

Generating a physics-based quantitatively-accurate model of an electrically-heated Rijke tube with Bayesian inference

Matthew P. Juniper*, Matthew Yoko

Cambridge University Engineering Department, Cambridge, CB2 1PZ, UK

Abstract

We perform 7000 experiments at 175 stable operating points on an electrically-heated Rijke tube. We pulse the flow and measure the acoustic response with eight probe microphones distributed along its length. We assimilate the experimental data with Bayesian inference by specifying candidate models and calculating their optimal parameters given prior assumptions and the data. We model the long timescale behaviour with a 1D pipe flow model driven by natural convection into which we assimilate data with an Ensemble Kalman filter. We model the short timescale behaviour with several 1D thermoacoustic network models and assimilate data by minimizing the negative log posterior likelihood of the parameters of each model, given the data. For each candidate model we calculate the uncertainties in its parameters and calculate its marginal likelihood (i.e. the evidence for that model given the data) using Laplace's method combined with first and second order adjoint methods. We rank each model by its marginal likelihood and select the best model for each component of the system. We show that this process generates a model that is physically-interpretable, as small as possible, and quantitatively accurate across the entire operating regime. We show that, once the model has been selected, it can be trained on little data and can extrapolate successfully beyond the training set. Matlab code is provided so that the reader can experiment with their own models.

*Corresponding author

Email address: mpj1001@cam.ac.uk (Matthew P. Juniper)

Keywords: Thermoacoustics, Data assimilation, Bayesian inference, adjoint methods, Laplace, saddle point

1. Introduction

The Rijke tube was one of Shon Ffowcs-Williams' demonstrations in his course on acoustics, which was given to the Engineering undergraduates at Cambridge University in the 1990's (Fig. 14 of [1]). An electrical heater is placed in the bottom half of a vertical tube, underneath which is a loudspeaker. When the heater is switched on, the tube begins to hum. When the loudspeaker is switched on, the humming stops. To the students, of which M. Juniper was one, it is a remarkable demonstration that noise can beget silence. To Shon Ffowcs-Williams, it was the big reveal for his lecture on the applications of *anti-sound* [1].

This application of anti-sound is now called *feedback control of thermoacoustic oscillations*, as implemented for example by Shon Ffowcs-Williams students Ann Dowling [2] and Maria Heckl [3]. Since the 1990's the simple phase shift controller has been upgraded to robust model-based control, adaptive control, and model-based active control. Feedback control has been achieved in practice at full scale [4, 5], activated through fuel modulation rather than a loudspeaker. It has not, however, been implemented routinely in industry because actuators would have to operate for several years, typically at 100 to 1000 Hz, without failing. Failure would risk thermoacoustic oscillations in $\mathcal{O}(100)$ MW ground-based gas turbines or $\mathcal{O}(10)$ MW aircraft engines, which could be catastrophic.

The trepidation towards feedback control of thermoacoustic oscillations in practical devices has motivated the search for passive control mechanisms [4]. An effective mechanism, which is widely-used in practice, is to fit acoustic dampers such as acoustic liners or Helmholtz resonators. Helmholtz resonators can, if space allows, be retro-fitted after an engine has been designed in order to damp modes identified in full engine tests. However, they add weight and another failure mode. The ideal scenario is therefore to manufacture a device that

is linearly stable across the entire operating regime without having to retro-fit Helmholtz resonators or to avoid particular operating regimes. The design of such a device can exploit the fact that the behaviour of thermoacoustic systems is exceedingly sensitive to small design changes [6], meaning that it may well be possible to stabilize all modes with small changes. In engines that cannot be fitted with Helmholtz resonators, small stabilizing design changes have always been discovered, albeit after extensive full-scale testing [7]. The challenge is to find these small changes during the early design phase, while also satisfying the primary requirements for high combustion efficiency, low pollution, and high altitude relight.

Gradient-based optimization combined with adjoint methods provide an efficient way to converge to stable designs [8, 9]. This requires, however, a quantitatively accurate model of the thermoacoustic behaviour of the system. Devising such a model is particularly challenging because the sensitivity mentioned above usually introduces significant systematic error into models and their parameters. Although it is possible to create *qualitatively*-accurate thermoacoustic models, it is therefore difficult to create a *quantitatively*-accurate model of a particular system, even from quantitatively-accurate models of its components [10]. For example, a highly detailed model of thermoacoustic oscillations of an electric heater in a tube [11], despite being carefully tuned to be quantitatively correct at one heater position, was only qualitatively correct at nearby heater positions [11, Figs. 5-5 to 5-8]. For large devices, accurate prediction of thermoacoustic behaviour is similarly challenging [12].

The solution proposed here is to combine physics-based modelling of the Rijke tube [3, 11, 13, 14] with statistical inference from thousands of experimental observations [15]. The remarkable recent success of data-driven approaches lies in their relentless focus on data, rather than on models, correlations, and assumptions that the research community has become used to. Rather than throw away these models entirely, however, we devise qualitatively-accurate physics-based candidate models of the components of a thermoacoustic system and then rigorously (i) tune their parameters by assimilating data from experiments; (ii)

quantify the uncertainties in each model's parameters; (iii) quantify the evi-
60 dence (the marginal likelihood) for each model; (iv) select the best model and
(v) repeat for the next component until the model of the thermoacoustic system
is complete. We use Laplace's method [15, 16] combined with adjoint methods
to first and second order [17, 18], which is technically more difficult to imple-
65 ment than methods such as Markov Chain Monte Carlo, but is thousands of
times faster, meaning that we can compare dozens of candidate models. This
paper is an extension of a conference paper [19] and Matlab code is provided so
that the interested reader can extend this study to new models [20].

Shon Ffowcs-Williams' vertical Rijke tube provides an ideal demonstration of
the application of physics-based Bayesian inference to develop a quantitatively-
70 accurate thermoacoustic model. On the one hand, a hot wire Rijke tube is easy
to automate, meaning that thousands of datapoints can be obtained cheaply[21].
On the other hand, a hot wire Rijke tube is difficult to model accurately. Firstly
the heat release rate is small, meaning that many visco-thermal dissipation
mechanisms are sufficiently large that they must be included in the model [13].
75 Secondly, as for all acoustic systems, the reflection coefficients at the ends of
the tube have a strong influence on the thermoacoustic behaviour. For a quan-
titatively accurate model, they must be measured in situ rather than taken
from models in the literature [22]. Thirdly, the heat release rate fluctuations at
the wire cannot be measured or observed directly. As we will show later, they
80 also need to be inferred in situ rather than taken from models in the literature
[3, 11, 14, 23, 24, 25], which are at best only qualitatively-accurate. Finally, as
Shon Ffowcs-Williams demonstrated for the applications of anti-sound, the hot
wire Rijke tube is an ideal starting point for new methods in thermoacoustics.
If we cannot create a quantitatively accurate model of a hot wire Rijke tube,
85 what hope is there for aircraft and rocket engines?

2. Physics-based modelling of the hot wire Rijke tube

There are several hundred papers that describe or model the Rijke tube. These are listed in the handful of reviews of this subject over the last 50 years [6, 26, 27, 28]. The physical mechanism that causes thermoacoustic instability was correctly described by Rayleigh [29]. Later, Chu [30] re-expressed this in an elegant and complete mathematical framework based on small perturbations to the governing equations. The thermoacoustic mechanism in the hot wire Rijke tube is described in [27, §3.3.1], which highlights the crucial role played by the time delay between velocity perturbations and heat release rate perturbations at the hot wire. If the hot wire is in the upstream (downstream) half of the tube, then this time delay causes the heat release rate to be slightly in phase (out of phase) with the acoustic pressure of the first acoustic mode, causing acoustic oscillations to be thermoacoustically driven (damped). Using analytical methods, Carrier [13] estimated the gain and phase of the heat release rate of a hot flat ribbon in a fluctuating air stream. Lighthill [24] performed a similar analysis for a hot circular cylinder and Merk [14] for a hot wire gauze. These analytical methods show that the phase lag is caused by heat conduction through the finite thickness boundary layers around the hot element and explain why the Rijke tube requires a through-flow. This is described by Bayly [31]: “If the blowing is too weak, the wires are surrounded by very thick jackets of stagnant air, and the fluctuations in the external flow have a comparatively small effect on the heat transfer. On the other hand, if the blowing is strong, the phase lag between the velocity fluctuations and the induced heat transfer fluctuations is small ... Although the heat transfer is more efficient with strong blowing, the absence of sufficient phase matching makes the amplification ineffective.”

The above analytical methods on simplified models, although qualitatively correct, are not quantitatively correct. Subsequent numerical simulations of heat transfer in an oscillating flow around a cylinder have shown that the time delay is two to three times greater than that predicted by Lighthill [32, 33], and that the gain and phase have more intricate dependence on Re and St than

can be derived analytically [25]. Instead, since the 1970's [31], researchers have tended to use simple linear relations such as the $n - \tau$ law, empirical alterations to King's law [34], or CFD simulations. This is shown in a detailed survey of heat transfer models from 1985 to 2017 [35, Table 3.1]. CFD simulations are
120 accurate but are too expensive for simple thermoacoustic models. All simple thermoacoustic models of the Rijke tube are therefore relying on heat release rate models that are, at best, only qualitatively accurate. The aim of this paper is to develop a general method that tunes these models with data such that they become quantitatively accurate with known uncertainties.

125 Most studies that aim to create quantitatively accurate models of the Rijke tube include viscous dissipation in the acoustic boundary layer along the inside wall of the tube [11, 13, 27, 32]. Viscous dissipation is greatest where the acoustic velocity fluctuation is greatest, which is at the ends of the Rijke tube. Some studies also include thermal dissipation in the acoustic boundary layer,
130 whose effect has a similar magnitude. Thermal dissipation is greatest where the acoustic temperature (equiv. pressure) fluctuation is greatest, which is at the centre of the Rijke tube. In Rayleigh [36, §348–350] and Kinsler et al. [37, §9.5] thermal dissipation is included through an added viscosity (note that [37, Eq.(9.34)] erroneously misses a factor of $\sqrt{\gamma}$ when adapting from [36, §350
135 Eq.(16)]). By modelling thermal dissipation as increased viscosity, one loses the important detail that the location of thermal dissipation differs from the location of viscous dissipation. In this paper we model both separately so that the two mechanisms can be disentangled.

All models of the Rijke tube must include heat transfer between the hot wire
140 and the surrounding air. The mechanism that causes heat transfer (molecular diffusion) is the same as that which causes momentum transfer, so there cannot be one without the other. The question is whether the momentum transfer is so small that it can be neglected. This momentum transfer manifests itself as an acoustic pressure drop across the heater that is, in the linear regime,
145 proportional to the velocity at the heater. This pressure drop is considered by [11, p84], [13, Eq.4.4], [38, Eq.34] [39, §3.3] but is always set to zero on the

grounds that it is small compared with the acoustic pressure amplitude. This simplifies the analysis but asserts that the drag from an object blocking the duct is negligible compared with the drag from the acoustic boundary layers
150 of the duct, which is questionable. In this paper we show that neither can be neglected.

All studies of the Rijke tube agree on the strong influence of the time or phase delay between velocity fluctuations and subsequent heat release rate fluctuations at the heater. This time delay can influence the frequency and the growth rate
155 by similar absolute values, although the relative influence on the growth rate is larger because the growth rate is close to zero, while the frequency is close to the resonant frequency of the tube. Indeed, after his famous quote about vibrations being encouraged, Rayleigh [29] writes: “If the air be at its normal density at the moment when the transfer of heat takes place, the vibration is
160 neither encouraged nor discouraged, but the pitch is altered.” An efficient data-driven approach will therefore use the frequency drift as well as the growth rate drift in order to infer the time delay. This requires an accurate measurement of the speed of sound in the tube, which we perform in §3.1.

3. Experimental configuration

The experimental configuration is a vertical Rijke tube containing an electric
165 heater, detailed in [40, 41, 42]. The electric heater (Fig. 1) sets up natural convection within the tube, which is essential for the thermoacoustic mechanism [13]. The stainless steel tube is 1 m long, has internal diameter 47.4 mm and wall thickness 1.7 mm. These dimensions are almost identical to those in Heckl
170 [3] (1 m long with diameter 44.6 mm), although that tube was horizontal rather than vertical, meaning that suction was required to create a flow through the tube (of 0.6 ms^{-1}), which is about the same as ours. Our heater consists of two concentric annular ceramic plates with inner diameter 31.6 mm and outer diameter 47.0 mm (Fig. 1). Flow passes through the central hole in
175 each annulus and over nichrome wire with diameter 0.6 mm, which is wrapped

Table 1: The number of experimental measurements (40) assimilated at each operating point. Column 1 describes the type of experiment: C1 and C2 denote experiments with the prongs and heater attached but switched off; C3 denote experiments with the prongs only; H denotes experiments with the prongs and heater attached and switched on. Column 2 contains a check mark if the thermocouples were present. Column 3 contains the heater power in Watts. Columns 4 to 23 contain the number of experimental measurements taken at the heater position shown in the second row, where E denotes the empty tube, in which the heater and prongs were removed.

			Heater position (cm)																				
	T	Q	E	5	10	15	20	25	30	35	40	45	50	55	60	65	70	75	80	85	90	95	
C1	–	0	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–	
C2	–	0	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40
C3	–	0	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40
H	✓	0	40	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–
H	✓	7	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–
H	✓	15	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–
H	✓	30	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–
H	✓	50	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–
H	✓	80	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–
H	✓	130	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–
H	✓	180	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–




Figure 1: (a) Top view, (b) side view, and (c) isometric view of the heater, which consists of two identical concentric annular ceramic plates, each wound with nichrome wire. It is held in place by two threaded prongs (shown), which enter from the upstream end of the tube. The dimensions are $d = 47.0$ mm, $d_i = 31.6$ mm, $d_w = 0.6$ mm, $t = 5.0$ mm, $h = 5.0$ mm, $d_p = 3.0$ mm. The nichrome wire is supplied by by two fabric-coated copper wires (not shown), which each have diameter 4.0 mm.

around the plates. The design of the annular plates is the same as that in Heckl [3], although the dimensions differ slightly. The electric heater is placed at 19 different positions from the bottom end of the tube (Tab. 1). Eight probe microphones record the pressure near the inner surface of the tube from $x_m/L = 0.25$ to 0.95 in steps of 0.1. Eight thermocouples are placed from $x_t/L = 0.2$ to 0.9 in steps of 0.1. The ambient temperature is measured with a thermocouple at the inlet of the tube.

The heater power is set to 0, 7, 15, 30, 50, 80, 130, 180 Watts, for 62.5 minutes at each power. Every 15 seconds, a loudspeaker at the base of the tube forces the system close to its resonant frequency for 6 seconds, referred to in this paper as a ‘ping’. The probe microphones measure the response throughout the tube at 10 kHz during the forced and the decaying period.

The experiment exhibits a long timescale (typically 1000 seconds), at which the tube heats up and reaches steady state, and a short timescale (typically 1 second), at which the acoustic oscillations decay. At the long timescale we assimilate the thermocouple and sound speed measurements. At the short timescale

we assimilate the decay rate measurements, assuming that the long timescale flow is quasi-steady. The quasi-steady assumption is reasonable because of the large difference in the timescales.

195 In this study we assimilate data only from experiments that are thermoacoustically stable. In previous studies on the same rig [21, 40] we measured the linear growth rates of unstable thermoacoustic oscillations by stabilising the system with feedback control and then switching off the control. The decision to assimilate linear decay rates from the stable system makes no difference to
200 the assimilation process demonstrated in this paper, but greatly simplifies the automation of the experiments. This decision also makes no difference to the thermoacoustic driving mechanism, which is active whether or not it is stronger than visco-thermal damping mechanisms and acoustic radiation from the ends of the tube. Our aim is to assimilate experimental measurements of the sta-
205 ble system in order to construct a physics-based model of the thermoacoustic driving mechanism. In this paper, we compare this model against experimental results from the stable system and confirm that this model is accurate over this range of heater powers. Because the model of the thermoacoustic mechanism is physics-based, it can extrapolate successfully to higher heater powers, whether
210 or not the system is linearly unstable at those heater powers. While we do not check that in this paper, Fig. 4(a) of [21] shows that the thermoacoustic linear growth rate increases smoothly with heater power from negative growth rates (stable system) to positive growth rates (unstable system). This shows that there is no abrupt change in the thermoacoustic driving mechanism as the
215 system changes from stable to unstable, implying that a physics-based model developed in the stable regime should extrapolate to the unstable regime. In other words, being physics-based rather than physics-agnostic, the assimilated model will be able to correctly predict positive growth rates, even if it has assimilated data only from negative growth rates. This makes the current study
220 relevant to unstable thermoacoustic systems, and also to the design of automated experiments on combustion systems: one can develop models of the heat release rate fluctuations at the flame (e.g. FTF, $n - \tau$) in stable configurations,

which are safer, and then extrapolate them to unstable configurations.

3.1. Measurement of the local speed of sound

225 The acoustic frequency at the short timescale depends strongly on the local speed of sound in the tube. A previous study [41, 42] inferred the local speed of sound from the local gas temperature measurements. We suspect, however, that radiation from the heater introduces excessive systematic error into this measurement. In this study we therefore measure the local speed of sound with
230 the probe microphones. Before each ping, the loudspeaker sends an impulse down the tube, known in this paper as a ‘click’. The 8 probe microphones measure the response along the tube at 62.5 kHz. We find the phase shift that maximises the cross-correlation function between the microphones (i.e. brings the measured impulses optimally in phase). This process yields the time at
235 which the impulse arrives at each microphone and therefore an estimate of the local sound speed. This local sound speed differs by up to 5% from that inferred from the thermocouple measurements. This local sound speed is fed into a model of the long timescale flow via an Ensemble Kalman Filter, described in the next section. In this study we infer the local temperature inside the tube from this
240 model. The only thermocouple we use is that which measures the ambient temperature.

The channels of the acquisition system have slightly different delays. We measure these delays by performing several clicks at ambient sound speed. We then subtract these delays from the measurements, thus removing a significant
245 source of systematic error. We find that this delay depends on the sampling frequency so we perform a second calibration for the ping experiments, which were sampled at 10 kHz.

4. Overview of the physics-based models

4.1. 1D pipe flow model for the long timescale flow

250 The long timescale flow is an unsteady 1D flow conjugate heat transfer model extended from a previous study [41] to include conductive cooling of the tube




Figure 2: (a) Velocity at the heater, U_h , as a function of heater position, X_h , and heater power, Q_h . (b) Local speed of sound, $c(x)$, as a function of heater position, X_h , at $Q_h = 180$ W. These results are calculated from the 1D pipe flow model (§4.1) for the long timescale flow after assimilating data from the click experiments (§3.1) and the ambient thermocouple.

through contact with the support structure, conductive heating of the tube through thermal contact with the heater, heat loss from the wires upstream of the heating element, and variation of Nusselt number along the inner wall. Four parameters characterize (i) the inviscid drag coefficient of the heater, (ii) the Nusselt number on the outer surface of the tube, (iii) the thermal resistance at the tube mounts, and (iv) the proportion of supplied power that conducts through the tube wall. Two further parameters characterize the Nusselt number distribution inside the tube. For a given configuration, the model outputs the velocity of the convective flow, the local temperature in the gas, and the local temperature on the inner and outer walls of the tube.

One hundred realisations of the long timescale model are iterated as an ensemble. The parameters of each ensemble member are randomly sampled from a uniform prior distribution. At each assimilation step, the ensemble forecast and experimental measurements are supplied to an Ensemble Kalman Filter (EnKF) [43]. These measurements consist of the ambient temperature, the click arrival time at each microphone, and the measured heater power. The EnKF returns the expected values and variances of the system’s state and the model’s parameters. As more data become available, the state and parameters converge to constant values with high certainty. Figure 2 shows the velocity at the heater and the local sound speed, as functions of the heater position and heater power. As X_h increases, the length of the column of hot air above the heater decreases, so the buoyancy force driving the flow decreases, causing the velocity at the heater U_h to decrease (Fig. 2a) and, at a given heater power, the temperature jump across the heater to increase (2b). Although these trends should be smooth, there is clear evidence of measurement and model error in Fig. 2: (a) U_h is expected to decrease smoothly with X_h but does not; (b) the maximum of $c(x)$ is expected to increase smoothly with X_h but does not at $X_h = 0.75$ m. The main source of error is the fact that the number of microphones downstream of the heater decreases as X_h increases, so there are fewer measurements of the downstream sound speed during the click test. It is impractical to remove all this error. Instead, we group the measurement and

model errors and check a-posteriori that the combined error is small compared with the trends in the model predictions. We note that U_h is derived from several measured quantities, whose errors accumulate into the error in U_h , and that the model assumes uniform flow at the heater location, which is a simplification. We must therefore be aware that U_h contains more systematic uncertainty and model error than the other variables.

4.2. Thermoacoustic Network Model for the short timescale flow

Acoustic waves are modelled as forward-travelling waves, $f(t - x/c)$, and backward-travelling waves, $g(t + x/c)$, in N acoustic elements within the tube [6, 41]. The speed of sound, c , and density, ρ , in each element are extracted from the long timescale model. In element i , the pressure is $p_i = f_i + g_i$ and the velocity is $u_i = (f_i - g_i)/(\rho_i c_i)$. At the interfaces between the acoustic elements, the waves in adjacent elements are related through jump conditions for the momentum and energy equations. The linear influence of all components of the network model can be expressed in terms of local linear feedback from velocity or pressure into the momentum or energy equations [44, §VI(A)]. Labelling these feedback coefficients k_{**} , the jump conditions are:

$$p_{i+1} - p_i = -k_{mu} u_i - k_{mp} p_i \quad (1)$$

$$u_{i+1} - u_i = -k_{eu} u_i - k_{ep} p_i \quad (2)$$

For example, viscous drag is modelled as local feedback from the velocity into the momentum equation (k_{mu}), thermal drag is modelled as local feedback from the temperature (equiv. pressure) into the energy equation (k_{ep}), and the heat release rate from the wire is modelled as local feedback from the velocity into the energy equation (k_{eu}). Local feedback from the pressure into the momentum equation, k_{mp} , is not required for this study. In §6 these local feedback coefficients are derived from candidate physics-based models and are expressed in terms of those models' parameters. Wave reflection at the upstream and downstream ends of the tube are modelled by complex reflection coefficients, R_u and R_d .

310 The equations are converted to the frequency domain through modal decom-
 positions such as $f(t - x/c) = Fe^{st}e^{-sx/c}$, where $-is$ is the complex angular
 frequency of oscillations. There are $2N$ unknown amplitudes (N each for F
 and G) and $2N$ constraints ($2(N - 1)$ jump conditions and 2 reflection condi-
 tions). This creates a nonlinear eigenvalue problem for the eigenvalue s , which
 315 is solved with Newton iteration. We calculate the corresponding eigenfunction
 $P(x, s)$ from F and G . We then use first order [17] and second order [18] adjoint
 methods to obtain the first and second derivatives of s and P with respect to
 all the local feedback coefficients $k_{\star\star}$ and reflection coefficients R_{\star} . We then
 combine these derivatives to obtain the first and second derivatives of s and P
 320 with respect to the model parameters of the components of the network model.
 Details can be found in comments in the Matlab code supplied with this paper
 [20].

5. Data assimilation with Bayesian Inference

We denote each physics-based model as \mathcal{H}_i with a set of variable parameters,
 325 **a**. We label the data as D , which consists of a complex angular frequency z and
 a complex pressure measurement Q at each microphone location.

5.1. Maximum a Posteriori parameter estimation

First we assume that each model is true and infer its parameters from the
 data D [15]. The posterior probability of the parameters **a** is:

$$P(\mathbf{a}|D, \mathcal{H}_i) = \frac{P(D|\mathbf{a}, \mathcal{H}_i)P(\mathbf{a}|\mathcal{H}_i)}{P(D|\mathcal{H}_i)} \quad (3)$$

At this level of inference the denominator of (3) is ignored and the *maximum
 posterior likelihood* of the model parameters, \mathbf{a}_{MP} , is found by maximizing the
 330 numerator.

The data, D , consists of nine complex numbers for each ping: the eigen-
 value, z (whose real part is the growth rate and whose imaginary part is the
 angular frequency), and the Fourier-decomposed complex pressure, Q , at the

eight microphone locations. A cost function, \mathcal{J} , is defined as the negative log of
 335 the posterior likelihood. If all distributions are assumed to be Gaussian then \mathcal{J}
 is the sum of the squares of the discrepancy between the model parameters and
 their prior estimates, weighted by the confidence in the prior estimates, added to
 the sum of the squares of the discrepancies between the model predictions and
 the experimental measurements, weighted by the confidence in the experimental
 340 measurements:

$$\begin{aligned}
 \mathcal{J} &= -\log \{P(D|\mathbf{a}, \mathcal{H}_i)P(\mathbf{a}|\mathcal{H}_i)\} \\
 &= (s_r(\mathbf{a}) - z_r)^T C_{sr}^{-1} (s_r(\mathbf{a}) - z_r) \dots \\
 &+ (s_i(\mathbf{a}) - z_i)^T C_{si}^{-1} (s_i(\mathbf{a}) - z_i) \dots \\
 &+ (P_r(\mathbf{a}) - Q_r)^T C_{pr}^{-1} (P_r(\mathbf{a}) - Q_r) \dots \\
 &+ (P_i(\mathbf{a}) - Q_i)^T C_{pi}^{-1} (P_i(\mathbf{a}) - Q_i) \dots \\
 &+ (\mathbf{a} - \mathbf{a}_f)^T C_{aa}^{-1} (\mathbf{a} - \mathbf{a}_f) \tag{4}
 \end{aligned}$$

where: \mathbf{a} is a column vector containing the parameter values; $s_r(\mathbf{a})$ is the model's
 growth rate; z_r is the measured growth rate in the experimental configuration
 represented by parameters \mathbf{a} ; C_{sr} is the covariance matrix of the growth rate
 measurements; $s_i(\mathbf{a})$ is the model's frequency; z_i is the measured frequency; C_{si}
 345 is the covariance of the frequency measurements; $P_r(\mathbf{a})$ and $P_i(\mathbf{a})$ are the real
 and imaginary components of the pressure predictions at each microphone; Q_r
 and Q_i are the real and imaginary components of the pressure measurements;
 C_{pr} and C_{pi} are the covariance matrices of the pressure measurements; C_{aa} is
 the prior covariance matrix of the parameters; \mathbf{a}_f is a column vector containing
 350 the prior estimates of the parameter values.

The first derivative of the cost function \mathcal{J} with respect to the parameters,
 \mathbf{a} , is derived via a first order Taylor expansion of (4). It is expressed in terms
 of the first derivatives of s and P with respect to the parameters. The deriva-
 tives of s and P are calculated from the network model using first order adjoint
 355 methods [17]. A BFGS (Broyden-Fletcher-Goldfarb-Shanno) gradient-based op-
 timization algorithm is used to find the parameter values, \mathbf{a}_{MP} , that minimize

the cost function \mathcal{J} , given the prior covariance C_{aa} , and the measurement variances $C_{sr}, C_{si}, C_{pr}, C_{pi}$. This is known as the *maximum a posteriori estimation* (MAP).

360 If the measurement uncertainties are unknown, for example if an un-measured or unknown factor is affecting the measurements, then the measurement variances $C_{sr}, C_{si}, C_{pr}, C_{pi}$ can be allowed to float. The derivative of \mathcal{J} with respect to $C_{sr}, C_{si}, C_{pr}, C_{pi}$ is easily calculated from (4). These derivatives can be included in the BFGS optimization algorithm in order to find the values of
 365 $C_{sr}, C_{si}, C_{pr}, C_{pi}$ and \mathbf{a}_{MP} at which \mathcal{J} is minimized.

5.2. Parameter error estimation with Laplace's Method

The MAP method gives the most likely parameters, \mathbf{a}_{MP} , given a model, \mathcal{H}_i , but does not yet give the uncertainties in the parameters. This is found with the Laplace Approximation, which is also known as the Saddle Point Method. In order to estimate these uncertainties we re-use the assumption that $P(\mathbf{a}|D, \mathcal{H}_i)$ is Gaussian around \mathbf{a}_{MP} and we define its inverse covariance matrix around this point as \mathbf{A} :

$$-\log \{P(\mathbf{a}|D, \mathcal{H}_i)\} = \frac{1}{2}(\mathbf{a} - \mathbf{a}_{\text{MP}})^T \mathbf{A}(\mathbf{a} - \mathbf{a}_{\text{MP}}) + \text{const} \quad (5)$$

By inspection, \mathbf{A} is simply the Hessian of \mathcal{J} :

$$A_{ij} = \frac{\partial^2 \mathcal{J}}{\partial a_i \partial a_j} \quad (6)$$

The second derivative of the cost function \mathcal{J} with respect to the parameters, \mathbf{a} , is derived via a second order Taylor expansion of (4) around \mathbf{a}_{MP} . This is expressed in terms of the first and second derivatives of s and P with respect
 370 to the parameters. The second derivatives, like the first derivatives, are found with adjoint methods [18]. It is worth mentioning that the contribution of the second derivatives to \mathbf{A} is usually an order of magnitude smaller than the contribution of the first derivatives. It is also worth mentioning that the approximate Hessian created during the BFGS optimization algorithm is usually an excellent
 375 approximation to \mathbf{A} . This means that the second derivatives of s and P , which are expensive to calculate, often do not need to be calculated.

5.3. Model Comparison with Laplace's Method

At the second level of inference, we wish to infer which model is most plausible given the data [15]. By Bayes' theorem the posterior probability of each model is:

$$P(\mathcal{H}_i|D) \propto P(D|\mathcal{H}_i)P(\mathcal{H}_i) \quad (7)$$

On the RHS, the second term is our prior estimate for the probability of model \mathcal{H}_i , while the first term is the denominator of (3), which is known as the *evidence* or *marginal likelihood*. Assuming that each candidate model has equal probability $P(\mathcal{H}_i)$ then the models are ranked by evaluating the evidence:

$$P(D|\mathcal{H}_i) = \int_{\mathbf{a}} P(D|\mathbf{a}, \mathcal{H}_i)P(\mathbf{a}|\mathcal{H}_i)d\mathbf{a} \quad (8)$$

Assuming furthermore that $P(\mathcal{H}_i|D)$ is Gaussian, this is given by:

$$P(D|\mathcal{H}_i) \approx P(D|\mathbf{a}_{\text{MP}}, \mathcal{H}_i)P(\mathbf{a}_{\text{MP}}|\mathcal{H}_i) (\det(\mathbf{A}/2\pi))^{-1/2} \quad (9)$$

where \mathbf{A} is the Hessian calculated in (6). This evidence, which is also known as the Marginal Likelihood, is calculated cheaply with Laplace's Method. The most likely model is that with the largest Marginal Likelihood. The larger the number of parameters, \mathbf{a} , in the model, \mathcal{H}_i , the larger the Hessian, \mathbf{A} . This tends to penalize over-elaborate models unless they fit the data extremely well.

The Marginal Likelihood, like the MAP point, depends on the measurement uncertainty. Laplace's method also gives the gradient of the Marginal Likelihood with respect to the measurement uncertainties. This gradient can then be used to find the Maximum Marginal Likelihood (MML) of a model, given some data, allowing the measurement noise to float. In this paper we calculate only the MAP parameter values because the MML parameter values turn out to be almost identical.

6. Creating the quantitatively-accurate physics-based model

In this section we construct a quantitatively accurate model of the Rijke tube by assembling several component models, each of which is the candidate

model with the highest marginal likelihood given the experimental data. We assimilate data from the cold experiments first and then assimilate data from the hot experiments. The network model (§4.2) of the empty tube contains $N = 40$ elements with equal length. When the heater is added to the model, the acoustic element at the heater location is split into two elements, with jump conditions (1,2) at the heater. The same procedure is followed when the 8 thermocouples are added. The network model therefore contains between 40 and 50 acoustic elements, N , depending on the configuration. We have checked that the model predictions do not change significantly with N for $N > 40$. In the Matlab code supplied with this paper, N is held in `param.N` and can be changed.

6.1. Thermo-viscous boundary layer drag and reflection coefficients

We model the viscous drag in the boundary layers inside the tube by using Stokes' solution [36, §347(7)]. In the frequency domain, the tangential wall shear stress, τ_w , acting on the air is:

$$\hat{\tau}_w = -\mu \left. \frac{\partial \hat{u}}{\partial y} \right|_w = \left(\frac{\omega \rho \mu}{2} \right)^{1/2} (1 + i) \hat{u} = (\omega \rho \mu)^{1/2} \sqrt{i} \hat{u} \quad (10)$$

This stress acts on the perimeter area $\pi D \delta x$, where δx is the length of an element in the network model. The pressure on either side of this element acts on a cross-sectional area $\pi D^2/4$. The viscous boundary layer drag is therefore modelled through a jump condition in the pressure (1) via a local feedback coefficient $k_{mu} = \delta x (4/D) (\omega \rho \mu)^{1/2} \sqrt{i}$.

Rayleigh [36, §348 – 350(16)] models the thermal drag by adding $(a/b - b/a)\sqrt{\alpha}$ to the viscosity, where $a \equiv \sqrt{\gamma p/\rho}$, $b \equiv \sqrt{p/\rho}$, and α is the thermal diffusivity. This simple model is not appropriate, however, if visco-thermal drag is modelled through *local* feedback because the thermal drag is proportional to the wall temperature gradient, which is largest at pressure antinodes, while the viscous drag is proportional to the wall velocity gradient, which is largest at velocity antinodes. Instead, here we model the thermal drag as local feedback from the acoustic temperature into the energy equation (2) using the method in [44, §X.D]. This gives $k_{ep} = \delta x (\gamma - 1) / (\rho c)^2 (4/D) (\omega \rho \mu)^{1/2} Pr^{-2/3} \sqrt{i}$. For

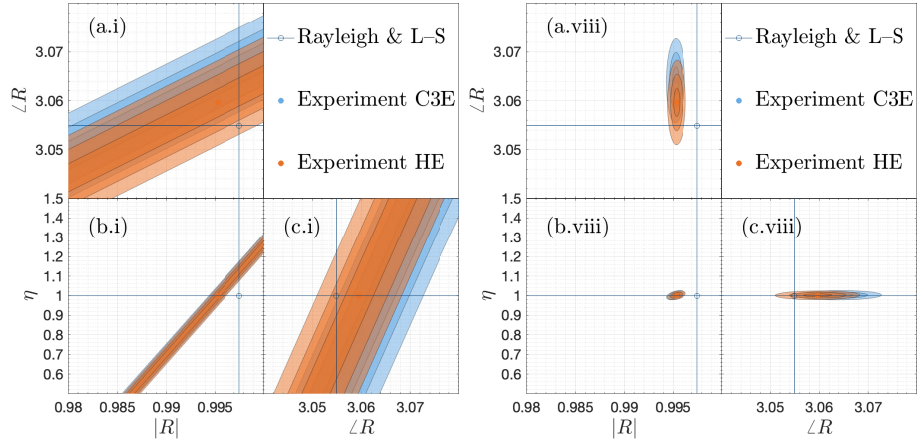


Figure 3: MAP Posterior probability distributions $P(\mathbf{a}_{\text{MP}}|D, \mathcal{H})$ for $\mathbf{a} = (|R|, \angle R, \eta)$ when assimilating data from the two empty tube experiments labelled E in Table 1. The contours show 1, 2, and 3 standard deviations from the assimilated expected values for (a) $\mathbf{a} = (|R|, \angle R)$, (b) $\mathbf{a} = (|R|, \eta)$, and (c) $\mathbf{a} = (\angle R, \eta)$. (\star .i) The data contains decay rate and frequency only, measured by the microphone at $x/L = 0.75$. (\star .viii) The data contains the decay rate, frequency, and relative pressures of all 8 microphones. This shows that the posterior probability distributions collapse only if the data contains sufficient information for that model. The values predicted by Rayleigh [36] and Levine-Schwinger [22] are shown for comparison.

the assimilation process, we multiply k_{mu} and k_{ep} by a real constant η , which would be unity if this model were perfect.

420 The ends of the tube are identical so, when the tube is empty, we assume that their reflection coefficients are the same: $R_u = R_d = R$. For both empty tube experiments (column E in Table 1), Levine-Schwinger’s (LS) calculations [22] give $|R|_{\text{LS}} = 0.9975$ and $\angle R_{\text{LS}} = 3.0550$. If we set $\eta = 1.0$ and assimilate R directly from the experimentally-measured growth rates and frequencies of
425 the two empty tube experiments we obtain $|R| = 0.9953$ and 0.9952 , and $\angle R = 3.0597$ and 3.0641 . Alternatively if we set $R = R_{\text{LS}}$ and assimilate η we obtain $\eta = 1.1235$ and $\eta = 1.1198$. It might be tempting to accept a 0.23% discrepancy in $|R|$ or a 12% discrepancy between the boundary layer model and experiment, but careful use of Bayesian inference enables us to improve on both values.

430 For demonstration, we will now infer R and η simultaneously, first when there is insufficient information in the data, and second when there is sufficient information. We set one standard deviation of the measurement uncertainties to be 0.1 rad s^{-1} for the decay rate and 1 rad s^{-1} for the frequency. For the prior expected values we set $R = R_{\text{LS}}$ and $\eta = 1.0$. We set large prior variances.

435 Figure 3(a) shows the first inference problem, in which we attempt to infer $|R|$, $\angle R$ and η from the decay rates and frequencies only. On the one hand, $|R|$ and η both strongly influence the decay rate but weakly influence the frequency. It should therefore be impossible to infer $|R|$ and η simultaneously because there is only one useful piece of information: the decay rate. On the other hand
 440 $\angle R$ strongly influences the frequency but weakly influences the decay rate. It should therefore be possible to infer $\angle R$ quite accurately because the frequency information is useful and is broadly unaffected by $|R|$ and η . Figure 3(a) shows the posterior expected values and co-variances between pairs of parameters, in the form of rings denoting 1, 2, and 3 standard deviations from the expected
 445 values. The bottom-left frame contains a long thin diagonal ellipse, showing as expected that the data can be explained by a wide range of $|R|$ and η , where the main source of uncertainty in one is the value of the other. The other frames contain fatter ellipses, showing as expected that the data can be explained by a rather small range of $\angle R$, where the main source of uncertainty is the measurement uncertainty rather than the values of $|R|$ and η .

450 Figure 3(b) shows the posterior expected values and covariances when we attempt to infer R and η from the decay rates, frequencies, and the relative amplitudes and phases of all 8 microphones. Now there is sufficient information in the data that it can only be explained by small ranges of $|R|$, $\angle R$, and η . We
 455 see that η is remarkably close to 1.0, indicating that Rayleigh's model performs well for this experiment, but that the inferred R deviates slightly from R_{LS} . This slight deviation is to be expected because our tube is not infinitely-thin and the exterior volume is not infinitely large, which are both ingredients of LS's model. We calculate k_{LS} such that $R = k_{\text{LS}} \times R_{\text{LS}}$ and hardwire this empirical
 460 coefficient into the model for R used in the rest of this paper.

Table 2: Description and performance of seven models for the drag and blockage caused by the heater prongs. Column 2 contains the number of parameters, \mathbf{a} , in each model. Columns 3 and 4 describe how the viscous prong drag, k_{mu_p} , and thermal prong drag, k_{ep_p} , are modelled: as the tube perimeter boundary layer drag multiplied by a real parameter, a complex parameter, or zero. Column 5 describes how the blockage of the prongs, k_{eu_h} , is modelled: as a real parameter or zero. Column 6 contains the log best fit likelihood (BFL) per datapoint, where $\text{BFL} = P(D|\mathbf{a}_{\text{MP}}, \mathcal{H}_i)$; a more positive $\log(\text{BFL})$ implies a better fit to the data. Column 7 contains the log Occam Factor (OF) per datapoint, where $\text{OF} = P(\mathbf{a}_{\text{MP}}|\mathcal{H}_i) (\det(\mathbf{A}/2\pi))^{-1/2}$; a more negative $\log(\text{OF})$ means that the parameter space has collapsed more when the data arrive. On the one hand this indicates that the model has learnt well from the data. On the other hand this can indicate that the prior parameter space is large and that the model has too many parameters. If the model is good then a small OF will be accompanied by a large BFL. Column 8 contains the log Marginal Likelihood (ML) per datapoint, where $\text{ML} = \text{BFL} \times \text{OF}$ (9); a more positive $\log(\text{ML})$ implies higher evidence for the model, given the data.

Model	Params	k_{mu_p}	k_{ep_p}	k_{eu_h}	$\log(\text{BFL})$	$\log(\text{OF})$	$\log(\text{ML})$
1	1	real		zero	-0.3635	-0.3561	-0.7197
2	2	complex		zero	-0.3575	-0.5567	-0.9141
3	2	real	real	zero	-0.3552	-0.4151	-0.7702
4	4	complex	complex	zero	+0.9452	-1.2319	-0.2867
5	1	zero	zero	real	-3.7064	-0.1736	-3.8799
6	2	real		real	+0.6705	-0.5767	+0.0938
7	3	real	real	real	+0.7010	-0.8191	-0.1181

6.2. Drag and blockage by the heater prongs

The heater is held by two 3 mm diameter threaded prongs and fed by two 4 mm diameter fabric-coated copper wires. These are inserted from the upstream end ($x = 0$). The black dots in Fig. 4 shows the measured growth rates and frequencies as the prongs are inserted without the heater attached. On physical grounds, we may suspect that these variations are caused by the extra thermo-viscous drag and blockage caused by the prongs and wires. We may not be sure, however, which physical effects need to be included in a low order model. On the one hand, we need to include all physical effects having a significant influence. On the other hand, if we include too many, it will be impossible

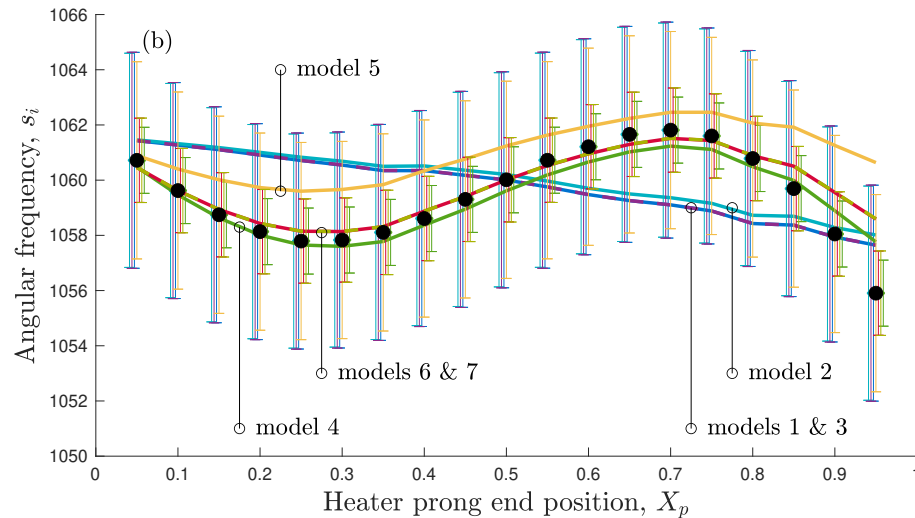
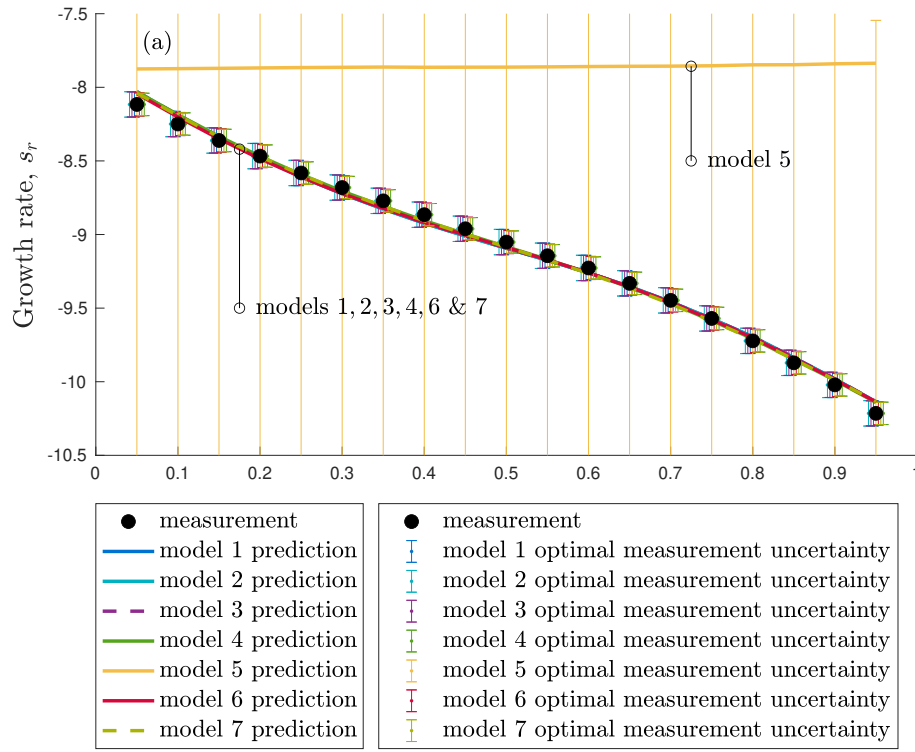


Figure 4: (a) Growth rate, s_r , and (b) angular frequency, s_i , as the heater prongs are inserted through the tube from $x = 0$ to $x = X_p$. The black dots show the experimental measurements. The coloured lines show the predictions from the models in Table 2 after assimilating the experimental data such that each model's parameters are at their MAP point. Coloured error bars: the measurement uncertainty that maximizes the peak of the MAP point for each model. For each measurement (black dot) the error bars are ordered from model 1 on the left to model 7 on the right.

to distinguish between them with the available data, as shown in Fig. 3(a). In this section we propose seven reasonable physics-based models, assimilate the data into those models, and use the marginal likelihood to identify the most likely model given the data.

475 We denote P_p as the sum of the perimeters of the prongs and heater wires, and P_t as the perimeter of the tube. At each axial location containing the prongs, their viscous drag and thermal drag are modelled by multiplying the viscous drag and thermal drag of the tube’s boundary layer at that location by P_p/P_t and then following the procedures outlined in the next paragraph. The
 480 effect of their blockage is modelled as an area change at the heater position x_h . We follow the analysis in [45, §4.4.2], which allows the effect of the area change to be expressed as local feedback from u into the energy equation: $k_{eu_h} = 1 - (\rho_1 A_1)/(\rho_2 A_2)$.

Model 1 multiplies the local viscous drag (k_{mu}) and local thermal drag (k_{ep})
 485 by the same real parameter to account for a possible change in the magnitude of drag between the tube walls and the prong walls due to the different surface roughness and materials. Model 2 multiplies the viscous drag and thermal drag by the same complex parameter to account for possible changes in the magnitude and phase of each type of drag. Model 3 multiplies the viscous
 490 drag by one real parameter and the thermal drag by another real parameter to account for possible independent changes in the magnitudes of the viscous and thermal drags. Model 4 multiplies the viscous drag by one complex parameter and the thermal drag by another complex parameter to account for possible independent changes in the magnitude and phases of the viscous and thermal
 495 drags. Model 5 multiplies k_{eu_h} by one real parameter to account for the blockage but contains no visco-thermal drag from the prongs. Model 6 combines models 1 and 5. Model 7 combines models 3 and 5. There is no limit to the number of models that can be proposed, but seven is enough for our purposes.

The parameters are assimilated allowing the measurement uncertainties to
 500 float in order to account for systematic error in the measurement and the model. This gives the fairest comparison of the marginal likelihoods because, if the mea-

surement uncertainties are fixed to their small measured values, then simple models (which tend to have larger model error) become overwhelmingly unlikely even if they miss just a few of the datapoints. Figure 4 shows the model
505 predictions vs experimental measurements, while Table 2 summarises the details of each model and shows their best fit likelihood (BFL), Occam Factor (OF), and marginal likelihood (ML). Models 1, 2, 3, and 5 have large negative values of $\log(\text{BFL})$ and $\log(\text{ML})$, meaning that they fit the data badly and are not supported by the data. Models 1, 2, and 3 (no blockage) cannot match the
510 experimental frequency, which implies that a successful model must include the prong blockage. Model 5 (no drag) cannot match the experimental growth rate, which implies that that a successful model must include the prong drag. Models 4 (complex thermo-viscous drag but no blockage) and model 6 and 7 (real drag and blockage) have large positive values of $\log(\text{BFL})$, meaning that they fit the
515 data well. Of these, model 6 has the largest value of ML because it achieves this good fit with just two parameters, while model 7 contains three parameters and model 4 contains four parameters. It is reassuring that the model with the most evidence (the highest marginal likelihood) is the simplest model that contains the expected physics. While a human may have been able to identify this model
520 by hand, this Bayesian framework provides a rigorous and quantifiable measure of its suitability compared with other plausible candidate models.

6.3. Drag and blockage by the heater itself

The experiments in rows C1 and C2 of Table 1 were performed on different days with the heater attached but switched off and with no thermocouples
525 present. The blue circles in Fig. 5 show the growth rates and frequencies as the heater is traversed through the tube. By comparing them with the purple squares, which are for the prong only experiments, it can be seen that the heater significantly changes the growth rate and frequency when at the ends of the tube and slightly changes the growth rate when at the centre of the tube. In this
530 section, we propose physical reasons for this and compare the evidence for three corresponding physical models.

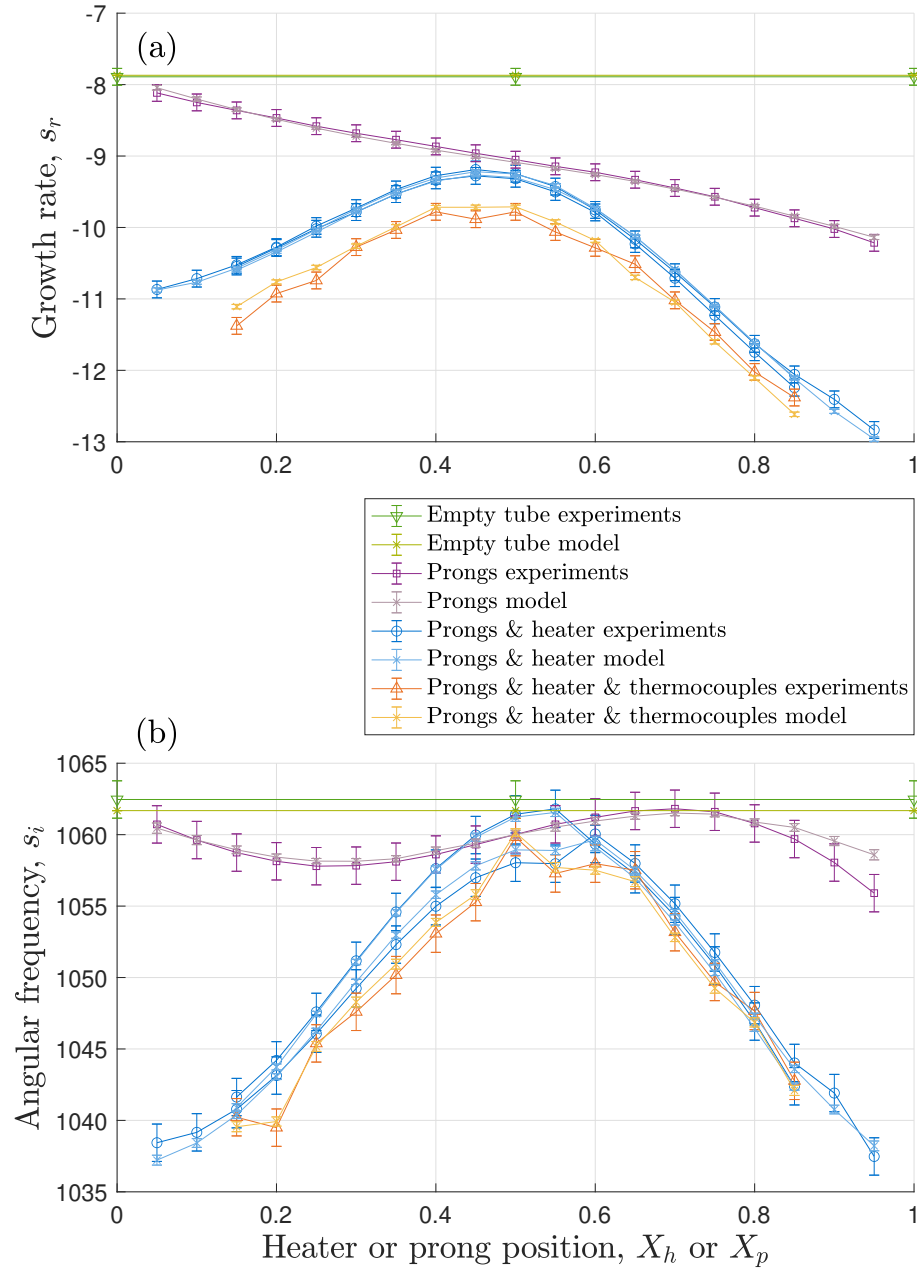


Figure 5: Experimental measurements and model predictions of (a) the growth rate, s_r , and (b) the angular frequency, s_i , for all the cold experiments (rows C1, C2, C3, and H at $Q = 0$ in Table 1). The model combines R and η from Fig. 3, model 6 from Table 2 for the prongs, model 3 from Table 3 for the heater, and thermo-viscous drag for the thermocouples. The experiments with the heater in place were performed on different days at different ambient temperatures, which accounts for the jagged lines that are particularly visible in the growth rate.

Table 3: As for Table 2 but assimilating models for the drag and blockage caused by the heater (Fig. 1). Column 3 describes whether k_{mu_h} is imaginary or complex: if imaginary, the heater is modelled as an orifice plate without viscous drag; if complex, the model also includes viscous drag. Column 4 describes whether k_{ep_h} is zero or complex: if complex, the model includes thermal drag.

Model	Params	k_{mu_h}	k_{ep_h}	log(BFL)	log(OF)	log(ML)
1	1	imag	none	-3.7472	-0.0949	-3.8420
2	2	complex	none	-1.6648	-0.1947	-1.8595
3	4	complex	complex	+0.4209	-0.4047	+0.0162

The heater (Fig. 1) consists of two orifice plates, one behind the other. From [45, §4.4.3], the pressure drop across a single orifice plate with length L and cross-sectional area A_d in a pipe of radius A equals $\Delta p = -\rho L_{\text{Leff}}(A/A_d)(du/dt)$,
 535 due to the inertia of the air in the orifice plate. The length L_{Leff} equals $L + 2\delta$ where δ is typically $(8/3\pi)(A_d/\pi)^{1/2}$. The local feedback mechanism is therefore from the velocity into the momentum equation. When converted to the frequency domain, the local feedback coefficient is $k_{mu_h} = \rho L_{\text{Leff}}(A/A_d)s$ where s is the complex angular frequency, whose imaginary part is much greater than
 540 its real part.

Nichrome wire is wound around each orifice plate. It is reasonable to suppose that the viscous drag from the wire is non-zero because momentum transfer from the wire relies on the same physical mechanism as heat transfer, and heat transfer will be crucial in later sections. The question is whether the data shows
 545 that the contributions to the viscous drag coefficient (k_{mu_h}) and thermal drag coefficient (k_{ep_h}) are negligible compared with the contribution to k_{mu_h} from the orifice. (Note that the orifice and the viscous drag both contribute to the same local feedback coefficient, although the orifice affects only the imaginary component.)

550 We evaluate three models. Model 1 models the orifice alone, by assimilating an imaginary value of k_{mu_h} and setting the other local feedback coefficients ($k_{mp_h}, k_{ep_h}, k_{eu_h}$) to zero. Model 2 is the same as model 1 but allows k_{mu_h} to

be complex, which models the orifice and the viscous drag, but not the thermal drag. Model 3 models the orifice and the wire’s visco-thermal drag by allowing
555 complex non-zero values of k_{mu_h} and k_{ep_h} . Table 3 shows the BFL, OF, and ML of all three models. Of these, the data supports model 3. Although not shown here, model 1 fails because it cannot model the influence of the heater on the growth rate at the ends of the tube, where the acoustic velocity is high and the viscous drag is important. Model 2 fails because it cannot model the
560 influence of the heater on the growth rate at the centre of the tube, where the acoustic temperature fluctuation is high and the thermal drag is important. Model 3 models the orifice blockage, the viscous drag, and the thermal drag. It is worth noting that the imaginary part of k_{mu_h} is six times greater than its real part, indicating that the orifice blockage is significantly more influential than
565 the viscous drag.

The experiments in row H0 of Table 1 were performed on different days with the heater attached but switched off and with thermocouples in place. The orange triangles in Fig. 5 show the growth rates and frequencies of these experiments. By comparing them with the blue circles, which are for the same
570 experiments without the thermocouples in place, it can be seen that the thermocouples significantly affect the growth rate and slightly affect the frequency. We model the thermocouples as identical local thermo-viscous drag elements and assimilate their complex visco-thermal drag coefficients k_{mu_t} and k_{ep_t} . Further details are not reported here but can be found in the Matlab code [20]. Figure 5
575 shows the experimental measurements and calibrated model predictions from all the cold experiments, using calibrated Levine-Schwinger for the reflection coefficients, Rayleigh’s model for the tube boundary layer, model 6 from Table 2 for the prongs, model 3 from Table 3 for the heater, and assimilated visco-thermal drag coefficients for the thermocouples.

580 6.4. Inferring the fluctuating heat release rate at the heater

Having carefully developed a quantitatively-accurate model of the elements of the Rijke tube from the cold experiments, we now assimilate k_{eu_h} , which

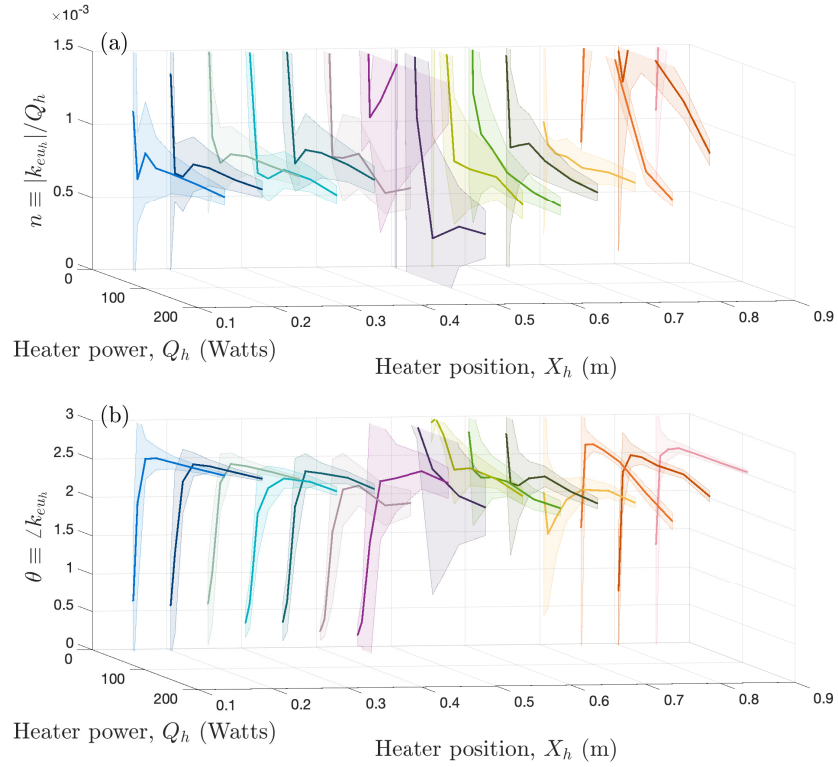


Figure 6: Local feedback coefficient from velocity fluctuations to heat release rate fluctuations at the heater, k_{eu_h} , inferred independently at each heater power and heater position. The solid lines show the expected values. The patches show ± 2 standard deviations. (a) $|k_{eu_h}|/Q_h$; (b) $\angle k_{eu_h}$. These values were inferred from the experiments in rows H of Table 1, with R and η from Fig. 3, model 6 from Table 2 for the prongs, model 3 from Table 3 for the heater, and thermo-viscous drag for the thermocouples.

is the local feedback from the velocity to the heat release rate at the heater. Figure 6 shows the expected values of $k_{eu_h} \pm 2$ standard deviations calculated independently at each heater power and heater position. The parameter k_{eu_h} is inferred from the thermoacoustic behaviour of the system. The thermoacoustic effect is strongest when the heater is placed around $X_h = 0.25$ and $X_h = 0.75$, and is weakest when placed around $X_h = 0.5$. The uncertainty in k_{eu_h} should therefore be smallest when the heater is placed around $X_h = 0.25$ and $X_h = 0.75$ and greatest when placed around $X_h = 0.5$. This can be observed clearly in Fig. 6. Away from $X_h = 0.5$ and for heater powers above 15 Watts, the values of $|k_{eu_h}|/Q_h$ and $\angle k_{eu_h}$ are, with little uncertainty, re-assuringly independent of the heater position and heater power. The biggest deviation is seen when the heater is near the top end of the tube ($X_h \sim 0.8$), and, by inspecting the two right-most temperature profiles in Fig. 2, we suspect this is due to a systematic error in the base flow calculation.

For comparison, Fig. 7 shows k_{eu_h} calculated from the same data, but without including models for the thermo-viscous drag and blockage of the heater and prongs. It is tempting to ignore the thermo-viscous drag and blockage of the heater and prongs, as in [11, 13, 38, 39] because they are difficult to model and seem *a priori* to have little influence. With these ignored, however, $|k_{eu_h}|/Q_h$ and $\angle k_{eu_h}$ seem to change significantly with heater power and heater position. Most worryingly, $\angle k_{eu_h}$ seems to change discontinuously around $X_h = 0.5$. There is, of course, no physical justification for this apparent dependence. It arises simply because k_{eu_h} is being used to accommodate deficiencies elsewhere in the model. A machine learning algorithm applied to a neural network would learn this dependence but, consequently, would only be able to interpolate between previous observations. The advantage of our physics-based Bayesian approach is that it reveals bad models, such as that used to create Fig. 7, and forces the researcher to develop good models, which tend to have simple physically-interpretable behaviour, such as that used to create figure 6.

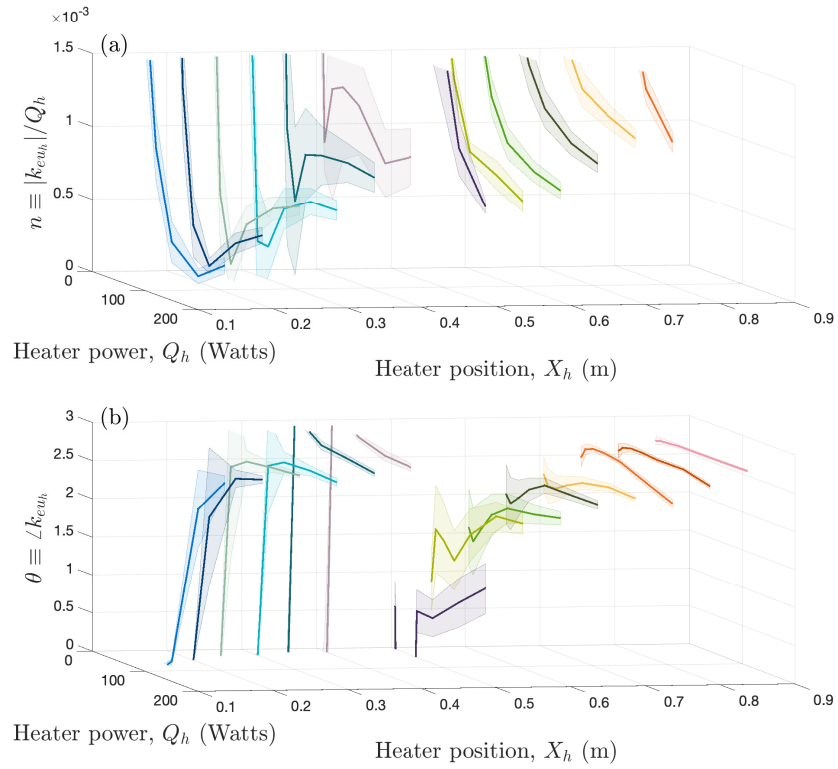


Figure 7: As for Fig. 6, with R from Levine-Schwinger [22], $\eta = 1$ in Rayleigh's model and neglecting blockage and thermo-viscous drag from the prongs and heater.

6.5. *Modelling the fluctuating heat release rate at the heater*

We now propose nine candidate models for the complex local feedback coefficient k_{eu_h} . These models are listed in Table 4, where k_i with numerical subscripts are the model parameters. Models 1, 2, 7, 8, & 9 model $|k_{eu_h}|$ as $k_1 \times Q_h$, taking inspiration from Fig. 6, in which this quantity is seen to be nearly uniform over the operating regime. Model 3 models $|k_{eu_h}|$ as $k_1 \times Q_h^{k_3}$, to see whether the added flexibility of the power law in Q_h increases the marginal likelihood of the model. Models 4, 5, & 6 model $|k_{eu_h}|$ as $k_1 \times Q_{\text{King}}$, where Q_{King} is calculated from King’s law [23, Eq(33)]:

$$Q_{\text{King}} = \frac{1}{(Re Pr \pi/2)^{-0.5} + 2} \frac{Q_h}{U_h} \quad (11)$$

Models 1, 3, 4, & 9 model $\angle k_{eu_h}$ as constant k_2 , taking inspiration from Fig. 6, in which this is nearly uniform over most of the operating regime. Models 2, 5, 7, & 8 model $\angle k_{eu_h}$ as $k_2 \times \text{Im}(s)$, which corresponds to a constant time delay model in which $\tau = k_2$. Model 6 models $\angle k_{eu_h}$ as $k_2 \times \text{Im}(s)\tau_L$ where $\tau_L \equiv 0.2d_{\text{wire}}/U_h$ is the time delay calculated by Lighthill [24]. If this is accurate then k_2 should assimilate to 1.

Models 1, 2, 3, 4, 5, & 6 assume that the thermo-viscous drag at the heater does not change when the heater element becomes hot. Model 7 allows the viscous drag to drift from its cold value in proportion to Q_h . Models 8 and 9 allow the viscous and thermal drags to drift from their cold values in proportion to Q_h . There is no limit to the number of models that can be proposed and the interested reader is encouraged to propose their own models in the Matlab code that accompanies this paper [20].

Of the models with constant thermo-viscous drag at the hot wire (models 1 to 6), model 2 has the highest marginal likelihood. The extra flexibility of model 3 turns out not to be beneficial. Modelling the heat release rate with King’s law is marginally less accurate, but this could be due to systematic errors in U_h calculated with the long timescale code in §4.1. The marginal likelihood increases significantly when the thermo-viscous drag at the hot wire is allowed to increase with the heater power in models 7, 8, & 9. This has little effect on the

Table 4: As for Table 2 but assimilating models for the hot heater. Column 2 contains the number, N , of parameters, which are labelled $k_1 \dots k_N$. Column 3 contains the model for $|k_{eu_h}|$, where Q_h is the heater power and Q_{King} comes from (11). Column 4 contains the model for $\angle k_{eu_h}$ where $\text{Im}(s)$ is the frequency and $\tau_L = 0.2d_{\text{wire}}/U_h$ is Lighthill's time delay. Columns 5 and 6 contain the viscous and thermal drag coefficients where models 7, 8 and 9 are allowed to deviate from their cold values, k_c .

Model	Params	$ k_{eu_h} $	$\angle k_{eu_h}$	k_{mu_h} ($k_c = \text{cold value}$)	k_{ep_h} ($k_c = \text{cold value}$)	$\log(\text{BFL})$	$\log(\text{OF})$	$\log(\text{ML})$
1	2	$k_1 \times Q_h$	k_2	k_c	k_c	-4.7055	-0.1016	-4.8071
2	2	$k_1 \times Q_h$	$k_2 \times \text{Im}(s)$	k_c	k_c	-4.5805	-0.1030	-4.6836
3	3	$k_1 \times Q_h^{k_3}$	k_2	k_c	k_c	-4.6477	-0.1385	-4.7862
4	2	$k_1 \times Q_{\text{King}}$	k_2	k_c	k_c	-4.7218	-0.1058	-4.8275
5	2	$k_1 \times Q_{\text{King}}$	$k_2 \times \text{Im}(s)$	k_c	k_c	-4.6956	-0.1072	-4.8028
6	2	$k_1 \times Q_{\text{King}}$	$k_2 \times \text{Im}(s)\tau_L$	k_c	k_c	-5.7721	-0.1019	-5.8741
7	4	$k_1 \times Q_h$	$k_2 \times \text{Im}(s)$	$k_c + (k_3 + ik_4) \times Q_h$	k_c	-3.4889	-0.2655	-3.7544
8	6	$k_1 \times Q_h$	$k_2 \times \text{Im}(s)$	$k_c + (k_3 + ik_4) \times Q_h$	$k_c + (k_5 + ik_6) \times Q_h$	-3.2512	-0.3964	-3.6476
9	6	$k_1 \times Q_h$	k_2	$k_c + (k_3 + ik_4) \times Q_h$	$k_c + (k_5 + ik_6) \times Q_h$	-3.6496	-0.3837	-4.0333

frequency but a significant effect on the growth rate when the heater is placed near the ends of the tube, where the acoustic velocity is greatest. Comparing models 2, 7 & 8 we see that the extra viscous drag is more influential than the
635 extra thermal drag, but that including both is best. Comparing models 1 & 2 and models 8 & 9 shows that there is more evidence for models with constant time delay τ than constant phase angle. Model 8 is the best model, but models 7 and 9 are also good.

640 Figure 8 compares the experimental measurements with the model predictions of a model similar to those used by Refs.[11, 13, 38, 39]. This model has Rayleigh's thermo-viscous boundary layers, Levine-Schwinger's reflection coefficients, and model 2 for the heat release rate k_{eu_h} , which was the best model when the thermo-viscous drag from the heater is not included. The most likely
645 coefficients of model 2 are assimilated from the data but the agreement is not good. This model is under-damped so, although it matches the frequency rea-

sonably well, it always over-predicts the growth rate. Indeed at high powers and $X_h \approx 0.25$, this model predicts a positive growth rate. It is acceptable for the model to predict a positive growth rate in general, because the model is physics-based and can therefore extrapolate beyond the experimental regime. In this case, however, the model has not extrapolated beyond the experimental regime, so the positive growth rates are revealing model error.

Figure 9 shows the model predictions when the researcher has noticed that the first model is under-damped and has chosen to model the extra damping as if it were all caused by the thermo-viscous boundary layer. This model therefore assimilates the coefficient η as well as k_{eu_h} . The assimilated constant η is 1.67, which corresponds to a 67% increase in the boundary layer drag. This seems to be a reasonable approach and the model is able to approximate the growth rates and frequencies over the operating range, but it is certainly not quantitatively accurate.

Figure 10 shows the predictions for a model that contains the components with the highest marginal likelihoods: model 8 from Table 4, R and η from Fig. 3, model 6 from Table 2 for the prongs, model 3 from Table 3 for the heater, and thermo-viscous drag for the thermocouples. This model is qualitatively and quantitatively accurate over the entire operating range, as desired.

6.6. Extrapolation

The process of selecting a model from a range of candidate models requires a large amount of data. The more parameters each model has, the more data is required in order to choose between them. Once a good model has been selected, however, it can be trained on sparse data. Figure 11 compares model predictions with experimental measurements when data from just eight operating points are assimilated into the model in Fig. 10. The results are almost indistinguishable from Fig. 10, for which 105 operating points were assimilated. We appreciate that this argument is somewhat circular: the model was chosen because it could explain all the data with few parameters, so it is unsurprising that data from a few points, when assimilated into the model, can extrapolate to the other

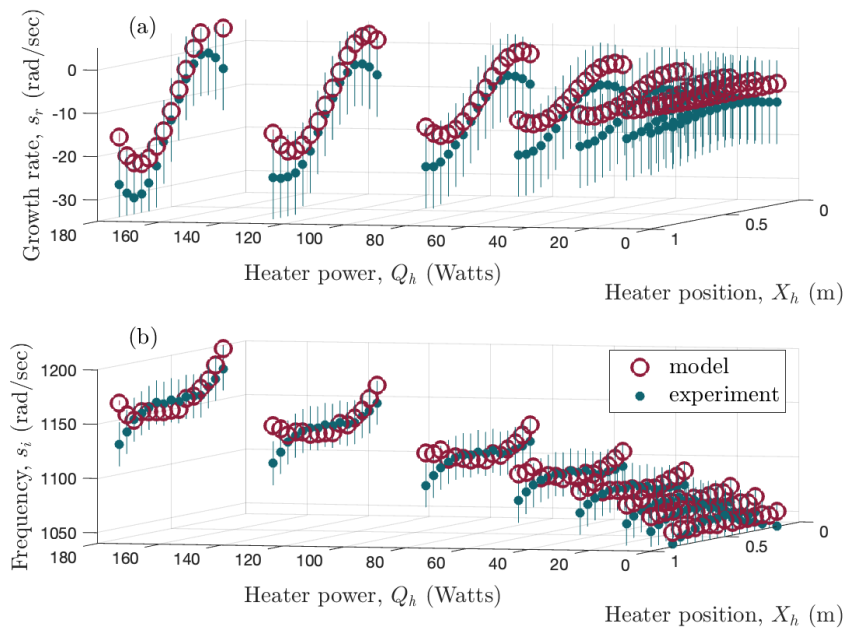


Figure 8: Model predictions of (a) growth rate and (b) frequency vs. experimental measurements assimilated with model 2 from Table 4 for the heat release rate, Rayleigh’s model for the thermo-viscous boundary layer, and Levine-Schwinger [22] for the reflection coefficients. The error bars on the experimental measurements show two standard deviations of the measurement uncertainty required to maximize the a posteriori likelihood, given this model.

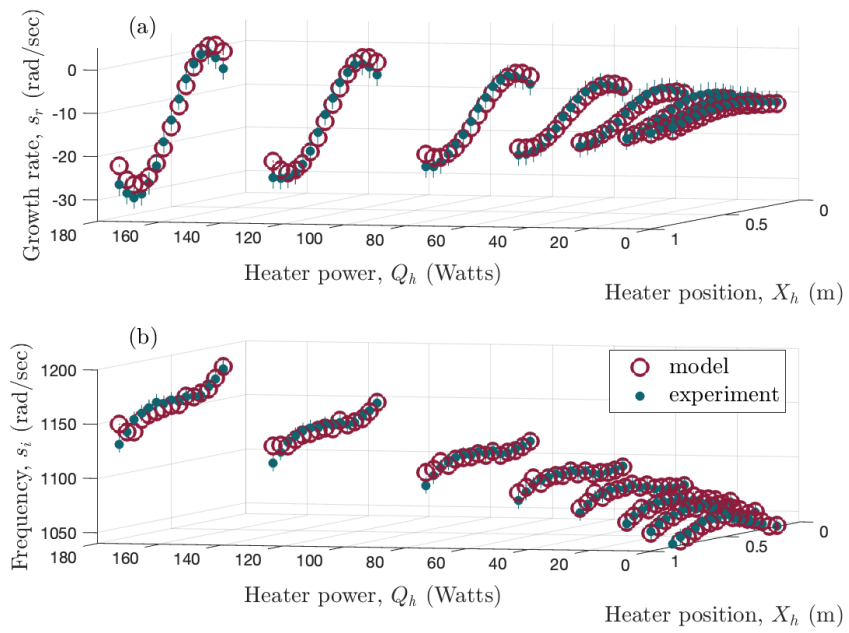


Figure 9: As for Fig. 8 but with Rayleigh's thermo-viscous dissipation multiplied by η , where η is assimilated (to 1.67) to maximize the MAP.

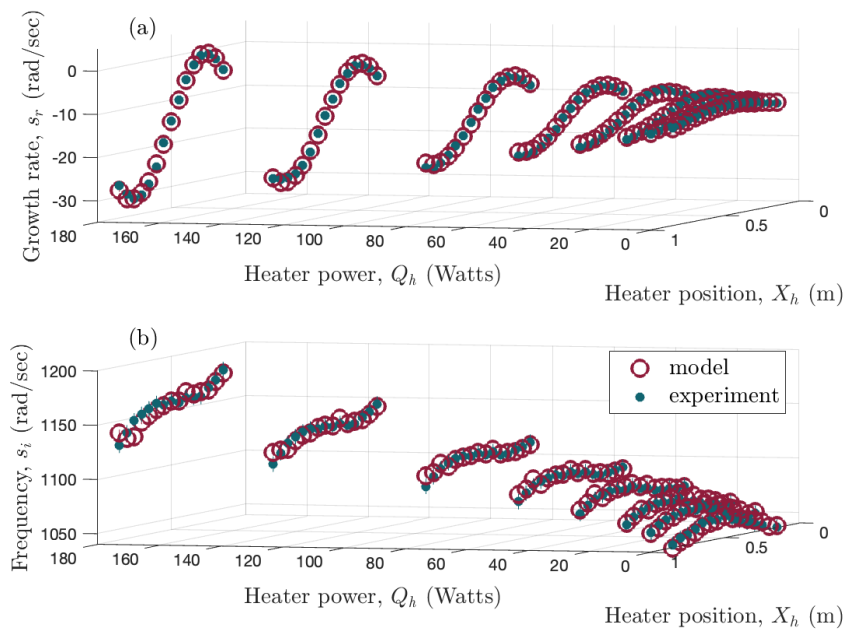


Figure 10: As for Fig. 8 but using the component models with the highest marginal likelihoods: model 8 from Table 4, R and η from Fig. 3, model 6 from Table 2 for the prongs, model 3 from Table 3 for the heater, and thermo-viscous drag for the thermocouples.

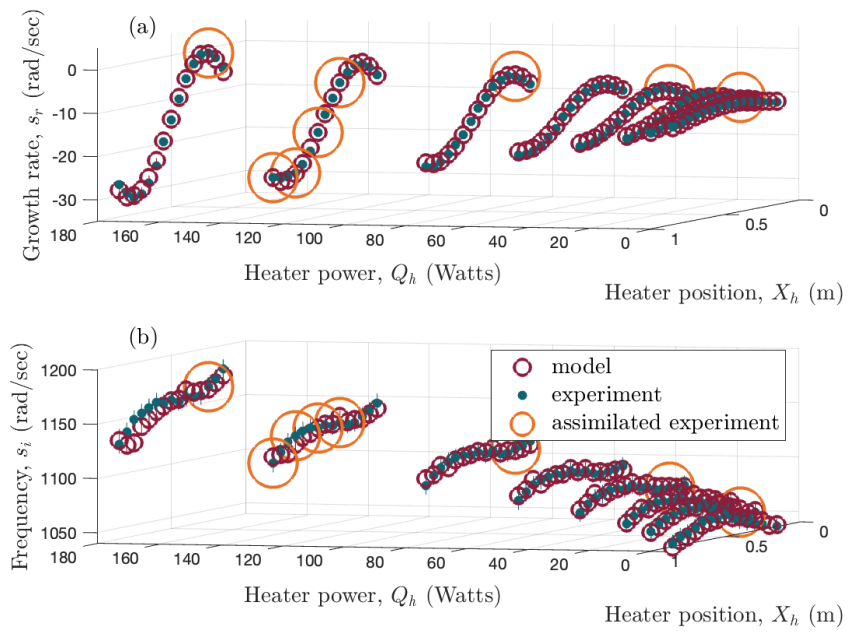


Figure 11: As for Fig. 10 but assimilating data from just eight of the 175 operating points.

points. Nevertheless, this feature shows why assimilating data into physics-based models with few parameters is more desirable than assimilating data into physics-agnostic models with many parameters: much less data is required.

680 *6.7. The required number of measurement devices*

Having shown that a model can be trained on measurements at just a few operating points, it is natural to ask how many measurement devices are required. The answer to this question depends on the stage of experimental design. When designing, setting up, and commissioning the rig, all information is
685 useful. At this stage, the biggest problem is systematic measurement error that the experimentalist has not yet discovered. One sign of unidentified systematic measurement error is that a high marginal likelihood can only be achieved with a model that seems unphysical, such as when one model parameter requires an unphysical dependence on another. For example, in earlier experiments on this
690 rig, our model would only fit the measurements well when the heater time delay was allowed to be a linear function of the heater position [42]. We investigated possible reasons for this with CFD simulations of the flow over the heater, but none was satisfactory. We concluded that the model parameters were shifting in order to absorb some as yet unidentified systematic measurement error. In order
695 to identify this systematic measurement error, different measurements were required. When we measured the sound speed directly with the microphones, we discovered that the systematic error was in the thermocouple measurements, probably due to radiation from the heater. When many different measurement devices are available, it is easier to identify and eliminate systematic measurement
700 error. In this case, we rejected all the thermocouple measurements except that of the ambient air.

At the next stage, the biggest problem is model error. The researcher may not know which physical phenomena are most influential in a given experiment, so has to select models by finding those with the highest marginal likelihood. As
705 mentioned in section 6.6, this requires a great deal of data, and every accurate measurement is useful. For example, in our early models, we did not include

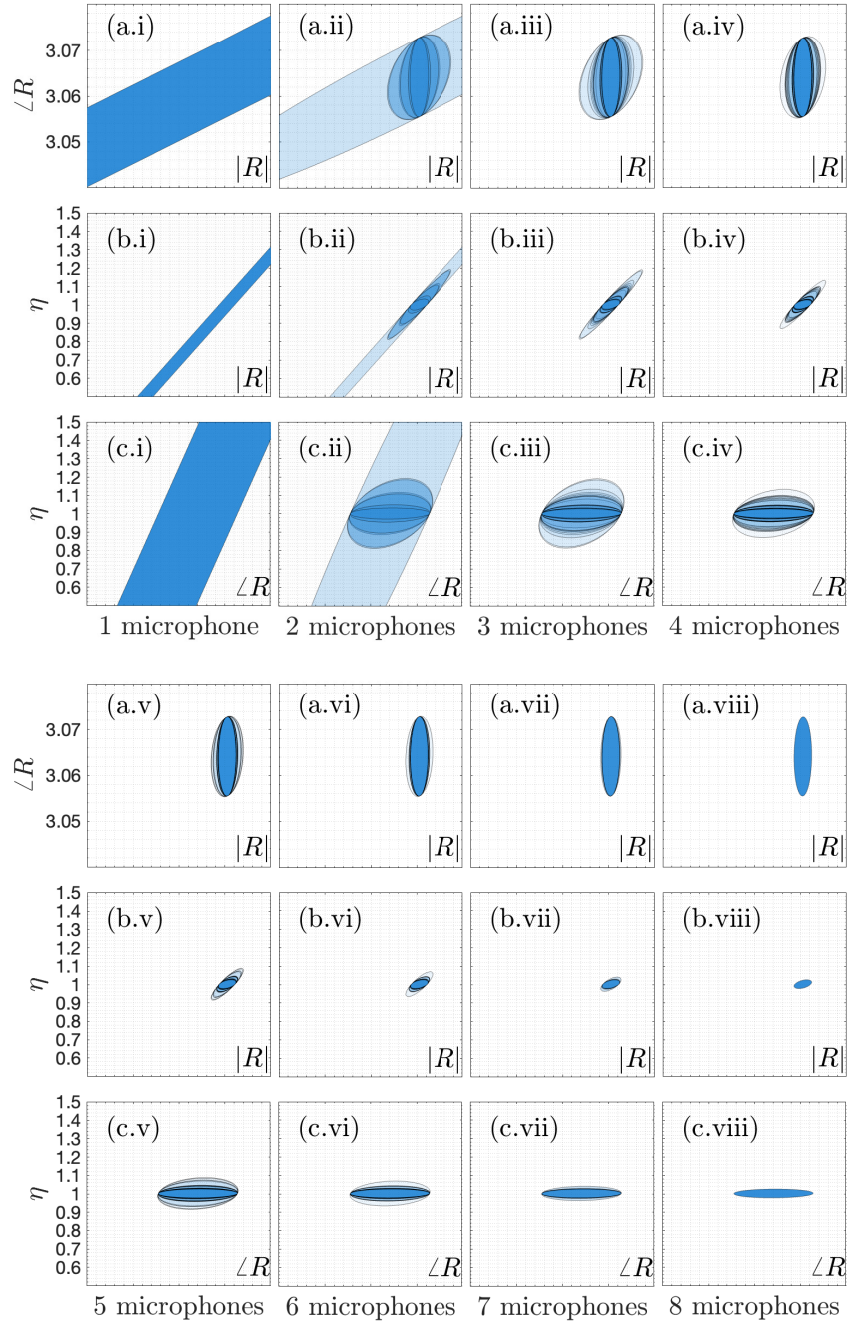


Figure 12: As for Fig. 3 but showing only the 3 standard deviation contour of the empty tube experiment labelled C3E in Table 1 as the number of assimilated microphones increases from 1 ($\star.i$) to 8 ($\star.viii$), where \star denotes (a) $\mathbf{a} = (|R|, \angle R)$, (b) $\mathbf{a} = (|R|, \eta)$, and (c) $\mathbf{a} = (\angle R, \eta)$. In each frame, contours are plotted for every 4th permutation of microphones that includes the reference microphone, which is placed at $x/L = 0.75$. There is therefore 1 combination of 1 and 8 microphones, 7 combinations of 2 and 7 microphones, 21 combinations of 3 and 6 microphones, and 35 combinations of 4 and 5 microphones. The uncertainties decrease as the number of microphones increases.

viscous and thermal drag from the tube boundary layers, assuming that their influence could be neglected or absorbed into the acoustic reflection coefficients. Because we had 8 microphones, however, we were able to show that their measurements could not be consistent with a model that excluded this visco-thermal drag but were highly consistent with a model containing this visco-thermal drag (Fig. 3b). If we had had just 2 microphones, we would not have been aware of this model error.

Once systematic measurement error and model error have been reduced as far as is practical, the number of required measurements depends only on the accuracy desired and the skill of the experimentalist. For example, Fig. 12 shows 3 standard deviations of the MAP posterior probability distributions of $\angle R$, $|R|$, and Rayleigh's parameter, η , when assimilating measurements from 1 to 8 microphones. These are the intermediate steps from Fig. 3(*.i) (1 microphone) to Fig. 3(*.viii) (8 microphones). In each figure, the microphone at $x/L = 0.75$ is the reference microphone, and all possible combinations of the other microphones are plotted on top of each other (e.g. there are 7 different configurations with 2 microphones, 21 with 3 microphones, 35 with 4 microphones etc.) With 1 microphone, measuring just the decay rate and frequency, it is impossible to disentangle $\angle R$, $|R|$, and η . With 2 microphones, the parameters' uncertainties depend strongly on which pair of microphones is used. A skilled experimentalist with reliable equipment can achieve nearly the same uncertainty with 2 microphones (the central ellipses in Fig. 12(*.ii)) as can be achieved with 8 microphones (the single ellipses in 12(*.viii)). As the number of microphones increases beyond 2, the uncertainty rapidly diminishes for all combinations of microphones. The extra microphones are therefore an excellent mitigation against imperfect experimental design.

In conclusion, two microphones are sufficient if they have minimal systematic error and if the model into which they are assimilated is physically accurate. More microphones are better, however, because this redundancy enables the experimentalist to identify systematic measurement error, to become aware of model error, and to accommodate microphone placements that, whether by

accident or design, are imperfect.

7. Conclusion

740 We have shown that Bayesian inference provides a powerful framework for
the assimilation of data into physics-based models. Using this framework we
have assembled a physics-based model of a thermoacoustic experiment, compo-
nent by component. For each component we have proposed several candidate
models and have then selected the component model with the highest marginal
745 likelihood, given the data. The selected component model is always the sim-
plest model that contains all the physics necessary to explain the data. The
final assembled model is therefore as small as possible, quantitatively accurate,
and physically interpretable. The model extrapolates successfully because it is
physics-based. Much data is required to select the model but, once selected,
750 little data is required to train it.

Computer-aided design optimization, whether gradient-based or not, re-
quires models to be quantitatively-accurate. While this is particularly clear
in thermoacoustics, where systems are notoriously sensitive to small changes
and therefore contain significant systematic error, this applies to all areas of
755 engineering. For the Rijke tube, we have successfully used experimental data
to create a quantitatively-accurate physics-based model, which is a significant
improvement on other attempts in the literature e.g. [11].

Our next steps are to repeat the experiments and data assimilation for a
flame in a Rijke tube with a more elaborate acoustic network. We will learn
760 on-the-fly by running the assimilation code alongside the experiments. This will
allow us to choose the next datapoint such that it contains as much information
as possible and thereby to reduce the total number of experiments. The methods
in this paper are applied most readily to physics-based low order models. They
extend, however, to CFD, which is the subject of current work in flow imaging
765 [46].

Finally, there is no limit to the number of physics-based candidate models

that can be proposed to fit the data. All the figures in this paper have been created with the Matlab code [20], which is well commented. The reader is invited to adapt this code by proposing their own models, and to improve on

770 ours.

CRedit author statement

Matthew Juniper: Conceptualization, Methodology, Software, Validation, Formal analysis, Data Curation, Writing – Original Draft, Writing – Review & Editing, Visualization, Supervision. Matthew Yoko: Investigation, Writing –
775 Review & Editing.

Acknowledgements

M. Yoko acknowledges financial support for his PhD from the Cambridge Trusts, the Skye Foundation, and the Oppenheimer Memorial Trust. M. Juniper acknowledges Shon Ffowcs-Williams for his characteristically frank opinion about Ph.D subjects. As an undergraduate M. Juniper had expressed an
780 interest in doing a Ph.D in Turbulence and asked Shon Ffowcs-Williams for advice. “There is a department up the road,” he said nodding towards the Department of Applied Maths and Theoretical Physics, “which has researched little else for seventy years. Do you know what they have achieved? ... What
785 makes you think you’d be any different?”

References

- [1] J. E. Ffowcs-Williams, Review Lecture: Anti-sound, Proc. Roy. Soc. A 395 (1984) 63–88. doi:10.1098/rspa.1984.0090.
 - [2] A. P. Dowling, A. S. Morgans, Feedback control of combustion oscillations, Ann. Rev. Fluid Mech. 37 (1) (2005) 151–182. doi:10.1146/annurev.fluid.36.050802.122038.
- 790

- [3] M. A. Heckl, Active control of the noise from a Rijke tube, *J. Sound Vib.* 124 (1) (1988) 117–133. doi:10.1016/S0022-460X(88)81408-1.
- [4] J. R. Seume, N. Vortmeyer, W. Krause, J. Hermann, C. Hantschk, P. Zangl,
795 D. Vortmeyer, A. Orthmann, Application of active combustion instability
control to a heavy duty gas turbine, *J. Eng. Gas Turb. Power* 120 (4) (1998)
721–726. doi:10.1115/1.2818459.
- [5] J. M. Cohen, N. M. Rey, C. A. Jacobson, T. J. Anderson, Active control of
combustion instability in a liquid-fueled low-NOx combustor, *J. Eng. Gas*
800 *Turb. Power* 121 (2) (1999) 281–284. doi:10.1115/1.2817118.
- [6] M. P. Juniper, R. I. Sujith, Sensitivity and nonlinearity of thermoacoustic
oscillations, *Ann. Rev. Fluid Mech.* 50 (2018) 661–689. doi:10.1146/
annurev-fluid-122316-045125.
- [7] J. C. Oefelein, V. Yang, Comprehensive review of liquid-propellant com-
805 bustion instabilities in F-1 engines, *J. Prop. Power* 9 (5) (1993) 657–677.
doi:10.2514/3.23674.
- [8] J. G. Aguilar, M. P. Juniper, Adjoint methods for elimination of ther-
moacoustic oscillations in a model annular combustor via small geome-
try modifications, in: *ASME Turbo Expo, 2018*, p.75692. doi:10.1115/
810 GT2018-75692.
- [9] J. Aguilar, M. P. Juniper, Thermoacoustic stabilization of a longitudinal
combustor using adjoint methods, *Phys. Rev. Fluids* 5 (8). doi:10.1103/
PhysRevFluids.5.083902.
- [10] H. C. Mongia, T. J. Held, G. C. Hsiao, R. P. Pandalai, Challenges and
815 progress in controlling dynamics in gas turbine combustors, *J. Prop. Power*
19 (5) (2003) 822–829. doi:10.2514/2.6197.
- [11] K. I. Matveev, Thermoacoustic instabilities in the Rijke tube: experiments
and modeling, Ph.D. thesis, California Institute of Technology (2003).

- [12] T. C. Lieuwen, A. Banaszuk, Background noise effects on combustor stability, *J. Prop. Power* 21 (1) (2005) 25–31. doi:10.2514/1.5549. 820
- [13] G. F. Carrier, The mechanics of the Rijke tube, *Quart. App. Math.* 12 (4) (1955) 383–395. doi:10.1090/qam/69698.
- [14] H. J. Merk, Analysis of heat-driven oscillations of gas flows (ii) on the mechanism of the Rijke tube phenomenon (1957). *App. Sci. Res. A.* 6 (1957) 402–420 doi:10.1007/bf03185045. 825
- [15] D. J. C. MacKay, Information theory, inference, and learning algorithms, C.U.P. 2003.
URL <http://www.inference.phy.cam.ac.uk/mackay/itila/>
- [16] H. Yu, Inverse problems in thermoacoustics, Ph.D. thesis, Cambridge 830 (2020).
- [17] L. Magri, M. P. Juniper, Sensitivity analysis of a time-delayed thermoacoustic system via an adjoint-based approach *J. Fluid Mech.* 719, 183–202, (2013) doi:10.1017/jfm.2012.639.
- [18] L. Magri, M. Bauerheim, M. P. Juniper, Stability analysis of thermoacoustic nonlinear eigenproblems in annular combustors. Part I Sensitivity, *J. Comp. Phys.* 325 (2016) 395–410. doi:10.1016/j.jcp.2016.07.032. 835
- [19] M. P. Juniper, M. Yoko, Data Assimilation with Laplace’s Methods in Thermoacoustics In: *Symposium on Thermoacoustics in Combustion*, 6–10 Sept 2021, Munich.
- [20] M. P. Juniper, M. Yoko, Code supporting ‘Generating a physics-based quantitatively-accurate model of an electrically-heated Rijke tube with Bayesian inference’, Univ. Cambridge (2022) doi:10.17863/CAM.84141. 840
- [21] G. Rigas, N. P. Jamieson, L. K. B. Li, M. P. Juniper, Experimental sensitivity analysis and control of thermoacoustic systems, *J. Fluid Mech.* 787 (2016) R1 1–11. doi:10.1017/jfm.2015.715. 845

- [22] H. Levine, J. Schwinger, On the radiation of sound from an unflanged circular pipe *Phys. Rev.* 73(1) (1948) 383–406. doi:10.1103/PhysRev.73.383.
- [23] L. V. King, On the convection of heat from small cylinders in a stream of fluid: determination of the convection constants of small platinum wires with applications to hot-wire anemometry *Phil. Trans. Royal Soc. A* 214 (1914) 373–432. doi:10.1098/rsta.1914.0023.
- [24] M. J. Lighthill, The response of laminar skin friction and heat transfer to fluctuations in the stream velocity, *Proc. Roy. Soc. A.* 224 (1156) (1954) 1–23. doi:10.1098/rspa.1954.0137.
- [25] A. Witte, W. Polifke, Dynamics of unsteady heat transfer in pulsating flow across a cylinder, *Int. J. Heat Mass Transf.* 109 (2017) 1111–1131. doi:10.1016/j.ijheatmasstransfer.2017.02.072.
- [26] K. T. Feldman, Review of the literature on Rijke thermoacoustic phenomena, *J. Sound Vib.* 7 (1) (1968) 83–89. doi:10.1016/0022-460X(68)90159-4.
- [27] R. L. Raun, M. W. Beckstead, J. C. Finlinson, K. P. Brooks, A review of Rijke tubes, Rijke burners and related devices, *Prog. Energy Comb. Sci.* 19 (1993) 313–364. doi:10.1016/0360-1285(93)90007-2.
- [28] G. Bisio, G. Rubatto, Sondhauss and Rijke oscillations - thermodynamic analysis, possible applications and analogies, *Energy* 24 (2) (1999) 117–131. doi:10.1016/S0360-5442(98)00090-5.
- [29] J. W. S. B. Rayleigh, The explanation of certain acoustical phenomena, *Nature* 18 (455) (1878) 319–321. doi:10.1038/018319a0.
- [30] B.-T. Chu, On the energy transfer to small disturbances in fluid flow (Part I), *Acta Mechanica* 1 (3) (1965) 215–234. doi:10.1007/BF01387235.

- [31] B. J. Bayly, Onset and equilibration of oscillations in general Rijke devices, *J. Acoust. Soc. Amer.* 79 (3) (1986) 846–851. doi:10.1121/1.393475.
- [32] S. Mariappan, R. I. Sujith, Modelling nonlinear thermoacoustic instability
875 in an electrically heated Rijke tube, *J. Fluid Mech.* 680 (2011) 511–533.
doi:10.1017/jfm.2011.176.
- [33] F. Selimefendigil, S. Föller, W. Polifke, Nonlinear identification of unsteady
heat transfer of a cylinder in pulsating cross flow, *Comp. Fluids* 53 (1)
(2012) 1–14. doi:10.1016/j.compfluid.2011.08.012.
- 880 [34] M. Heckl, Non-linear acoustic effect in the Rijke tube, *Acustica* 72 (1990)
63–71.
- [35] A. Witte, Dynamics of unsteady heat transfer and skin friction in pulsating
flow across a cylinder, Ph.D. thesis, T. U. Munich (2018).
- [36] J. W. S. Rayleigh, *The theory of sound volume 2*, Macmillan 1896 and
885 Dover 1945
- [37] L. E. Kinsler, A. B. Coppens, J. V. Sanders, A. R. Frey *Fundamentals of
acoustics*, Wiley 1982
- [38] C. Nicoli, P. Pelce, One-dimensional model for the Rijke tube, *J. Fluid
Mech.* 202 (83) (1989) 83–96. doi:10.1017/S0022112089001102.
- 890 [39] G. Gelbert, J. P. Moeck, C. O. Paschereit, R. King, Feedback control of
unstable thermoacoustic modes in an annular Rijke tube, *Cont. Eng. Pract.*
20 (8) (2012) 770–782. doi:10.1016/j.conengprac.2012.03.016.
- [40] N. Jamieson, G. Rigas, M. Juniper, Experimental sensitivity analysis via
a secondary heat source in an oscillating thermoacoustic system, *Int. J.
895 Spray Comb. Dyn.* 9 (4) 230–240. doi:10.1177/1756827717696325.
- [41] F. Garita, H. Yu, M. Juniper, Assimilation of experimental data to create
a quantitatively accurate reduced-order thermoacoustic model, *J. Eng. Gas
Turb. Power* 143 (2). doi:10.1115/1.4048569.

- [42] F. Garita, Physics-based statistical learning in thermoacoustics, Ph.D thesis, U. Cambridge (2021).
900
- [43] G. Evensen Data Assimilation, Springer (2009) doi:10.1007/978-3-642-03711-5.
- [44] M. Juniper, Sensitivity analysis of thermoacoustic instability with adjoint Helmholtz solvers, Phys. Rev. Fluids 3 110509 (2018). doi:10.1103/PhysRevFluids.3.110509.
905
- [45] S. W. Rienstra, A. Hirschberg An Introduction to Acoustics, Eindhoven U. Tech. (2006).
URL <https://research.tue.nl/en/publications/an-introduction-to-acoustics>
- 910 [46] A. Kontogiannis, S. V. Elgersma, A. J. Sederman, M. P. Juniper Joint reconstruction and segmentation of noisy velocity images as an inverse Navier–Stokes problem, J. Fluid Mech. *in print* (2022)