line and patch. The blue line and patch show the FTF we learn by assimilating the data into the physics-based model.

We first note that all three time delays are positive, which is required for the model to be physically valid. This is an important check because a systematic error in the model could drive the parameters to non-physical values and, in some cases, still achieve a reasonable fit to the data.

The first time delay is consistent with the convective time scale between the injector dump plane and the flame, and is accompanied by a positive interaction index. Both of these observations suggest that the first $n-\tau$ pair captures the effect of fluctuating turbulence intensity generated at the sudden expansion, which is then convected to the flame front and modulates the turbulent flame speed (Schuermans *et al.*, 2004).

The second time delay is consistent with the distance between several structures in the injector (swirl vanes and struts) and the flame, and is accompanied by a negative interaction index. This suggests that the second $n-\tau$ pair is related to equivalence ratio fluctuations driven by vorticity shed off of these structures. Detailed simulations on premixed swirling flames have shown that equivalence ratio fluctuations produce two peaks in the impulse response with alternating signs (Huber and Polifke, 2009), yet we do not find an accompanying positive peak. This will be explained shortly.

The third time delay is too long to match any convective scales in the rig, suggesting that this delay corresponds to circulation fluctuations. Similarly to the equivalence ratio fluctuations, Karmarkar *et al.* (2022) used simulations to show that a swirl-stabilized flame responds to tangential perturbations with a positive and negative peak of equal magnitude, but we only see a negative peak in this work.

For both the equivalence ratio and circulation fluctuations, we note that we can only learn the superimposed effects of these time delays from the microphone measurements. Because of this, if any of the peaks have similar time delays, these will be lumped together in a single time delay with their interaction indices summed. The summation and

TABLE III. Summary of the $n-\tau$ pairs learned for the three discrete time delay models. The time delays have been nondimensionalized by the injector bulk velocity (u_i) and the distance from the injector inlet to the flame (L_{i-f}): $\hat{\tau} = \tau u_i / L_{i-f}$.

	Delay 1		Delay 2		Delay 3	
n	0.88	±0.03	−0.41	±0.04	−0.30	±0.03
$\hat{\tau}$	0.33	±0.0064	1.02	±0.02	1.73	±0.02

partial cancellation of these overlapping delays is probably why we only find three unique time delays from this data.

Finally, we assess the low-frequency limit, which is commonly used as a test of the validity of an FTF. Polifke and Lawn (2007) demonstrate that the low-frequency limit depends on how the fuel is introduced into the combustor, and discuss two limiting cases. If the fuel injection is sensitive to acoustic perturbations, the acoustic fluctuations drive a fluctuation in the fuel mass flow rate, resulting in a low-frequency limit gain of one and a phase of zero (Polifke and Lawn, 2007). However, if the fuel injection is insensitive to acoustic perturbations, the fuel mass flow rate is fixed, and the acoustic velocity perturbations drive equivalence ratio fluctuations, resulting in a low-frequency limit gain of zero (Polifke and Lawn, 2007). In this case, the phase is meaningless, but tends to $\pm\pi/2$ in the limit.

The injector fitted to SCARLET for this study introduces the fuel via several air-blast atomizers, which inject liquid fuel at relatively high pressure. Therefore, it is expected that this injector would more closely resemble the second case, so we expect to find a low-frequency limit close to zero. Reassuringly, we see in Fig. 8 that the FTF we learn from the data has a low-frequency limit that approaches zero, with error bars that almost include zero. This insight gives us the option to add the constraint that the low-frequency limit must be zero, reducing the number of parameters in our flame model from six to five. While this is not necessary for learning the parameters of the flame model, and has negligible impact on the FTF, it reduces the amount of data required to learn a model that maps between operating conditions and flame response (Yoko *et al.*, 2026).

3. Identifiability of the parameters

The small uncertainty in the inferred parameters, shown in Table III, indicates that the data contains ample information about the model parameters, but does not necessarily guarantee that the six parameters are uniquely identifiable from the microphone measurements. We demonstrate identifiability in two ways: by inspecting the pairwise joint posterior distributions for evidence of highly correlated parameters, and by inspecting the eigenvalues of the information-gain matrix to identify any directions in parameter space that are poorly informed by the data (MacKay, 1999).

The pairwise joint posterior distributions of the six flame model parameters are shown in Fig. 9. While some parameters are moderately correlated, the correlations are not sufficient to prevent unique identification, which would manifest as the ellipsoids collapsing towards a line.

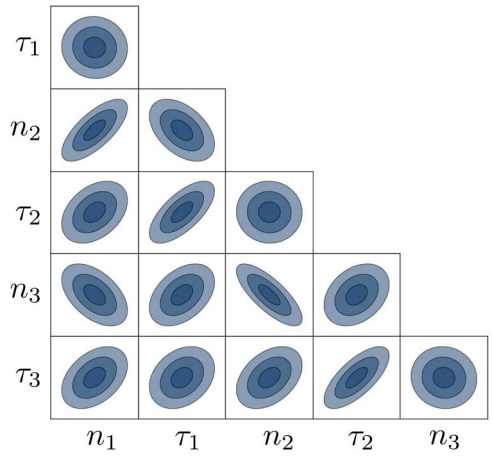

FIG. 9. Pair-wise joint posterior distributions of the six flame model parameters. Each panel shows the joint distribution between a pair of parameters, revealing the correlation between the parameters. The rings denote 1–3 standard deviations around the MAP estimate.

We define the information-gain matrix $\boldsymbol{A} = C_{a|D} C_a^{-1}$, whose eigenvalues lie between zero and one, where a value close to zero indicates a parameter direction that is well-informed by the data (MacKay, 1999). The largest eigenvalue of $\boldsymbol{A}$ is $\mathcal{O}(10^{-4})$, confirming that all parameter directions are well-informed by the data and that the parameters are therefore uniquely identifiable.

4. Impact of errors in the acoustic parameters

When we discussed the results of inferring the model parameters of the non-reacting rig in Fig. 4, we noted that the prior model performs surprisingly well. Based on this, it may seem as if it would be safe to neglect the inference step for the non-reacting rig and move directly to inferring flame model parameters. To demonstrate the importance of carefully inferring the acoustic parameters from non-reacting data, we repeat the process of inferring the flame model parameters while fixing the acoustic parameters to their prior values (i.e., those that produced the gray line in Fig. 4).

Figure 10 shows the resulting two-source method FTFs from the model and data. We see that a seemingly small systematic error in the non-reacting model results in a significant error in the inferred flame response. This is because the flame model parameters are being used to compensate for the errors in the acoustic parameters. The magnitude of this error is largely driven by the poor modelling of the reflection of the upstream travelling wave (r+ in Fig. 4). The over-prediction of this wave reflection in the non-reacting case will be worsened by the impedance mismatch introduced by the temperature discontinuity in the reacting case, amplifying the systematic error and driving the FTF to extreme values in an attempt to match the microphone measurements. We also note that we have not incorporated this systematic error in our uncertainty quantification, so while

https://doi.org/10.1121/10.0043969

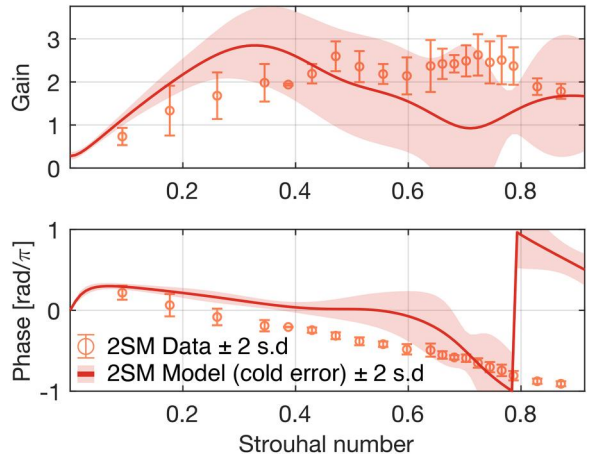

FIG. 10. Demonstration of the impact of a small error in the non-reacting model. We plot the FTF that we learn when neglecting the calibration of the non-reacting model. We fix the acoustic parameters to their prior values (gray line in Fig. 4) and learn the FTF from reacting data. We once again process the data and model using the two-source method to produce comparable FTFs. The data is shown as orange error bars, and the model is shown as a red line and patch.

the error bars on the model prediction are large, they are still underestimated. This result demonstrates the importance of carefully and rigorously constructing the model and inferring the acoustic model parameters from non-reacting data before attempting to learn the flame model parameters from reacting data.

VI. COMPUTATIONAL COST AND INDUSTRIAL APPLICATION

The computational cost of this method is driven by the cost of evaluating the likelihood, which requires solving the network model and its adjoint at each frequency and forcing direction for which we have data. This scales with (i) the size of the network, (ii) the amount of data, and (iii) the number of parameters that we are learning. Additionally, if any of the parameters influence the mean flow, then the adjoint of the nonlinear mean flow solver must be evaluated at each iteration, which increases the cost.

For the SCARLET rig, the network model comprises 32 ducts connected by 19 jump conditions and three boundary conditions. We have data at 20 frequencies from upstream and downstream forcing, requiring 40 forward and adjoint model evaluations to evaluate the likelihood and its gradient. For the non-reacting case, we learn six parameters, each of which affects the mean flow. It takes 50 s on a single core of a 2 GHz laptop CPU to find the posterior parameter distribution and calculate the model ranking metrics. For the reacting case, we also learn six parameters, but none affect the mean flow, reducing the cost to 8 s. These costs are negligible compared with the time required to collect the data, meaning that the method can be run online alongside experimentation.

In a typical workflow, the cold model is calibrated once for a given fuel injector and chamber geometry, at a computational cost of $\mathcal{O}(1\,\text{min})$. Following this, the flame model can be learned for several operating conditions, at a computational cost of $\mathcal{O}(10\,\text{s})$ per condition. This provides data from which the mapping between operating condition and flame model parameters can be built, such that the FTF for a specific injector can be estimated at untested conditions (Yoko *et al.*, 2026).

The calibrated model itself has several uses. As demonstrated in this paper, it allows an unscaled FTF to be extracted from microphone measurements in a rig where the two-source method is not valid. If the rig replicates engine geometry and operating conditions, the calibrated acoustic parameters should transfer directly to an engine model, as supported by the insensitivity to local flow conditions demonstrated in the SCARLET rig. Similarly, as a first estimate, the calibrated flame model can also be ported to an engine model, though this would neglect changes to the FTF from flame-flame interaction and different confinement when moving from a single-can to an annular chamber. Such an engine model can provide early indications of instability risk, which can be continually updated as more data becomes available from higher-TRL rigs or engine operation. While the methodology in this paper demonstrated assimilating data from forced experiments, the method is also capable of assimilating data from unforced experiments, or even a combination of both.

VII. CONCLUSION

In this paper, we demonstrate the application of Bayesian model selection and parameter inference to a model of a complex industrial thermoacoustic test rig. By rigorously declaring our prior state of uncertainty, then using data to update our knowledge about the uncertain aspects of our model, we develop a low-order model that fits the data under both non-reacting and reacting conditions. In doing so, we simultaneously infer quantities that we want to know but could not measure directly, most importantly, the flame response.

We use Bayesian model comparison to rank several plausible models for (i) the swirl injector, (ii) the perforated liner, and (iii) the flame. We find that the swirl injector is best modelled as an area decrease and increase with full pressure recovery, and a fluctuating drag force which must be included as a source in both the momentum and energy equations. This is an adaptation of the popular $L-\zeta$ model, with L set to zero, as has previously been applied when modelling the SCARLET rig (Eder *et al.*, 2025). The perforated liner is best modelled using Howe's Rayleigh conductivity model (Howe, 1979) with a correction factor for finite liner thickness (Jing and Sun, 1999). Finally, we find that, for this flame, the best flame model is a discrete time delay model with just three $n-\tau$ pairs. Following this careful process of ranking and selecting the best models for each component of the rig, we show that we can create a low-order model of the rig that fits the data with high precision over the full frequency range that was tested, for both the non-reacting and reacting cases.

By studying the inferred FTF, we confirm the findings of a recent paper, which demonstrated that the two-source method introduces artefacts in the FTF when it is applied to complex combustor geometries, such as the SCARLET rig. However, we show that our inference framework overcomes this problem and improves the estimate of the true FTF. Moreover, because we assimilate data into a physics-based model of the rig, the calibrated model can interpolate accurately and extrapolate with known uncertainty bounds.

Finally, we demonstrate that the $n-\tau$ pairs that we learn are physically meaningful and can be correlated to physical phenomena inside the rig, suggesting that it might be possible to build a physics-based model that predicts the FTF at other operating conditions at which the dominant physics remains the same. This is the topic of ongoing work. However, in a companion paper to this, we demonstrate how we can make immediate progress by training a small probabilistic machine learning model to learn the mapping between operating conditions and flame model parameters, the cost of which can be minimized by using Bayesian optimal experiment design (Yoko *et al.*, 2026).

In future work, we will extend this method to learning the flame response from a full-annular test rig. These rigs typically lack excitation systems, so we will learn the thermoacoustic growth rate and frequency from the combustion noise (Noiray and Denisov, 2017). We can then assimilate this data into a model of the combustor to learn the flame response at the dominant eigenfrequency. Without forcing the rig, we will not be able to learn a frequency-resolved FTF. However, by repeating this at several operating conditions, we can produce an operating map describing how the flame response changes with engine operating conditions.

ACKNOWLEDGMENTS

This work is funded by the European Union, under Grant Agreement No. 101102004 within the HEAVEN project. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or Clean Aviation Joint Undertaking. Neither the European Union nor Clean Aviation JU can be held responsible for them. The test rig was funded with financial support by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) as part of the AG-Turbo program "COORETEC–AG Turbo 2020," which is acknowledged as per resolution of the German Federal Parliament under Grant No. 03ET2012B.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data that support the findings of this study are available from Rolls-Royce. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors upon reasonable request and with the permission of Rolls-Royce.

APPENDIX A: NETWORK MODEL BOUNDARY AND JUMP CONDITIONS

1. Solution method

We decompose each flow variable into steady and fluctuating components: $\phi = \overline{\phi} + \phi'$. The fluctuating component can be expressed in terms of three waves: forward (f) and backward (g) travelling acoustic waves, and a convected entropy wave (e), the strengths of which must be solved for. The travelling waves are used to construct the three primary acoustic variables: pressure (p'), velocity (u'), and entropy (ζ'), as well as, where convenient, the acoustic density perturbations (ρ'),

$$p' = f + g, \tag{A1}$$

$$u' = \frac{f-g}{\overline{\rho}\overline{c}}, \tag{A2}$$

$$\zeta' = e, \tag{A3}$$

$$\rho' = \frac{f+g}{\overline{c}^2} - e\frac{\overline{\rho}}{c_p}, \tag{A4}$$

where $\overline{c}$ is the sound speed and c_p is the specific heat.

Each jump and boundary condition provides a set of equations for the steady and acoustic flows in the network. For the steady flow, the resulting system of nonlinear equations is linearized and solved using Newton iteration. For the acoustic flow in the forced case, the resulting system of equations is linear and can be solved directly. In the unforced case, the resulting system produces a nonlinear eigenvalue problem, which is solved using Newton iteration on the eigenvalue.

Unless otherwise specified, we present the governing equations in terms of the total quantities (i.e., mean + acoustic) for conciseness. The equations can then be linearized to extract the mean and acoustic equations. Additionally, to keep the equations compact, we present them in terms of the conserved quantities: mass flow rate ($\dot{m}$), momentum (M), enthalpy (H), and energy (E), rather than the primary acoustic variables. The relations we use between the conserved quantities and the primary acoustic variables are as follows:

$$\dot{m}' = A(\rho'\overline{u} + \overline{\rho}u'), \tag{A5}$$

$$M' = \dot{m}'\overline{u} + \overline{\dot{m}}u' = A(\rho'\overline{u}^2 + 2\overline{\rho u}u'), \tag{A6}$$

$$H' = c_p\overline{T}\left(\frac{p'}{\overline{p}} - \frac{\rho'}{\overline{\rho}}\right) + \overline{u}u', \tag{A7}$$

$$E' = \dot{m}'\overline{H} + \overline{\dot{m}}H'. \tag{A8}$$

FIG. 11. Upstream boundary.

https://doi.org/10.1121/10.0043969

FIG. 12. Downstream boundary.

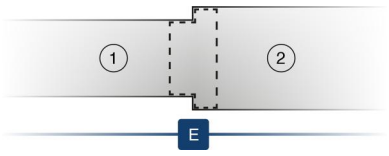

FIG. 14. Sudden expansion.

We solve for two additional variables in the mean flow equations, which are the mixture fraction (Z) and a binary progress variable (C). We assume instantaneous and complete combustion, so we do not admit perturbations of the progress variable ($C' = 0$ everywhere). While we could solve for the perturbations of mixture fraction in the acoustic equations, they have a negligible influence on all jump conditions except for the flame. Rather than solving for them globally, we introduce their effects locally through the FTF in the flame jump condition.

Through the mixture fraction and progress variable, we compute the local gas composition considering six component species: Ar, CO_2, H_2O, N_2, O_2, and $C_{11}H_{22}$, the last of which is used as a surrogate for jet-A1. Using thermodynamic data from *HyChem* (Wang *et al.*, 2018; Xu *et al.*, 2018), we compute the local thermodynamic properties of the gas in each duct, as well as the steady heat release rate at the flame using a one-step global reaction mechanism.

2. Boundary conditions

We provide two options for boundary conditions: one for upstream terminations and the other for downstream terminations.

Upstream boundary: At an upstream boundary (Fig. 11), we solve two equations: one to impose zero incoming entropy perturbations, and the other to set the appropriate acoustic boundary condition, depending on whether the termination is unforced or forced.

Under unforced conditions, we solve

$$\zeta_1' = 0, \tag{A9a}$$

$$f_1 = R_u g_1. \tag{A9b}$$

Under forced conditions, we solve

$$\zeta_1' = 0, \tag{A9c}$$

$$f_1 = 1, \tag{A9d}$$

where R_u is the upstream acoustic reflection coefficient.

For the mean flow, we specify the mass flow rate, temperature and mixture fraction at the inlet.

Downstream boundary: At a downstream boundary (Fig. 12), we solve one equation to apply the appropriate boundary condition depending on whether the termination is unforced or forced.

Under unforced conditions, we solve

$$g_1 - R_d f_1 + R_e \zeta_1' = 0, \tag{A10a}$$

Under forced conditions, we solve

$$G_1 e^{k_g L} = 1, \tag{A10b}$$

where R_d is the downstream acoustic reflection coefficient, and R_e is the coefficient governing the conversion of entropy fluctuations into acoustic waves.

For the mean flow, we specify the pressure at the outlet.

3. Jump conditions

In general, at a discontinuity, we solve the integral conservation laws of mass, momentum, and energy, assuming 1D, inviscid flow. We include volumetric source terms to account for additional mass injection, force, heat or mechanical work (Chu, 1965; Juniper, 2018),

$$\frac{\mathrm{D}\rho}{\mathrm{D}t} = -\rho \nabla \cdot u + \epsilon \delta \dot{m}, \tag{A11a}$$

$$\rho \frac{\mathrm{D}u}{\mathrm{D}t} = -\nabla p + \epsilon \delta f, \tag{A11b}$$

$$\rho T \frac{\mathrm{D}\zeta}{\mathrm{D}t} = \dot{q} + \epsilon \delta (\dot{q} + \dot{w}), \tag{A11c}$$

where $\epsilon\delta\phi$ denotes volumetric source terms in mass flow ($\dot{m}$), force (f), heat release ($\dot{q}$), and work ($\dot{w}$).

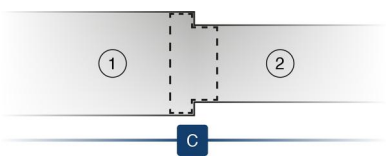

FIG. 13. Sudden contraction.

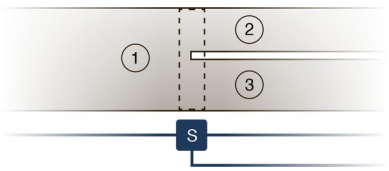

FIG. 15. Branch split.

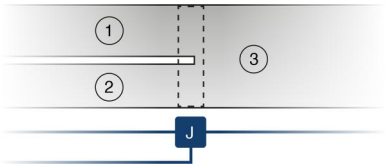

FIG. 16. Branch junction.

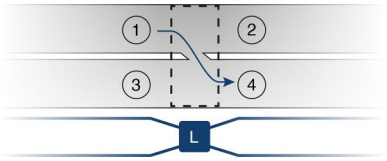

FIG. 18. Perforated liner.

For isentropic flow elements, we replace momentum and energy conservation with conservation of entropy and stagnation enthalpy (Dowling and Stow, 2003).

Sudden contraction: We assume that the flow responds to a sudden contraction (Fig. 13) isentropically,

$$\dot{m}_1 = \dot{m}_2, \tag{A12a}$$

$$\zeta_1 = \zeta_2, \tag{A12b}$$

$$H_1 = H_2. \tag{A12c}$$

Additionally, in the mean flow equations, the mixture fraction and progress variable are conserved

$$\overline{Z}_1 = \overline{Z}_2, \tag{A12d}$$

$$\overline{C}_1 = \overline{C}_2. \tag{A12e}$$

Sudden expansion: At a sudden expansion (Fig. 14), we solve conservation of mass, momentum and energy. We introduce a pressure recovery factor, R, which is used to calculate the axial pressure-force on the walls,

$$\dot{m}_1 = \dot{m}_2, \tag{A13a}$$

$$M_1 = M_2 + A_2(p_2 - p_1) + \frac{1}{2}(A_2 - A_1)\rho_1 R u_1^2, \tag{A13b}$$

$$E_1 = E_2. \tag{A13c}$$

Additionally, in the mean flow equations, the mixture fraction and progress variable are conserved,

$$\overline{Z}_1 = \overline{Z}_2, \tag{A13d}$$

$$\overline{C}_1 = \overline{C}_2. \tag{A13e}$$

Branch split: When one duct splits into several (Fig. 15), we have a jump condition with one inlet and two (or more) outlets. A branch split is almost always accompanied by an area decrease, so we assume that the flow is isentropic,

$$\dot{m}_1 = \dot{m}_2 + \dot{m}_3, \tag{A14a}$$

$$\zeta_1 = \zeta_2 = \zeta_3, \tag{A14b}$$

$$H_1 = H_2 = H_3. \tag{A14c}$$

Additionally, in the mean flow equations, the mixture fraction and progress variable are conserved,

$$\overline{Z}_1 = \overline{Z}_2 = \overline{Z}_3, \tag{A14d}$$

$$\overline{C}_1 = \overline{C}_3 = \overline{C}_3. \tag{A14e}$$

Branch junction: When many ducts join into one (Fig. 16), we have a jump condition with two (or more) inlets and one outlet. To close the equations, we assume that there is pressure continuity among the inlet ducts at the junction point,

$$p_1 = p_2, \tag{A15a}$$

$$\dot{m}_1 + \dot{m}_2 = \dot{m}_3, \tag{A15b}$$

$$M_1 + M_2 = M_3 + A_3(p_3 - p_1), \tag{A15c}$$

$$E_1 + E_2 = E_3. \tag{A15d}$$

Additionally, in the mean flow equations, we apply conservation of fuel mass flow rate to calculate the downstream mixture fraction, and apply a Boolean OR operation to propagate the progress variable,

$$\dot{m}_1\overline{Z}_1 + \dot{m}_2\overline{Z}_2 = \dot{m}_3\overline{Z}_3, \tag{A15e}$$

$$\overline{C}_1 \quad \overline{C}_2 = \overline{C}_3. \tag{A15f}$$

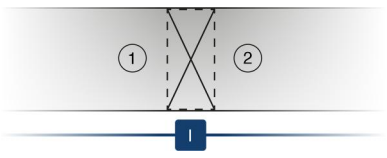

FIG. 17. Swirl injector.

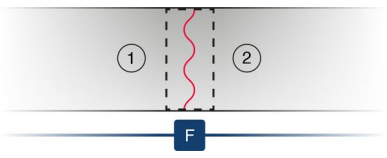

FIG. 19. Flame.

https://doi.org/10.1121/10.0043969

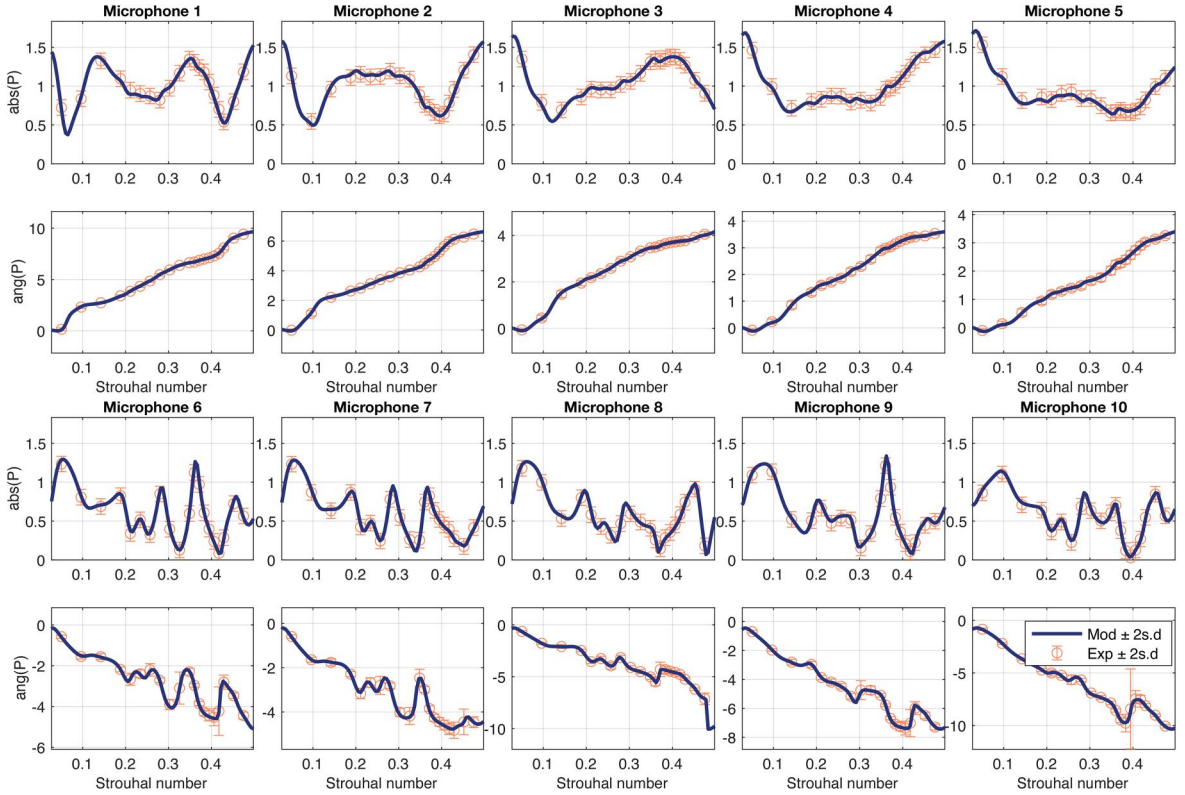

FIG. 20. Measured vs modelled pressure at each of the microphones. Upstream forcing in the non-reacting rig.

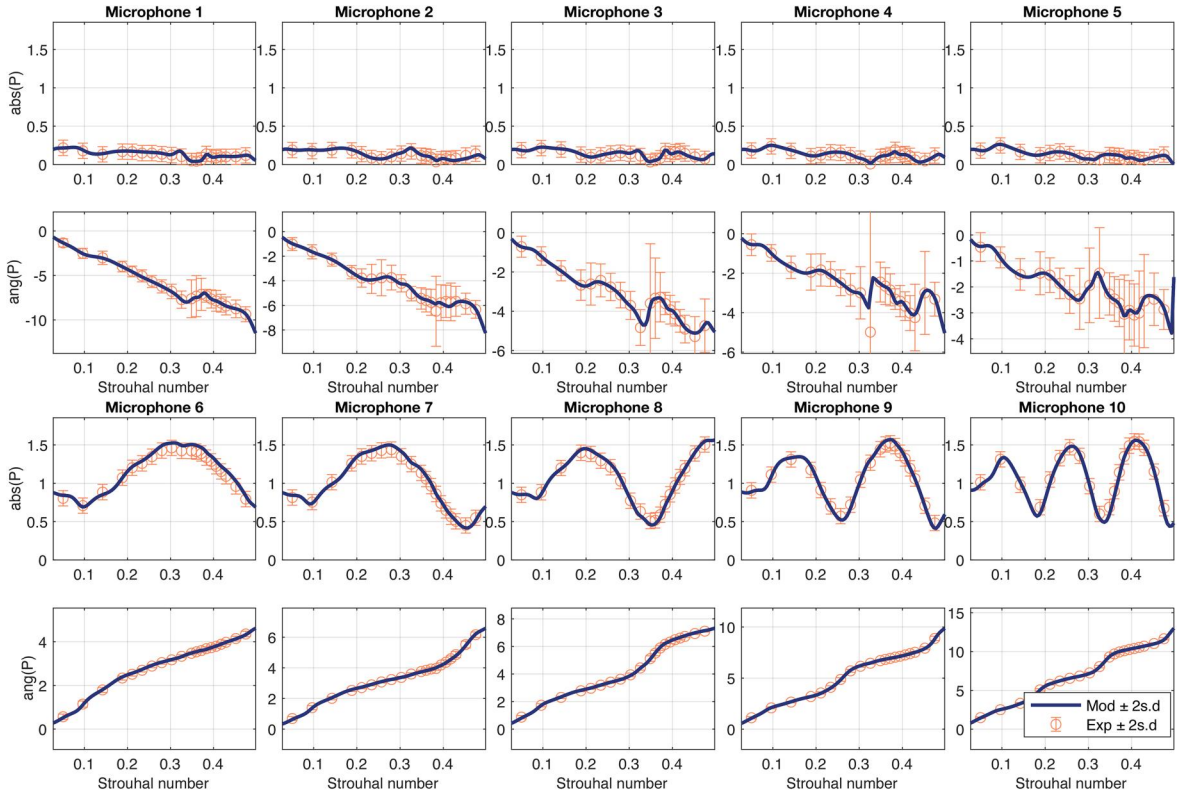

FIG. 21. Measured vs modelled pressure at each of the microphones. Downstream forcing in the non-reacting rig.

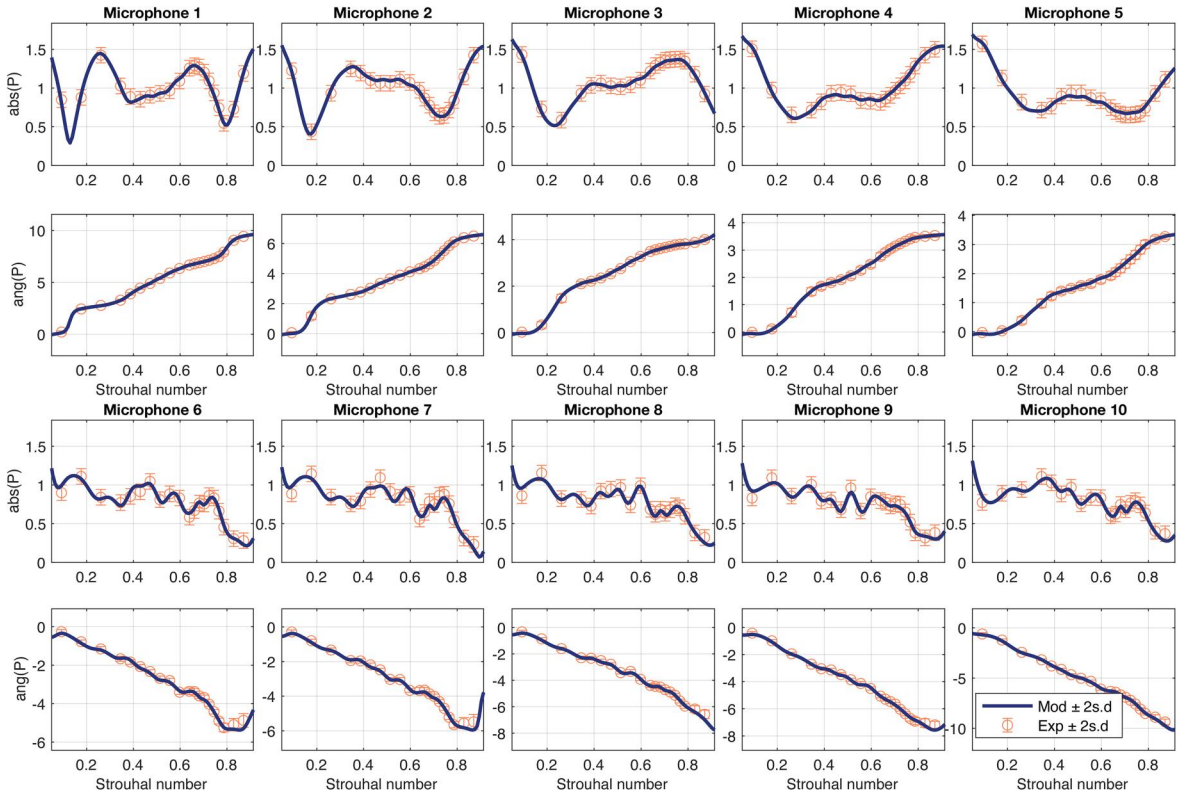

FIG. 22. Measured vs modelled pressure at each of the microphones. Upstream forcing in the reacting rig.

Swirl injector: A swirl injector (Fig. 17) is modelled as a localized momentum and energy jump between two area changes, which are added as separate jump conditions. The injector jump condition allows us to (i) account for drag on the swirl vanes, (ii) account for inertia of the fluid slug in the injector, and (iii) introduce a source term for the injected fuel ($\dot{m}_f$),

$$\dot{m}_1 + \dot{m}_f = \dot{m}_2, \tag{A16a}$$

$$M_1 - \rho_1 A_1 \left(\frac{1}{2} C_d u_1^2 + L_{\text{eff}} \frac{\mathrm{d}u_1}{\mathrm{d}t} \right) = M_2 + A_2 (p_2 - p_1), \tag{A16b}$$

$$E_1 - \frac{1}{2} \rho_1 A_1 C_d u^3 = E_2. \tag{A16c}$$

Additionally, in the mean flow equations, we apply a mixing law to calculate the downstream mixture fraction, and the progress variable is conserved,

$$\frac{\dot{m}_1 \overline{Z}_1 + \dot{m}_f}{\dot{m}_1 + \dot{m}_f} = \overline{Z}_2, \tag{A16d}$$

$$\overline{C}_1 = \overline{C}_2. \tag{A16e}$$

Perforated liner: A perforated liner (Fig. 18) typically links two parallel or concentric ducts, giving two inlet ducts and two outlet ducts. We assume that the extraction of mass from the suction side (1–2) is an isentropic process, while the injection of mass on the blowing side (3–4) is not.

The liner model is based on Howe's Rayleigh conductivity model, which relates the perturbed mass flow rate through a cooling hole to the difference in acoustic pressure across the liner panel,

$$\dot{m}'_h = \frac{K_r}{s} (p'_2 - p'_1), \tag{A17}$$

where K_r is the Rayleigh conductivity and s is the complex frequency. The Rayleigh conductivity for an infinitely thin panel is

$$K_{r0} = \frac{2R_h \left(\frac{\pi}{2} I_1(\mathrm{St}) e^{-\mathrm{St}} + i K_1(\mathrm{St}) \sinh(\mathrm{St}) \right)}{\mathrm{St} \frac{\pi}{2} I_1(\mathrm{St}) \sinh(\mathrm{St}) - i \mathrm{St} K_1(\mathrm{St}) \cosh(\mathrm{St})}, \tag{A18}$$

where I_1 and K_1 are modified Bessel functions of the first and second kind, $\mathrm{St} = -2is(R_h/u_h)$ is the Strouhal number based on the radius of the hole, R_h, and the vortex convection velocity, which Howe approximates as half the jet velocity, u_h.

Finally, Jing and Sun (1999) propose a thickness correction to Howe's model, giving the corrected Rayleigh conductivity,

https://doi.org/10.1121/10.0043969

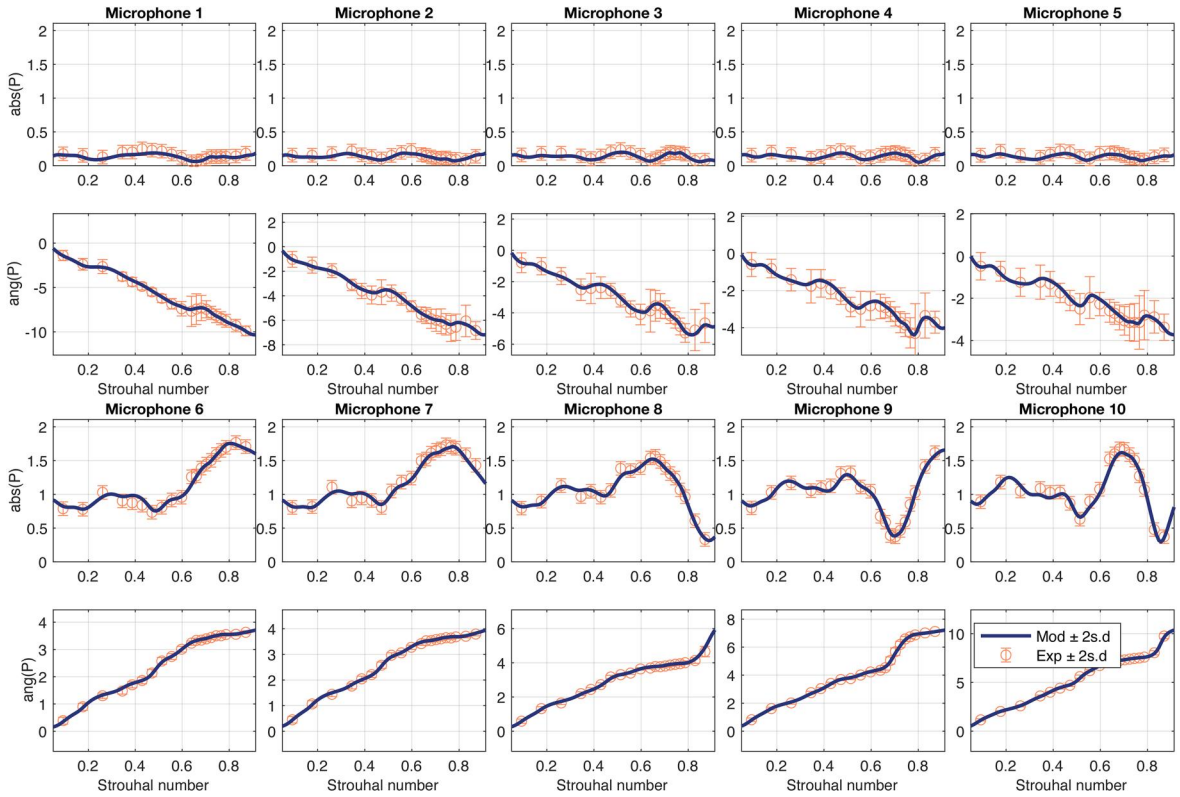

FIG. 23. Measured vs modelled pressure at each of the microphones. Downstream forcing in the reacting rig.

$$K_r = K_{r0}\left(1 + \frac{K_{r0}L_h}{\pi R_h^2}\right)^{-1}. \tag{A19}$$

With an expression for the fluctuating mass flow, we can form the jump condition. On the suction side, the acoustic equations are

$$\dot{m}_1 - \dot{m}_h = \dot{m}_2, \tag{A20a}$$

$$\zeta_1 = \zeta_2, \tag{A20b}$$

$$H_1 = H_2. \tag{A20c}$$

Additionally, in the mean flow equations, the mixture fraction and progress variable are conserved,

$$\overline{Z}_1 = \overline{Z}_2, \tag{A20d}$$

$$\overline{C}_1 = \overline{C}_2. \tag{A20e}$$

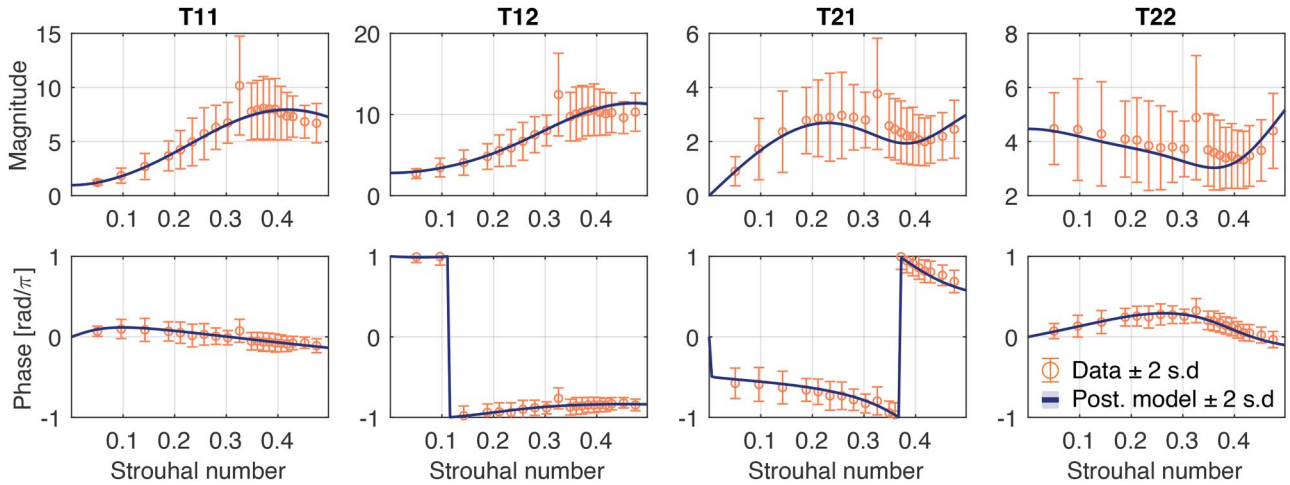

FIG. 24. Measured vs modelled transfer matrices for the non-reacting rig. The data is shown as orange error bars, and the model is shown as a blue line and patch.

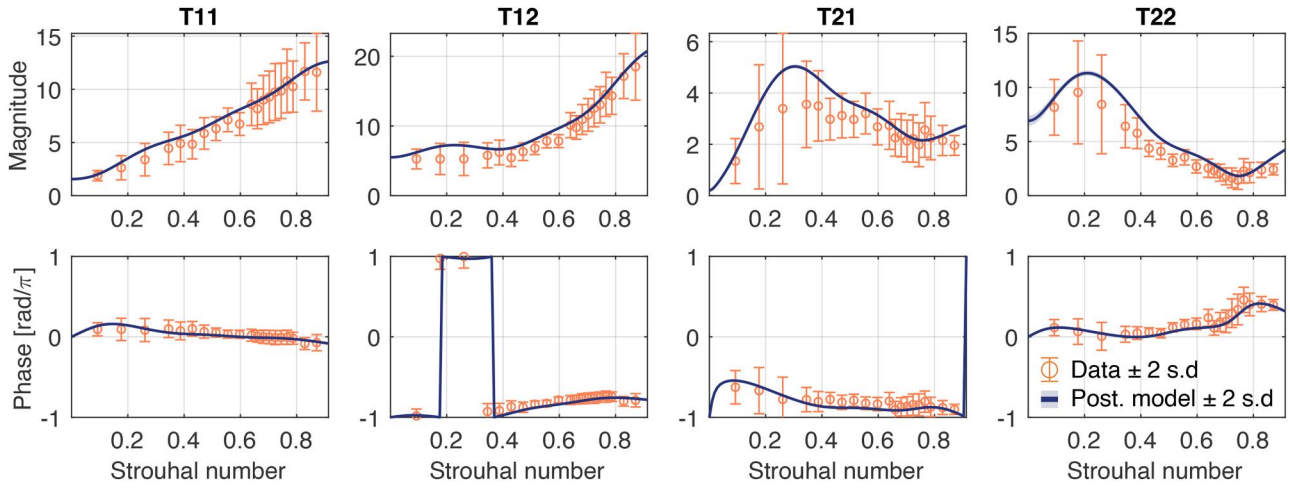

FIG. 25. Same as Fig. 24 but for the reacting rig.

On the blowing side, the acoustic equations are

$$\dot{m}_3 + \dot{m}_h = \dot{m}_4, \tag{A20f}$$

$$M_3 + \dot{m}_h u_h \cos(\theta_h) = M_4 + A_4(p_4 - p_3), \tag{A20g}$$

$$E_3 = E_4. \tag{A20h}$$

Additionally, in the mean flow equations, we apply a mixing law to account for the dilution of the mixture fraction from the air entering through the liner, and the progress variable is conserved,

$$\frac{\dot{m}_3 \overline{Z}_3}{\dot{m}_3 + \dot{m}_h} = \overline{Z}_4, \tag{A20i}$$

$$\overline{C}_3 = \overline{C}_4. \tag{A20j}$$

Flame: At the flame (Fig. 19), we introduce a volumetric heat release rate source in the energy equation, which is related to the FTF, $\mathcal{F}$,

$$\dot{m}_1 = \dot{m}_2, \tag{A21a}$$

$$M_1 = M_2 + A_2(p_2 - p_1), \tag{A21b}$$

$$E_1 + \overline{\dot{q}}_f \frac{u_1'}{\overline{u}_1} \mathcal{F} = E_2. \tag{A21c}$$

Additionally, in the mean flow equations, the mixture fraction is conserved, and the progress variable is set to one,

$$\overline{Z}_1 = \overline{Z}_2, \tag{A21d}$$

$$\overline{C}_2 = 1. \tag{A21e}$$

We can then propose any number of models for $\mathcal{F}$, for example, discrete distributed time delays $\mathcal{F} = \sum_{i=1}^{N} n_i e^{(-s\tau_i)}$, or Gaussian distributed time delays: $\mathcal{F} = \sum_{i=1}^{N} n_i e^{[-s\tau_i + (1/2)s^2\sigma_i^2]}$.

APPENDIX B: MODELLED VS MEASURED ACOUSTIC PRESSURE

Figures 20–23 compare the measured and modelled acoustic pressure at the ten microphone locations. The acoustic pressure is presented as the magnitude and phase of the Fourier decomposed signal.

APPENDIX C: MODELLED VS MEASURED TRANSFER MATRICES

In Figs. 24 and 25, we present the combustor transfer matrices obtained from the measurements and from the model, for both the non-reacting and reacting cases. We compute the error bars on the measured transfer matrices by propagating the uncertainty in the microphone measurements through the transfer matrix computation. We see that moderate uncertainty in the microphone measurements (see plots in Appendix B) drives large uncertainty in the transfer matrix, which is a consequence of the poor conditioning of the transfer matrix computation. By comparison, the scattering matrices in Figs. 4 and 7 maintain moderate uncertainty, which is why they are preferred for this paper.

Æsøy, E., Indlekofer, T., Gant, F., Cuquel, A., Bothien, M. R., and Dawson, J. R. (**2022**). "The effect of hydrogen enrichment, flame-flame interaction, confinement, and asymmetry on the acoustic response of a model can combustor," Combust. Flame **242**, 112176–112111.

Alemela, P. R., Fanaca, D., Hirsch, C., Sattelmayer, T., and Schuermans, B. (**2010**). "Determination and scaling of thermo acoustic characteristics of premixed flames," Int. J. Spray Combust. Dyn. **2**(2), 169–198.

Bauer, A. B. (**1977**). "Impedance theory and measurements on porous acoustic liners," J. Aircraft **14**(8), 720–728.

Birbaud, A. L., Durox, D., and Candel, S. (**2006**). "Upstream flow dynamics of a laminar premixed conical flame submitted to acoustic modulations," Combust. Flame **146**(3), 541–552.

Chu, B. T. (**1965**). "On the energy transfer to small disturbances in fluid flow (Part I)," Acta Mech. **1**(3), 215–234.

Cuquel, A., Durox, D., and Schuller, T. (**2013**). "Impact of flame base dynamics on the non-linear frequency response of conical flames," C. R. Mecanique **341**(1–2), 171–180.

Dowling, A. P., and Stow, S. R. (**2003**). "Acoustic analysis of gas turbine combustors," J. Propul. Power **19**(5), 751–764.
Ducruix, S., Durox, D., and Candel, S. (**2000**). "Theoretical and experimental determination of the flame transfer function of a laminar premixed flame," Proc. Combust. Inst. **28**, 765–773.
Eder, A. J., Merk, M., Hollweck, T., Fischer, A., Lahiri, C., Silva, C. F., and Polifke, W. (**2025**). "Model-based inference of flame transfer matrices from acoustic measurements in an aero-engine test rig," J. Eng. Gas Turbines Power **147**, 1–31.
Gatti M., Gaudron R., Mirat, C., Zimmer, L., and Schuller, T. (**2018**). "A comparison of the transfer functions and flow fields of flames with increasing swirl number," in *Proceedings of the ASME Turbo Expo 4B-2018*, pp. 1–12.
Howe, M. S. (**1979**). "On the theory of unsteady high Reynolds number flow through a circular aperture," Proc. R Soc. London Ser. A **366**(1725), 205–223.
Howe, M. S., Scott, M., and Sipcic, S. (**1996**). "The influence of tangential mean flow on the Rayleigh conductivity of an aperture," Proc. Roy. Soc. A **452**, 2303–2317.
Huber, A., and Polifke, W. (**2009**). "Dynamics of practical premixed flames, Part II: Identification and interpretation of CFD data," Int. J. Spray Combust. Dyn. **1**(2), 229–249.
Jaynes, E. (**2003**). *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, UK).
Jeffreys, H. (**1973**). *Scientific Inference*, 3rd ed. (Cambridge University Press, Cambridge, UK).
Jeffreys, H. (**1983**). *Theory of Probability*, 3rd ed. (Claredon Press, Oxford, UK).
Jing, X., and Sun, X. (**1999**). "Experimental investigations of perforated liners with bias flow," J. Acoust. Soc. Am. **106**(5), 2436–2441.
Juniper, M. P. (**2018**). "Sensitivity analysis of thermoacoustic instability with adjoint Helmholtz solvers," Phys. Rev. Fluids **3**(11) 110509.
Juniper, M. P., and Sujith, R. (**2018**). "Sensitivity and nonlinearity of thermoacoustic oscillations," Ann. Rev. Fluid Mech. **50**, 661–689.
Juniper, M. P., and Yoko, M. (**2022**). "Generating a physics-based quantitatively-accurate model of an electrically-heated Rijke tube with Bayesian inference," J. Sound Vib. **535**, 117096.
Karmarkar, A., Gupta, S., Boxx, I., Hemchandra, S., and O'Connor, J. (**2022**). "Impact of precessing vortex core dynamics on the thermoacoustic instabilities in a swirl-stabilized combustor," J. Fluid Mech. **946**, 1–32.
Kass, R. E., and Raftery, A. E. (**1995**). "Bayes factors," J. Am. Stat. Assoc. **90**(430), 773–795.
Kornilov, V. N., Schreel, K. R., and De Goey, L. P. (**2007**). "Experimental assessment of the acoustic response of laminar premixed Bunsen flames," Proc. Combust. Inst. **31**(1), 1239–1246.
MacKay, D. J. (**1999**). "Comparison of approximate methods for handling hyperparameters," Neural Comput. **11**(5), 1035–1068.
MacKay, D. J. C. (**2003**). *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, UK).
Mejia, D., Miguel-Brebion, M., and Selle, L. (**2016**). "On the experimental determination of growth and damping rates for combustion instabilities," Combust. Flame **169**, 287–296.
Munjal, M. L., and Doige, A. G. (**1990**). "Theory of a two source-location method for direct experimental evaluation of the four-pole parameters of an aeroacoustic element," J. Sound Vib. **141**, 323–333.
Noiray, N., and Denisov, A. (**2017**). "A method to identify thermoacoustic growth rates in combustion chambers from dynamic pressure time series," Proc. Combust. Inst. **36**(3), 3843–3850.
Nygård, H. T., and Worth, N. A. (**2021**). "Flame transfer functions and dynamics of a closely confined premixed bluff body stabilized flame with swirl," J. Eng. Gas Turbines Power **143**(4), 1–10.
Paschereit, C. O., Schuermans, B., Polifke, W., and Mattson, O. (**1999**). "Measurement of transfer matrices and source terms of premixed flames," in *Proceedings of the ASME Turbo Expo*.
Peterleithner, J., Stadlmair, N. V., Woisetschläger, J., and Sattelmayer, T. (**2016**). "Analysis of measured flame transfer functions with locally resolved density fluctuation and OH-chemiluminescence data," J. Eng. Gas Turbines Power **138**(3), 1–9.
Polifke, W. (**2020**). "Modeling and analysis of premixed flame dynamics by means of distributed time delays," Prog. Energy Combust. Sci. **79**, 100845.
Polifke, W., and Lawn, C. (**2007**). "On the low-frequency limit of flame transfer functions," Combust. Flame **151**(3), 437–451.
Rao, K., and Munjal, M. L. (**1986**). "Experimental evaluation of impedance of perforates with grazing flow," J. Sound Vib. **108**(2), 283–295.
Schuermans, B., Bellucci, V., Guethe, F., Meili, F., Flohr, P., and Paschereit, C. O. (**2004**). "A detailed analysis of thermoacoustic interaction mechanisms in a turbulent premixed flame," *Proceedings of the ASME Turbo Expo 2004*, pp. 539–551.
Schuermans, B., Guethe, F., and Mohr, W. (**2010**). "Optical transfer function measurements for technically premixed flames," J. Eng. Gas Turbines Power **132**(8), 1–8.
Schuermans, B., Polifke, W., and Paschereit, C. O. (**1999**). "Modeling transfer matrices of premixed flames and comparison with experimental results," in *Proceedings of the ASME Turbo Expo*.
Seybert, A. F., and Ross, D. F. (**1977**). "Experimental determination of acoustic properties using a two-microphone random-excitation technique," J. Acoust. Soc. Am. **61**(5), 1362–1370.
Shreekrishna, V. A., and Lieuwen, T. (**2013**). "Flame response to equivalence ratio fluctuations—Relationship between chemiluminescence and heat release," Int. J. Spray Combust. Dyn. **5**(4), 329–358.
Wang, H., Xu, R., Wang, K., Bowman, C. T., Hanson, R. K., Davidson, D. F., Brezinsky, K., and Egolfopoulos, F. N. (**2018**). "A physics-based approach to modeling real-fuel combustion chemistry—I. Evidence from experiments, and thermodynamic, chemical kinetic and statistical considerations," Combust. Flame **193**, 502–519.
Xu, R., Wang, K., Banerjee, S., Shao, J., Parise, T., Zhu, Y., Wang, S., Movaghar, A., Lee, D. J., Zhao, R., Han, X., Gao, Y., Lu, T., Brezinsky, K., Egolfopoulos, F. N., Davidson, D. F., Hanson, R. K., Bowman, C. T., and Wang, H. (**2018**). "A physics-based approach to modeling real-fuel combustion chemistry—II. Reaction kinetic models of jet and rocket fuels," Combust. Flame **193**, 520–537.
Yoko, M., and Juniper, M. P. (**2024a**). "Adjoint-accelerated Bayesian inference applied to the thermoacoustic behaviour of a ducted conical flame," J. Fluid Mech. **985**, A38.
Yoko, M., and Juniper, M. P. (**2024b**). "Inferring flame transfer functions of turbulent conical flames from pressure measurements," J. Eng. Gas Turbines Power **147**, 051007.
Yoko, M., Lahiri, C., Fischer, A., and Juniper, M. P. (**2026**). "Learning the relationship between operating condition and flame response from data," J. Acoust. Soc. Am. **159**(6), 5188–5198.