Material parameter estimation and hypothesis testing on a 1D viscoelastic stenosis model: Methodology

H.T. Banks, Shuhua Hu, and Z.R. Kenz

Center for Research in Scientific Computation

Center for Quantitative Sciences in Biomedicine

North Carolina State University, Raleigh, NC 27695-8212

Carola Kruse, Simon Shaw, and J.R. Whiteman

BICOM, Brunel University, Uxbridge, UB8 3PH, England

M.P. Brewin and S.E. Greenwald

Blizard Institute, Barts and the London School of Medicine and Dentistry,

Queen Mary, University of London, England

M.J. Birch

Clinical Physics, Barts and the London National Health Service Trust, England

October 31, 2012

Abstract

Non-invasive detection, localization and characterization of an arterial stenosis (a blockage or partial blockage in the artery) continues to be an important problem in medicine. Partial blockage stenoses are known to generate disturbances in blood flow which generate shear waves in the chest cavity. We examine a one-dimensional viscoelastic model that incorporates Kelvin-Voigt damping and internal variables, and develop a proof-of-concept methodology using simulated data. We first develop an estimation procedure for the material parameters. We use this procedure to determine confidence intervals for the estimated parameters, which indicates the efficacy of finding parameter estimates in practice. Confidence intervals are computed using asymptotic error theory as well as bootstrapping. We then develop a model comparison test to be used in determining if a particular data set came from a low input amplitude or a high input amplitude; this we anticipate will aid in determining when stenosis is present. These two thrusts together will serve as the methodological basis for our continuing analysis using experimental data currently being collected.

Mathematics Subject Classification: 62F12; 62F40; 65M32; 74D05.

Key words: viscoelastic model; sensitivity analysis; inverse problem; asymptotic theory; bootstrapping; model selection.

1 Introduction

Current methods to detect and locate arterial stenoses (blocked arteries) include somewhat invasive angiography and expensive CT scans. Neither procedure is particularly easy to administer, while the CT scan can localize hard plaques but not soft plaques [28]. Accordingly, there is interest in examining other methods to determine the existence and location of stenosed vessels. Previous work [2, 11, 12, 13, 15, 28, 34] focused on developing a sensor device to be used with a physical model of a chest cavity, and then developing a mathematical model to describe the medium in which a stenosis-generated acoustic signal is propagated to the chest surface. After an interregnum of roughly five years between that earlier work and our current efforts, we have returned to the early ideas and have reformulated the problem to some extent. This is motivated by companion experiments being conducted with novel acoustic phantoms built at Queen Mary, University of London (QMUL) and Barts & London NHS Trust (BLT) in England. Our viscoelastic model will be quite general, incorporating Kelvin-Voigt damping and internal variables in a hysteresis formulation, so as to provide maximum flexibility in these early stages of model development and analysis.

In this work, we continue to use mathematical modeling techniques in order to non-invasively determine the existence and location of any arterial stenoses, ideally through sensors placed only on the surface of the chest. To this end, we have developed novel experiments to produce one-dimensional pressure wave data that can be fit using the viscoelastic model developed here. While this work will be informed partly by test experimental devices which have been built, we note that all the data employed in this initial mathematical formulation will be simulated data. Broadly, then, the work here is intended to serve as a proof-of-concept of our mathematical and statistical inverse problem methodology, using simulated data very similar to the experimental data we will subsequently use with these methods along with incorporation of standard viscoelastic models (see, e.g., [6]).

Our goals for these initial efforts are twofold. We first focus on *developing methodologies for determining material parameters and analyzing data using a viscoelastic model*, as well as also *quantifying the uncertainty in the estimation procedure* through both bootstrapping and asymptotic error theories. As part of this latter goal, we also conduct model comparison testing to examine the viability of determining if data originated from a low-amplitude traction (e.g., resulting from normal blood flow) or high-amplitude traction (e.g., resulting from abnormal blood flow caused by a stenosis). This notion uses and enhances previous work (see, e.g., [2, 29, 30, 31, 32, 35, 36, 37]) which discussed the compression and shear waves which result from a stenosed vessel and some methods for measuring these waves, in particular the shear waves which experience slower transmission than the pressure waves. An ultimate goal of our wider research project will be a methodology to decide if a vessel is stenosed or not, and if so, possibly the extent and location of the stenosis. However, at this point we are still in the process of carrying out experiments to determine these differences in either test devices or live subjects. Thus, in these early efforts we make the (very) tentative assumption of representing the difference between normal vessels and stenosed vessels as a comparison between low and high shear input amplitudes, leaving the specifics of the actual system inputs to future work. We believe that this a reasonable first approximation to the anticipated experimental data. Overall, then, these two thrusts of material parameter estimation and model comparison tests represent a proof-of-concept for our future data-driven inverse problem efforts.

2 Viscoelastic model

In this section, we will examine a simplified one-dimensional viscoelastic model for an agar phantom, as pictured in Figure 1. This configuration is an approximation to the novel experimental devices

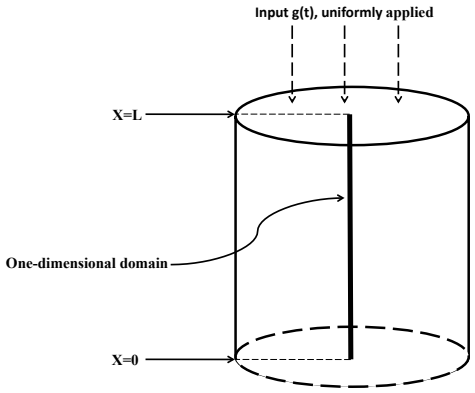

Figure 1: Setup of agar phantom, with sample one-dimensional domain denoted.

we are using at QMUL to gather experimental data. Development of general viscoelastic equations can be found throughout the literature; in particular, one may refer to [6] as a source of the model components discussed in this current work. As in the example in [6], we make simplifying assumptions that will result in a one-dimensional wave equation. If we assume a uniform force applied along the top of the phantom and radial symmetry within the phantom (in part to closely match the symmetrically constructed phantoms used at QMUL), then we can simplify the cylindrical physical domain to a one-dimensional domain and to finding the function $u(x,t)$ which represents the material response to, in this case, an applied stenosis-generated like force.

We will use a general acoustic pressure viscoelastic wave equation on a one-dimensional domain $\Omega = [0, L]$. For the purposes of our initial investigation here, all parameters will be considered constant (i.e., a homogeneous medium). This is not necessary but significantly simplifies our initial computations in the methodology development. Choosing a material initially at rest with a reflecting boundary at $x = 0$, an applied force $g(t)$ at the $x = L$ boundary, and no additional forcing terms we obtain the system for the displacement u given by

$$\begin{aligned} &\rho u_{tt} - \sigma_x = 0 \\ &u(0,t) = 0, \quad \sigma(L,t) = -g(t), \\ &u(x,0) = 0, \quad u_t(x,0) = 0. \end{aligned} \tag{1}$$

Here the stress $\sigma(x,t)$ is assumed to be described by

$$\sigma(x,t) = E_1 u_{xt}(x,t) + E_0 \int_0^t P(t-s)\frac{d}{ds}u_x(x,s)ds, \tag{2}$$

where the $E_1 u_{xt}$ term is the so-called Kelvin-Voigt damping, and P is some stress relaxation function. Development of this model is described in [6], as well as in standard viscoelastic theory [18, 20, 21, 22, 25, 27]. It is worth noting that the linear system (1) with (2) is found to give a reasonable approximation to the experimental data provided by QMUL. Thus, this linear model is deemed adequate for our investigation for the current experimental setting.

2.1 Existence and uniqueness of weak solution to system (1) with (2)

Let $\mathbb{H} = \mathcal{L}^2(0, L)$, $\mathbb{V} = \{\phi \mid \phi \in \mathcal{H}^1(0, L), \phi(0) = 0\}$, and $\mathbb{V}^*$ denote the topological dual space of $\mathbb{V}$. If we identify $\mathbb{H}$ with its topological dual $\mathbb{H}^*$, then $\mathbb{V} \hookrightarrow \mathbb{H} = \mathbb{H}^* \hookrightarrow \mathbb{V}^*$ is a Gelfand triple [1, 38]. Throughout this presentation $\langle \cdot, \cdot \rangle$ denotes the inner product in $\mathbb{H}$, and $\langle \cdot, \cdot \rangle_{\mathbb{V}^*, \mathbb{V}}$ represents the duality paring between $\mathbb{V}^*$ and $\mathbb{V}$ (again see [1, 38] for details).

Let $\mathcal{C}_w(0, T; \mathbb{V})$ denote the set of weakly continuous functions in $\mathbb{V}$ on $[0, T]$, and $\mathcal{L}_T = \{v : [0, T] \to \mathbb{H} \mid v \in \mathcal{C}_w(0, T; \mathbb{V}) \cap \mathcal{L}^2(0, T; \mathbb{V})$ and $v_t \in \mathcal{C}_w(0, T; \mathbb{H}) \cap \mathcal{L}^2(0, T; \mathbb{V})\}$. The weak solution for system (1) with (2) is defined in the following way.

Definition 2.1. *We say that $u \in \mathcal{L}_T$ is a weak solution of the system* (1) *and* (2) *if it satisfies*

$$\begin{aligned} &\rho\langle u_t(t), \eta_t(t)\rangle - \rho \int_0^t \langle u_s(s), \eta_s(s)\rangle ds + \int_0^t g(s)\eta(L, s)ds + E_1 \int_0^t \langle u_{sx}(s), \eta_x(s)\rangle ds \\ &+E_0 \int_0^t \left\langle \int_0^s P(s-\xi)\frac{d}{d\xi}u_x(\xi)d\xi, \eta_x(s) \right\rangle ds = 0 \end{aligned} \tag{3}$$

for any $t \in [0, T]$ and $\eta \in \mathcal{L}_T$. Here and elsewhere in this section $u(t)$ and $\eta(t)$ denote the functions $u(\cdot, t)$ and $\eta(\cdot, t)$, respectively.

As in [1, 13, 14, 38] we remark that this notation of the weak solution for system (1)-(2) agrees with the usual one in that it yields $u_{tt} \in \mathcal{L}^2(0, T; \mathbb{V}^*)$ with equation (1) holding in the sense of $\mathcal{L}^2(0, T; \mathbb{V}^*)$. To ensure the existence and uniqueness of a weak solution to system (1) with (2), we make the following assumptions on the force function g and stress relaxation function P:

(A1) $g \in L^2(0, T)$.

(A2) The function P is differentiable with respect to $t \in \mathbb{R}^+$, and there exist constants c_1 and c_2 such that $|P(t)| \le c_1$ and $|\dot{P}(t)| \le c_2$ for all $t \in \mathbb{R}^+$, where $\dot{P}$ denotes the differentiation of P with respect to t.

Theorem 2.2. *Under assumptions (A1) and (A2), the system* (1) *with* (2) *has a unique weak solution on any finite interval $[0, T]$.*

We remark that system (1) with (2) is a special case of the viscoelastic model presented in [13] if we change the left boundary condition $u(0, t) = 0$ in (1) to $\sigma(0, t) = g_l(t)$ with $g_l(t)$ being some force function and change the right boundary condition to be stress free (i.e., $g \equiv 0$). Specifically, the stress-strain relationship in [13] is described by the following nonlinear form

$$\sigma(x, t) = E_1 u_{xt}(x, t) + \int_0^t \Gamma(t - s; \mathcal{F})\frac{d}{ds}\sigma^e(u_x(x, s))ds. \tag{4}$$

Here σ^e is the so-called elastic response function, and Γ is defined by

$$\Gamma(t;\mathcal{F}) = \int_{\mathcal{T}} \gamma(t,\tau) d\mathcal{F}(\tau),$$

where $\mathcal{T} \subset (0,\infty)$, γ is a function of t and relaxation time τ, and $\mathcal{F}$ is a probability distribution function of relaxation time τ. The existence and uniqueness of weak solution for this general viscoelastic model was given in [13], and the continuous dependence of the weak solution on the probability distribution function $\mathcal{F}$ was also given in [13] under a Prohorov metric framework [1, 3, 9, 13, 33] on the space of probability distributions.

Note that assumptions (A1) and (A2) conform with the ones made in [13]. Hence, the arguments for Theorem 2.2 are similar to those given in [13, 14]. Hence, we only sketch the ideas, referring readers to [1, 13, 14, 38] for further details. Let $\{\psi_j\}_{j=1}^{\infty}$ be any linearly independent total subset of $\mathbb{V}$. We define the Galerkin approximation $u^m(t) = \sum_{j=1}^{m} \beta_j^m(t)\psi_j$ as the unique solution of

$$\rho\langle u_{tt}^m, \psi_j\rangle_{\mathbb{V}^*,\mathbb{V}} + E_1\langle u_{tx}^m, \psi_j'\rangle + E_0\left\langle \int_0^t P(t-s)\frac{d}{ds}u_x^m(s)ds, \psi_j'\right\rangle + g(t)\psi_j(L) = 0, \quad j = 1,2,\ldots,m,$$

on the interval $[0,T]$, where ψ_j' denotes the derivative of ψ_j with respect to x, $j = 1,2,\ldots,m$. We can argue that $\{u^m\}$ and $\{u_t^m\}$ are bounded uniformly in $\mathcal{L}^2(0,T;\mathbb{V})$. Then by the Banach-Alaoglu theorem we know that there exists a function $u \in \mathcal{L}^2(0,T;\mathbb{V})$ such that

$$u^m \to u \text{ weakly in } \mathcal{L}^2(0,T;\mathbb{V}),$$
$$u_t^m \to u_t \text{ weakly in } \mathcal{L}^2(0,T;\mathbb{V}).$$

In addition, the following convergence results can be proven by using the Ascoli-Arzela theorem (e.g., see [23, Theorem 3.6.4]) and Aubin's lemma (e.g., see [19, Lemma 8.4])

$$u^m \to u \text{ weakly in } \mathbb{V} \text{ uniformly in } t \in [0,T], \text{ i.e., } u^m \to u \text{ in } \mathcal{C}_w(0,T;\mathbb{V}),$$
$$u_t^m \to u_t \text{ weakly in } \mathbb{H} \text{ uniformly in } t \in [0,T], \text{ i.e., } u_t^m \to u_t \text{ in } \mathcal{C}_w(0,T;\mathbb{H}),$$
$$u_t^m \to u_t \text{ in } \mathcal{L}^2(0,T;\mathbb{H}).$$

We can then show that $u_x^m(t) \to u_x(t)$ in $\mathbb{H}$. Based on these convergence results, u can be easily shown to be a weak solution of system (1)-(2). The uniqueness of the weak solution can be established by using a standard technique which demonstrates that the difference between any two possible solutions must be zero [38].

2.2 System (1)-(2) with a special form of stress relaxation function

Note that the origin for time is assumed at the beginning of motion and loading. Hence, we can rewrite (2) into the following form

$$\sigma(x,t) = E_1 u_{xt}(x,t) + E_0\left(P(0)u_x(x,t) - \int_0^t \left(\frac{d}{ds}P(t-s)\right)u_x(x,s)ds\right). \tag{5}$$

For the remainder of the paper, the form of the stress relaxation function $P(t)$ is assumed to be a Prony series

$$P(t) = p_0 + \sum_{j=1}^{N_p} p_j e^{-t/\tau_j}, \tag{6}$$

where all the p_j are nonnegative numbers and the τ_j values are positive, and with N_p being a positive integer. This series is based on the assumption that relaxation in a viscoelastic material can be well represented by a discrete number of relaxation times τ_j. Without loss of generality, we will also enforce $P(0) = 1$. A result of this constraint is that $\sum_{j=0}^{N_p} p_j = 1$. It is worth noting here that this special form of P satisfies assumption (A2). Hence, system (1) and (5) with P given by (6) also has a unique weak solution on any finite time interval $[0, T]$.

If we replace $\frac{d}{ds}P(t-s)$ in (5) with the s-derivative of the Prony series at $t-s$, we obtain

$$\sigma(x,t) = E_1 u_{xt}(x,t) + E_0 \left(u_x(x,t) - \sum_{j=1}^{N_p} \int_0^t \frac{p_j}{\tau_j} e^{-(t-s)/\tau_j} u_x(x,s) ds \right).$$

We can reformulate the integrals related to each internal variable as differential equations which we can solve simultaneously with the main system (1). To this end, we define the "internal variables"

$$\epsilon^j = \int_0^t \frac{p_j}{\tau_j} e^{-(t-s)/\tau_j} u_x(s) ds.$$

Then the time derivative of ϵ^j is given by

$$\epsilon_t^j = \frac{p_j}{\tau_j} u_x(t) - \frac{1}{\tau_j} \int_0^t \frac{p_j}{\tau_j} e^{-(t-s)/\tau_j} u_x(s) ds.$$

Relating ϵ^j and ϵ_t^j allows us to model the internal variables dynamically as

$$\begin{aligned} &\tau_j \epsilon_t^j + \epsilon^j = p_j u_x, \\ &\epsilon^j(0) = 0. \end{aligned} \tag{7}$$

for $j = 1, 2, \ldots, N_p$. We can then write the overall stress-strain relationship (5) as

$$\sigma \;=\; E_1 u_{xt} + E_0 \left(u_x - \sum_{j=1}^{N_p} \epsilon^j \right). \tag{8}$$

Note that even though p_0 is an element in the Prony series for $P(t)$, once the series is substituted in the model the constant p_0 no longer appears. However, p_0 is still present in the sum-to-one constraint on *all* p_i values, but we can easily work with the alternate constraint that the remaining p_j terms must satisfy $\sum_{j=1}^{N_p} p_j \leq 1$.

The damping and internal variables provide us the future flexibility to match the model to data from the experimental devices, and also present an interesting question of identifiability of the damping and internal variable parameters which we will later discuss in depth. Note also that the authors in [2, 28, 34] give computational results (using essentially an equivalent model from a slightly different conceptual formulation involving a distribution of relaxation mechanisms-see (4)) showing that discrete relaxation times can model well the viscoelastic material responses of the type we consider in this work (namely, attempting to approximate the response of biological soft tissue as characterized in [2, 24]). In fact, in previous work no more than two discrete relaxation times were used, which has informed our decision to allow a maximum of two relaxation times.

Since the ultimate goal of the wider research project will be examining the traction into the chest cavity that results from a healthy artery experiencing a heartbeat as compared with an artery containing a stenosis, our nonzero boundary input $g(t)$ will here be represented by an approximation to a pulse traction. In order to ensure a smooth, compactly supported input, we implement the input function as a Van Bladel function which is a good approximation to expected perturbation inputs to our system [2, 36] . This smoothness is useful in order to get the maximum benefit from using high order numerics. The function used is

$$g(t) = \begin{cases} A \cdot \exp\left(\dfrac{|ab|}{t(t+a-b)}\right) & \text{if } t \in (0, b-a), \\ 0 & \text{otherwise,} \end{cases} \tag{9}$$

where A is some positive constant, a and b are some constants with $b > a$.

Parameter values used in simulations Motivated by the experimental data to which we intend to apply this methodology, we choose values for the system parameters which simulate low-amplitude (on the order of $0.1mm$) oscillatory motions. For data generation, we will use two internal variables ($N_p = 2$). The weights p_i for our two relaxation time model will be fixed as given below. The baseline material parameter values chosen for this work are as follows:

$$\begin{gathered} E_0 = 2.2 \times 10^5 \text{ Pa}, \quad E_1 = 40 \text{ Pa} \cdot \text{s}, \quad \rho = 1010 \text{ kg/m}^3, \quad L = 0.053 \text{ m} \\ \tau_1 = 0.05 \text{ s}, \quad \tau_2 = 10 \text{ s}, \quad p_1 = 0.3, \quad p_2 = 0.55. \end{gathered} \tag{10}$$

Note that the density $\rho = 1010$ kg/m^3 is the true density of the agar gel that is used in the medium for our experiments at QMUL, and $L = 0.053$ m is the true height of the phantom. These are parameter values which are directly taken to approximate the experimental device. The values for E_0 and E_1 and for the relaxation times are physically reasonable parameters based on a perusal of the viscoelastic materials research literature and are also informed by our early experiments with the agar gels. In the Van Bladel function, the values of a and b have an effect on pulse width as well as the amplitude, while the value of A only has an effect on the amplitude. Their baseline values are chosen as follows:

$$a = 6 \times 10^{-3}, \quad b = 20 \times 10^{-3}, \quad A = 6 \times 10^3, \tag{11}$$

where the values a and b allow for an effective pulse application time of 14 ms.

Direct problem Altogether, this description (1) along with (7)-(9) encompasses the one dimensional model for the displacement $u(x,t)$ that will be studied in both major sections of this work. In other words, system (1) along with (7)-(9) with given parameter values is the *direct or forward problem*. Throughout this work, we solve (1), (7)-(9) for $u(x,t)$ with given parameter values using a spectral continuous finite element method in space (Gauss-Lobatto nodes) and a discontinuous Galerkin method in time. The numerical scheme is specially tailored to allow for high order space-time discretization in order to control dispersion errors and will be documented fully in [26] along with its convergence properties.

3 Estimation of material parameters

In this section, we examine an inverse problem methodology for estimating material parameters (and thus gain a sense for our ability to characterize an individual's material properties) with given simulated observations of displacement at the $x = L$ position (where the simulated data is generated under various measurement noise conditions, see details blow). In addition to determining an *estimate* for material parameters, we also need to determine our *confidence in the estimation procedure*. To this end, we will compare two techniques for determining confidence intervals, specifically the asymptotic theory discussed in [5, 16] versus using bootstrapping as discussed in [5].

In reality, one will obtain a set of experimental data and then one needs to determine how many (if any) relaxation times are required to represent well the data. Thus, we will want to compare the performance of our estimation procedures in three models. In each model, we will always estimate E_0 and E_1 (assuming given values for ρ and L in (10) along with given values of a, b and A in (11)), but we will vary the number of relaxation times incorporated into the model. The three models to be used in inverse problems are as follows:

1. For a model with no relaxation times, we do not include any τ_i or corresponding p_i in the model. Thus, we estimate only $\bar{\theta} = (E_0, E_1)^T$.

2. In the case with one relaxation time, we incorporate a single internal variable (i.e., $N_p = 1$). For this case, $\bar{\theta} = (E_0, E_1, \tau_1)^T$ will be estimated, but the corresponding material weight p_1 is fixed to be $p_1 = 0.3$.

3. For the case of two relaxation times (i.e., $N_p = 2$), we will estimate $\bar{\theta} = (E_0, E_1, \tau_1, \tau_2)^T$ with the corresponding material weights fixed to be the values in (10), that is, $p_1 = 0.3$ and $p_2 = 0.55$.

Note that for the models with one and two relaxation times the corresponding weights p_i are fixed. Though in reality one would certainly need to estimate the weights p_i, we take the liberty here of assuming them to be known so we can focus on the general methodology and in particular the reliability in identifying the relaxation times.

Considering this set of models will allow us to follow what one would consider in practice, examining the results of adding/subtracting model features. It is worth noting here that for this particular set of models, the one-relaxation-time model (the second case) is not a special case of the two-relaxation-time model (the third case) as the material weight p_2 in the two-relaxation-time model is fixed, and that the zero-relaxation-time model (the first case) is not a special case of

the one-time-relaxation model as the material weight p_1 in the one-relaxation-time model is fixed. However, if we allow the corresponding material weights to be free (i.e., to be estimated along with relaxation times), then the zero-relaxation-time model is indeed a special case of the one-relaxation-time model, and the one-relaxation-time model is a special case of the two-relaxation-time model. We will therefore use the sensitivity equations and parameter estimation results as well as model selection criterion to suggest the number of relaxation times needed in practice.

3.1 Sensitivity of model output with respect to material parameters

Before discussing simulated data and actually solving the inverse problems, we wish to complete some analysis on the model around the "true" material parameter values in (10). We found through numerical simulations that changes in E_0 have a significant effect on the oscillation frequency of model output $u(L,t)$ as well as a minor effect on its peak heights, and that increases in damping, E_1, lead to the expected effects that the energy dissipates faster in the material (so the oscillation peak heights become smaller and the material experiences fewer small oscillations at later simulation times). We also found that relaxation times can allow the model flexibility in matching the periodic local "peaks" and "troughs" in the oscillation part of model output. However, the scale of changes they induce in the model output is minor (compared to the scale of the model output). For more information on the general effects of changing material parameters on the model output, the interested readers can refer to [7, Section 2.1].

In order to further quantify the model response to changes in parameters around the baseline values of (10), we will examine the sensitivity of the model output $u(L,t)$ with respect to material parameters. Note that since the values of parameters are on such a varying scale, we will actually work with the log-scaled versions of the material parameters we are attempting to estimate. In other words, if $\bar{\theta} = (E_0, E_1, \tau_1, \tau_2)^T$ is the vector of parameters to be estimated, we define $\theta = \log_{10}(\bar{\theta})$.

Sensitivity analysis has been widely used in inverse problem investigations (e.g., see [16] and the references therein for details) to identify the model parameters and/or initial conditions to which the model outputs are most sensitive and for which one can readily construct confidence intervals when they are estimated (i.e., which are the most reliably estimated values). To compute the sensitivity of the model output to each parameter, one needs to find sensitivity equations which describe the time evolution of the partial derivatives of the model state with respect to each parameter. Sensitivity equations in terms of the non-log-scaled parameters $\bar{\theta}$ are derived in [7, Appendix A].

We can use the sensitivity of model output to the non-log-scaled parameters to find the sensitivity of model output with respect to the log-scaled parameters, which will be of interest here. Using the chain rule, we find that

$$\frac{\partial u(L,t;10^{\theta})}{\partial \theta_i} = \bar{\theta}_i \ln(10) \frac{\partial u(L,t;\bar{\theta})}{\partial \bar{\theta}_i},$$

where θ_i and $\bar{\theta}_i$ are the ith elements of θ and $\bar{\theta}$, respectively.

The sensitivities of model output with respect to parameters $(\log_{10}(E_0), \log_{10}(E_1), \log_{10}(\tau_1), \log_{10}(\tau_2))$ are depicted in Figure 2. From this figure we see that model output is most sensitive to $\log_{10}(E_0)$, sensitive to $\log_{10}(E_1)$, less sensitive to $\log_{10}(\tau_1)$, and only minimally sensitive to $\log_{10}(\tau_2)$. The most interesting feature related to our study is the fact that the scale of sensitivity of model output to the first relaxation time is on the order of 10^{-5} whereas the sensitivity of model output to

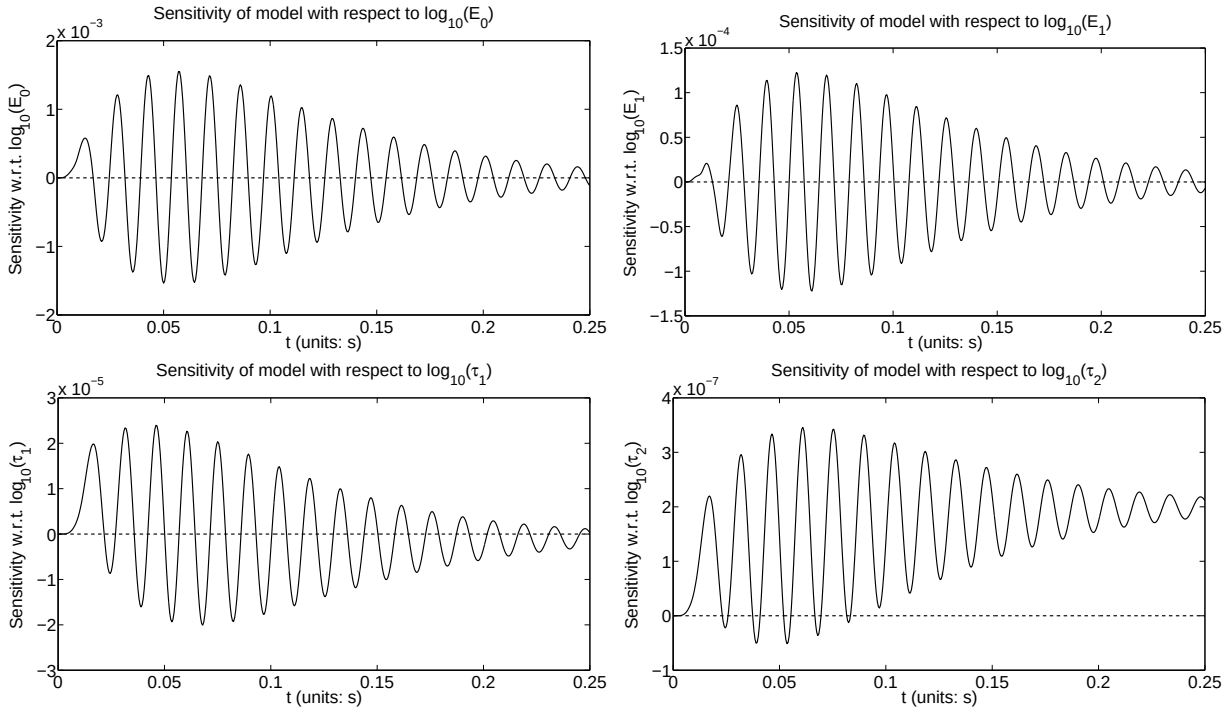

Figure 2: (upper left pane) Sensitivity of model output with respect to $\log_{10}(E_0)$; (upper right pane) Sensitivity of model output with respect to $\log_{10}(E_1)$; (bottom left pane) Sensitivity of model output with respect to $\log_{10}(\tau_1)$; and (bottom right pane) Sensitivity of model output with respect to $\log_{10}(\tau_2)$. All sensitivities are around the baseline parameters (10) and (11).

the second relaxation time is roughly two orders of magnitude smaller on the order of 10^{-7}. We will later see that, while we have difficulty estimating both relaxation times due to the small changes they induce in the model solution (as previously discussed), we especially have difficulty obtaining a reasonable estimate for τ_2 because the model is so much less sensitive to the second relaxation time than to the first. In addition, we observe from Figure 2 that at later times the model output is not particularly sensitive to all the material parameters except the second relaxation time. This further indicates that we may have trouble in estimating the second relaxation time with high additive noise data (which will be discussed later).

Armed with our knowledge of sensitivities of model output with respect to the material parameters around the true parameter values (10), and our knowledge of effects on the model solution of changing the parameters, we next describe the generation of our simulated data and discuss solution of the inverse problem.

3.2 Statistical model and inverse problem

We will work with simulated data for various noise levels generated at position $x = L$, namely data u_j corresponding to the model solution $u(L, t_j)$ at measurement time points t_j, $j = 0, 1, \ldots, n-1$. Then the *statistical model* (a model used to describe the observation process) is assumed to take

the following form

$$U_j = u\left(L, t_j; 10^{\theta_0}\right) + \mathcal{E}_j, \quad j = 0, 1, \ldots, n-1, \tag{12}$$

where $u(L, t_j; 10^{\theta_0})$ is the solution to (1) along with (7)-(9) at $x = L$ with a given set of true material parameter values θ_0 and the values of the rest of parameters given in (10) and (11). Here $\mathcal{E}_j$ denotes the measurement error (a random variable) at measurement time point t_j, $j = 0, 1, \ldots, n-1$. (It is worth noting that U_j, $j = 0, 1, 2, \ldots, n-1$, are also random variables due to the randomness of measurement errors.) For the current proof of concept discussion, we will assume the measurement errors $\mathcal{E}_j$, $j = 0, 1, \ldots, n-1$, are independent and identically distributed with mean zero ($E(\mathcal{E}_j) = 0$) and constant variance $\text{var}(\mathcal{E}_j) = \sigma_0^2$. We thus are assuming absolute additive error; this is necessary in order to use the hypothesis testing methodology later, and is reasonable as an initial error model for our proof of concept investigations. We do not make further assumptions about the distributions of the $\mathcal{E}_j$ in order to carry out the inverse problem methodology or the asymptotic analysis below. (To apply the AIC comparison methodology below there is the tacit assumption of normality on the $\mathcal{E}_j$.)

Under these assumptions for the measurement errors in the statistical model (12), the estimator $\hat{\Theta}$ of θ can be obtained by using the ordinary least squares method

$$\hat{\Theta} = \arg\min_{\theta \in Q} \sum_{j=0}^{n-1} \left[U_j - u(L, t_j; 10^{\theta})\right]^2, \tag{13}$$

where $Q \subset \mathbb{R}^{\kappa}$ is some viable admissible parameter set, assumed compact in $\mathbb{R}^{\kappa}$ with κ being the number of parameters requiring estimation. Thus, $\hat{\Theta}$ can be viewed as a minimizer that minimizes the distance between the data and the model where all observations are treated as of equal importance. Note that under different error assumptions, one would need to modify the cost function in (13) (a topic discussed in [16]) for an appropriate asymptotic parameter distribution theory to be valid.

Since $\hat{\Theta}$ is a random variable (inherited from the fact that U_j are random variables), we can define its corresponding realizations $\hat{\theta}$ by minimizing the cost function

$$\mathcal{J}(\theta) = \sum_{j=0}^{n-1} \left[u_j - u(L, t_j; 10^{\theta})\right]^2$$

over the set Q. That is, $\hat{\theta}$ is obtained by solving the following inverse problem

$$\hat{\theta} = \arg\min_{\theta \in Q} \sum_{j=0}^{n-1} \left[u_j - u(L, t_j; 10^{\theta})\right]^2, \tag{14}$$

Here u_j is a corresponding realization of U_j, and it is given by

$$u_j = u(L, t_j; 10^{\theta_0}) + \epsilon_j, \quad j = 0, 1, \ldots, n-1. \tag{15}$$

with ϵ_j being realizations of $\mathcal{E}_j$, $j = 0, 1, \ldots, n-1$. Note that the model solution $u(x, t)$ is continuously dependent on the model parameter θ. Hence, $\mathcal{J}$ is a continuous function of θ. Since

Q is assumed to be compact, the inverse problem (14) has a solution. Since for our studies we will be interested in perturbations around nominal values 10^{θ_0} of parameters and the corresponding solutions, our test problems will not in general suffer from serious ill-posedness and some type of stability or regularization techniques (Tikhonov regularization, regularization by discretization, etc., [10]) are not required for our studies. This will not necessarily be the case when using the inverse problem methods studied here with experimental data.

Estimating material parameters $\hat{\theta}$ from given sets of data with different noise levels, as well as *quantifying uncertainty* in our estimates, will be the key focus of our work in this section. We use the values for E_0, E_1, τ_1, and τ_2 in (10) in their log scaled form as the true values θ_0 used to simulate data. That is,

$$\theta_0 = (5.3424, 1.6021, -1.3010, 1)^T = (\log_{10}(2.2 \times 10^5), \log_{10}(40), \log_{10}(.05), \log_{10}(10))^T.$$

As previously discussed, we will find parameter estimates for models with zero, one, and two relaxation times in the model itself (and thus the number of parameters estimated changes). In all cases, the parameters belong to a viable compact set Q with the upper and lower bounds on parameters being taken (in educated guesses) as $\theta_{lb} = (-15, -15, -15, -15)^T$, $\theta_{ub} = (7.3010, 2.3010, 2, 2)^T$ for two relaxation times estimation, $\theta_{lb} = (-15, -15, -15)^T$, $\theta_{ub} = (7.3010, 2.3010, 2)^T$ for the one relaxation time estimation, and $\theta_{lb} = (-15, -15)^T$, $\theta_{ub} = (7.3010, 2.3010)^T$ for the no relaxation time estimation.

3.2.1 Data generation

We will simulate data using two relaxation times (and a question of interest later will be how many of those relaxation times we can recover) with the values of parameters given in (10) and (11). The measurement time points are taken at $t_j = 0.001j$, $j = 0, 1, \ldots, 250$. Thus, there are a total of $n = 251$ data points. We note that noiseless data has maximum amplitude on the order of 10^{-4} (depicted by the solid line in Figure 3), which was again motivated by the anticipated scale of results from the experimental device. This level informs the magnitude we choose for the additive noise. We represent a "low" noise level with $\sigma_0^2 = 5 \times 10^{-6}$, a "medium" noise level by $\sigma_0^2 = 10 \times 10^{-6}$, and a "high" noise level by taking $\sigma_0^2 = 20 \times 10^{-6}$. In Figure 3, we show plots corresponding to

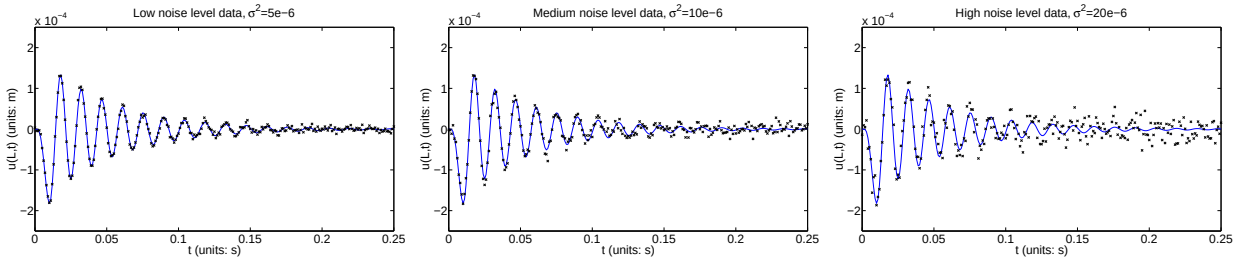

Figure 3: Simulated noisy data around the true parameter values. (left pane) Low noise level. (middle pane) Medium noise level. (right pane) High noise level.

the three levels of noisy simulated data against the system dynamics corresponding to the true parameters θ_0. Noise is assumed absolute for our initial investigations (though we may ultimately

need to explore relative noise once an error model is developed for our experimental data), and is added according to the error model (12). Low noise results in data mostly along the trajectory of the true model. Medium noise begins to obfuscate the later-time oscillations which have lost much of their earlier energy. High noise significantly affects the level of peaks and troughs from $t = 0.05$ forward. We thus obtain a series of increasingly difficult problems in obtaining material parameter estimates, though entirely expected since higher noise tends to significantly affect data features and presents a more difficult parameter estimation problem.

3.2.2 Parameter estimates obtained by using different optimization routines

In this section, we discuss different options for the optimization routine used to solve the inverse problem (14), and begin to gain a sense of the robustness of parameter estimation with respect to the optimization routine. Note that we expect to have some difficulty in relaxation time estimation, based on our earlier discussion on the model response to changes in relaxation times as well as model sensitivities. We do expect to obtain more accurate estimates for E_1, and very good estimates for E_0. To begin this discussion, we will examine parameter estimates for a model which incorporates two relaxation times.

The optimization routines we compare are all built-in Matlab routines. We use `fmincon` with active-set optimization, which treats the optimization as constrained nonlinear programming with our cost function $\mathcal{J}(\theta)$. We also examine the use of `lsqnonlin`, which is designed for nonlinear least squares data-fitting problems; our cost function is exactly the form of a nonlinear least squares function. We test both the Levenburg-Marquardt (LM) option and the trust-region-reflective (TRR) option. Note that the Levenburg-Marquardt algorithm does *not* allow bound constraints; we tried the routine out of curiosity, to see if it would produce unrealistic estimates of any parameters (it does at high noise levels).

Results from optimizing for θ by using different optimization routines are shown in Table 1. All optimization runs used the initial guess

$$\theta_{init} = (\log_{10}(1.8 \times 10^5), \log_{10}(60), \log_{10}(0.5), \log_{10}(20))^T = (5.2553, 1.7782, -0.3010, 1.3010)^T.$$

This table includes the parameter estimates $\hat{\theta}$, computation time (CPU) in seconds for that particular optimization run, and the residual sum of squares (RSS) defined as

$$\text{RSS} = \sum_{j=0}^{n-1} \left[u_j - u(L, t_j; 10^{\hat{\theta}}) \right]^2.$$

Overall, we observe from Table 1 that the routines do a good job of estimating E_0 (as we expected). The `lsqnonlin` routines tend to better estimate E_1. As for relaxation times, we begin to see a major flaw in use of the `fmincon` routine. It does not seem particularly sensitive to the relaxation times, and the resulting estimates of the relaxation times stay near the initial guess. The `fmincon` routine produced similar non-responsive results for different initial guesses. The `lsqnonlin` routines estimate both the relaxation times well in the presence of low noise. At medium noise, the routines estimate τ_1 well but not τ_2. At high noise, relaxation time estimation is poor. This will be quantified further in the following sections on error analysis.

Table 1: Estimation of material parameters at low, medium and high noise levels: Comparison between optimization routines (`lsq-TRR`=`lsqnonlin` with trust-region-reflective option, `lsq-LM`=`lsqnonlin` with Levenburg-Marquardt option).

Noise level	Routines	Estimated parameter values $\hat{\theta}$	CPU (s)	RSS
Low	`fmincon`:	$(5.3422, 1.6581, -0.3000, 1.3012)^T$	194.16	6.7618×10^{-9}
	`lsq-TRR`:	$(5.3425, 1.6046, -1.2297, 1.0046)^T$	347.09	6.2458×10^{-9}
	`lsq-LM`:	$(5.3425, 1.6044, -1.2309, 1.0316)^T$	613.47	6.2458×10^{-9}
Medium	`fmincon`:	$(5.3430, 1.6583, -0.2998, 1.3012)^T$	203.75	2.4435×10^{-8}
	`lsq-TRR`:	$(5.3433, 1.5889, -1.3269, 2.0000)^T$	241.51	2.3647×10^{-8}
	`lsq-LM`:	$(5.3433, 1.5893, -1.3252, 5.9303)^T$	608.27	2.3646×10^{-8}
High	`fmincon`:	$(5.3433, 1.6380, -0.2995, 1.3012)^T$	238.58	1.03291×10^{-7}
	`lsq-TRR`:	$(5.3433, 1.6361, -1.990, 0.2496)^T$	606.36	1.03257×10^{-7}
	`lsq-LM`:	$(5.3433, 1.6351, -0.02324, 3.6112 \times 10^{-4})^T$	1110.83	1.03248×10^{-7}
true values θ_0: $(5.3424, 1.6021, -1.3010, 1.0000)^T$				

Even though there might be some spurious computation times on desktop machines (due to other background programs), we still include them here in Table 1 to demonstrate typical optimization routine performance. Consistently, `fmincon` was the fastest routine. This is in part due to the fact that this routine alone of the three supports parallel computation, so on our multi-core desktop machines we were able to see a speed-up. However, the computation times for the trust-region-reflective `lsqnonlin` algorithms are reasonable. Using Levenburg-Marquardt consistently is the slowest method, and the results are not better than those using trust-region-reflective `lsqnonlin` algorithm.

As a result, we recommend using the trust-region-reflective `lsqnonlin` algorithm when trying to estimate relaxation times (this is also the routine that we use in the remainder of this paper). If the model does not contain relaxation times (i.e., only estimating E_0 and E_1), the speedup afforded by using `fmincon` may make that algorithm the one of choice. Figure 4 illustrates model fits to the data at different noise levels, where the model solution is calculated with the values of model parameters obtained through `lsqnonlin` TRR routine. We see in all cases that the model solution provides reasonable fits to the data.

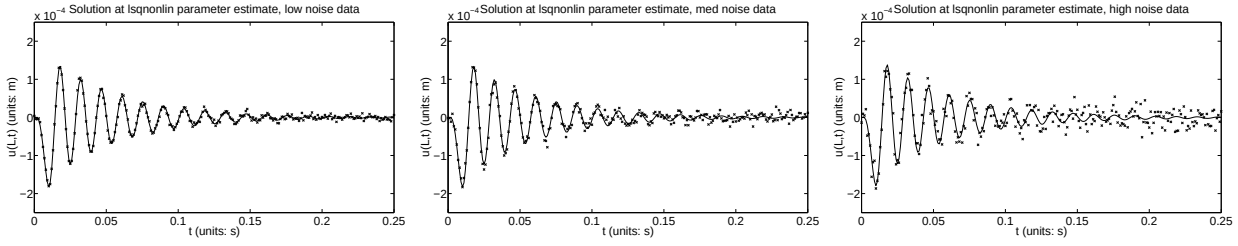

Figure 4: Data and two-relaxation-time model solutions at parameter estimates obtained using `lsqnonlin`, trust-region-reflective method, at different noise levels (see Table 1). (left pane) Low noise. (middle pane) Medium noise. (right pane) High noise.

3.2.3 Asymptotic error analysis

Most asymptotic error theory [5, 16] is described in the context of an ODE model example $\dot{z}(t) = f(z(t;\theta);\theta)$. However, we can use the PDE sensitivities of the model output with respect to each parameter in θ, namely $\dfrac{\partial u(L,t;10^{\theta})}{\partial \theta_i}$, in a similar manner to the ODE sensitivities in the asymptotic theory. The steps of the asymptotic theory error analysis are as follows (the theory for the following steps is described in [5, 16]).

1. Determine $\hat{\theta}$ by solving the inverse problem (14).

2. Compute the sensitivity equations to obtain $\dfrac{\partial u(L,t;10^{\hat{\theta}})}{\partial \theta_i}$ (as discussed in Section 3.1) for $i = 1, \ldots, \kappa$ where κ is the number of parameters being estimated. The sensitivity matrix $\chi(\hat{\theta})$ can then be calculated with its entries

$$\chi_{j,i}(\hat{\theta}) = \frac{\partial u(L,t_j;10^{\hat{\theta}})}{\partial \theta_i}, \quad j = 0, 1, \ldots, n-1, \text{ and } i = 1, \ldots, \kappa.$$

   Note that $\chi(\hat{\theta})$ is then an $n \times \kappa$ matrix. We can also obtain an estimate for the constant variance σ_0^2 as

$$\hat{\sigma}^2 = \frac{1}{n-\kappa} \sum_{j=0}^{n-1} \left(u_j - u(L,t_j;10^{\hat{\theta}}) \right)^2.$$

3. Asymptotic theory yields that the estimator $\hat{\Theta}$ is asymptotically (as sample size $n \to \infty$) normal with mean approximated by $\hat{\theta}$ and the covariance matrix approximated by

$$Cov(\hat{\Theta}) \approx \hat{\Sigma} = \hat{\sigma}^2 [\chi^T(\hat{\theta})\chi(\hat{\theta})]^{-1}.$$

4. The standard errors for each element in the parameter estimator $\hat{\Theta}$ can be approximated by

$$SE(\hat{\Theta}_i) = \sqrt{\hat{\Sigma}_{ii}}, \quad i = 1, 2, \ldots, \kappa,$$

   where $\hat{\Theta}_i$ is the ith element of $\hat{\Theta}$, and $\hat{\Sigma}_{ii}$ is the (i,i)th entry of the matrix $\hat{\Sigma}$. Hence, the endpoints of the confidence intervals for $\hat{\Theta}_i$ are given by

$$\hat{\theta}_i \pm t_{1-\alpha/2} SE(\hat{\Theta}_i)$$

   for $i = 1, 2, \ldots, \kappa$. Here $t_{1-\alpha/2}$ is a distribution value that is determined from a statistical table for Student's t-distribution based on the level of significance α (i.e., $\alpha = .05$ for a 95% confidence interval).

We will present results below in Tables 2-4 on the low, medium, and high noise data sets using zero, one, and two relaxation times, and using the routine `lsqnonlin` with trust-region-reflective option (interested readers can refer to [7] for the results obtained by using `fmincon`). We see throughout the tables that the problem of estimating the second relaxation time is fraught with

Table 2: TRR `lsqnonlin`: Parameter estimates, asymptotic standard errors (SE) and confidence intervals for zero-relaxation-time model (model 0), one-relaxation-time model (model 1) and two-times-relaxation model (model 2) obtained at low noise level.

Model	Params	True Value	Estimate	SE	95% Confidence Interval
0	$\log_{10}(E_0)$	5.3424	5.3422	4.9498×10^{-4}	(5.3413, 5.3432)
	$\log_{10}(E_1)$	1.6021	1.6651	0.005434	(1.6544, 1.6758)
1	$\log_{10}(E_0)$	5.3424	5.3425	0.01011	(5.3226, 5.3624)
	$\log_{10}(E_1)$	1.6021	1.6050	0.3167	(0.9811, 2.2288)
	$\log_{10}(\tau_1)$	-1.3010	-1.2317	2.2200	(-5.6041, 3.1407)
2	$\log_{10}(E_0)$	5.3424	5.3425	0.0101	(5.3225, 5.3624)
	$\log_{10}(E_1)$	1.6021	1.6046	0.3202	(0.9738, 2.2353)
	$\log_{10}(\tau_1)$	-1.3010	-1.2297	2.2369	(-5.635, 3.1761)
	$\log_{10}(\tau_2)$	1	1.0046	16.0237	(-30.5560, 32.5651)

Table 3: TRR `lsqnonlin`: Parameter estimates, asymptotic standard errors (SE) and confidence intervals for zero-relaxation-time model (model 0), one-relaxation-time model (model 1) and two-times-relaxation model (model 2) obtained at medium noise level.

Model	Params	True Value	Estimate	SE	95% Confidence Interval
0	$\log_{10}(E_0)$	5.3424	5.3429	9.2836×10^{-4}	(5.3411, 5.3448)
	$\log_{10}(E_1)$	1.6021	1.6653	0.0102	(1.6452, 1.6854)
1	$\log_{10}(E_0)$	5.3424	5.3433	0.01042	(5.3228, 5.3638)
	$\log_{10}(E_1)$	1.6021	1.6050	0.3167	(0.9811, 2.2288)
	$\log_{10}(\tau_1)$	-1.3010	-1.2317	2.2200	(-5.6041, 3.1407)
2	$\log_{10}(E_0)$	5.3424	5.3433	0.01045	(5.3227, 5.3639)
	$\log_{10}(E_1)$	1.6021	1.5889	0.3717	(0.8567, 2.3211)
	$\log_{10}(\tau_1)$	-1.3010	-1.3269	2.0383	(-5.3415, 2.6878)
	$\log_{10}(\tau_2)$	1	2.0000	156.5630	(-306.369, 310.369)

difficulty (the standard error is significantly higher than its estimated value), even though we know the simulated data came from a model incorporating two relaxation times. This could be predicted from our earlier examination of the sensitivities with respect to the second relaxation time, as well as the results for relaxation times seen when using different optimization routines. In addition, when estimating two relaxation times on high noise data (shown in Table 4) we see that the estimates for τ_1 and τ_2 are not close to the true parameter values; also, the standard error for τ_1 is much larger than in any other case. Thus, instead of merely having difficulty estimating a second relaxation time, in this estimation we now additionally have a less confidence on the estimate of τ_1.

We observe from Tables 2-4 that the standard errors for E_0 and E_1 increase significantly at all noise levels when moving from the zero-relaxation-time model to the one-relaxation-time model. We also found that at all noise levels the difference for the residual sum of squares is small among the zero-relaxation-time, one-relaxation-time, and two-relaxation-time models (see the third columns of Table 5). In addition, for each level noise data set, when we plot the model solutions corresponding

Table 4: TRR `lsqnonlin`: Parameter estimates, asymptotic standard errors (SE) and confidence intervals for zero-relaxation-time model (model 0), one-relaxation-time model (model 1) and two-times-relaxation model (model 2) obtained at high noise level.

Model	Params	True Value	Estimate	SE	95% Confidence Interval
0	$\log_{10}(E_0)$	5.3424	5.3433	1.7767×10^{-3}	(5.3398, 5.3468)
	$\log_{10}(E_1)$	1.6021	1.6452	0.0204	(1.6050, 1.6855)
1	$\log_{10}(E_0)$	5.3424	5.3433	0.01046	(5.3227, 5.3639)
	$\log_{10}(E_1)$	1.6021	1.6397	0.1526	(1.3391, 1.9403)
	$\log_{10}(\tau_1)$	-1.3010	-1.3253	2.0328	(-5.3291, 2.6785)
2	$\log_{10}(E_0)$	5.3424	5.3433	0.01046	(5.3227, 5.3639)
	$\log_{10}(E_1)$	1.6021	1.6361	0.1591	(1.3227, 1.9496)
	$\log_{10}(\tau_1)$	-1.3010	-0.1990	15.8821	(-31.4806, 31.0827)
	$\log_{10}(\tau_2)$	1	0.2496	12.0051	(-23.3958, 23.8951)

to the zero, one, and two relaxation time models, we found that they approximately lie on top of each other, and give good fits to the data. To gain further insight into which model should be chosen, we turned to some model selection criterion analysis.

Model selection criteria There are numerous model selection criteria in the literature that can be used to select a best approximating model from a prior set of candidate models. These criteria are based either on hypothesis testing or mean squared error or Bayes factors or information theory, and they all are based to some extent on the principle of parsimony (see [17]). It should be noted that some of these criteria can only be used for nested models (e.g., two models are said to be nested if one model is a special case of the other), but others can be used for both nested models and non-nested models.

Here we employ one of the most widely used model selection criteria – the Akaike information criterion (AIC). The AIC was developed by Akaike (in 1973) who formulated a relationship between the Kullback-Leibler information (used to measure the information lost when a model is used to approximate the true model) and the maximum value of the log likelihood function of the approximating model. As might be expected we find that the AIC value depends on the data set used. Thus, when we try to select a best model from a set of candidate models, we must use the same data set to calculate AIC values for each of the models. One of the advantages of the AIC is that it can be used to compare non-nested models (which is our case here). For the least squares case, it can be found (e.g., see [17, Section 2.2]) that if the observation errors are i.i.d normally distributed, then the AIC is given by

$$AIC = n \log\left(\frac{\mathrm{RSS}}{n}\right) + 2(\kappa + 1). \tag{16}$$

Here $\kappa + 1$ is the total number of estimated parameters including θ and the observation error variance. Given a prior set of candidate models, we can calculate the AIC value for each model, and the best approximating model is the one with minimum AIC value. It should be noted that the AIC may perform poorly if the sample size n is small relative to the total number of estimated

parameters (it is suggested in [17] that the sample size n should be at least 40 times the total number of estimated parameters $(\kappa + 1)$; note this is true for our investigations).

In practice, the absolute size of the AIC value may have limited use in supporting the chosen best approximating model, and one may often employ other related values such as Akaike differences and Akaike weights to further compare models. The Akaike difference is defined by

$$\Delta_i = AIC_i - AIC_{\min}, \quad i = 1, 2, \ldots R, \tag{17}$$

where AIC_i is the AIC value of the ith model in the set, $AIC_{\min}$ denotes the AIC value for the best model in the set, and R is the total number of models in the set. The larger Δ_i, the less plausible it is that the ith model is a good approximating model for given the data set. The Akaike weights are defined by

$$w_i = \frac{\exp(-\frac{1}{2}\Delta_i)}{\sum_{r=1}^{R} \exp(-\frac{1}{2}\Delta_r)}, \quad i = 1, 2, \ldots R. \tag{18}$$

These Akaike weights w_i can then be interpreted as the probability that ith model is the best approximating model (see [17]).

Table 5 presents residual sum squares (RSS), AIC values, AIC differences, and AIC weights obtained for the two-relaxation-time model, the one-relaxation-time model and the zero-relaxation-time model at low, medium and high noise levels using lsqnonlin. From this table, we see that on

Table 5: `lsqnonlin`: Residual sum of squares (RSS), AIC values, AIC differences (Δ) and AIC weights for zero-relaxation-time model (model 0), one-relaxation-time model (model 1) and two-relaxation-times model (model 2) obtained at low, medium and high noise levels.

Noise level	Model	RSS	AIC	Δ	AIC weights
low noise	0	7.0368×10^{-9}	-6.0927×10^{3}	27.8791	6.4125×10^{-7}
	1	6.2470×10^{-9}	-6.1206×10^{3}	0	7.2595×10^{-1}
	2	6.2458×10^{-9}	-6.1186×10^{3}	1.9483	2.7405×10^{-1}
medium noise	0	2.4674×10^{-8}	-5.7778×10^{3}	8.6863	9.4255×10^{-3}
	1	2.3646×10^{-8}	-5.7865×10^{3}	0	7.2527×10^{-1}
	2	2.3647×10^{-8}	-5.7845×10^{3}	2.0113	2.6531×10^{-1}
high noise	0	1.0337×10^{-7}	-5.4182×10^{3}	0	6.4303×10^{-1}
	1	1.0330×10^{-7}	-5.4164×10^{3}	1.8299	2.5756×10^{-1}
	2	1.0326×10^{-7}	-5.4145×10^{3}	3.7340	9.9406×10^{-2}

low and medium level noise data sets the one-relaxation-time model is the best with the probability to be chosen as the best model being more than 0.7 (see the Akaike weights in the last column of these two tables), and the zero-relaxation time model has almost no chance of being selected as the best. For the high noise data set, the zero-relaxation-time model is the best, with the probability of being chosen as the best being more than 0.6, while the two-relaxation-time model has little chance of being selected as the best model. It is worth noting here that we obtain similar conclusions from the results obtained by using `fmincon` (the interested readers can refer to [7] for details).

Summary remark Based on our analysis to this point, we can conclude that estimating two relaxation times is likely to be difficult. Adopting a model with zero or one relaxation times may be the most feasible approach. However, until we can confirm this approach by examining these methods on experimental data we believe that attempting all three options for including relaxation times in the viscoelastic model (zero, one, or two times) is advisable.

3.2.4 Bootstrapping error analysis

For ease of presentation, we reiterate here the algorithm described in [5], in the context of the current viscoelastic model under study.

1. Determine $\hat{\theta}^0$ by solving the inverse problem (14).

2. Define the *standardized residuals* (recall n is the number of data points, and κ is the number of parameters under consideration) to be

$$\bar{r}_j = \sqrt{\frac{n}{n-\kappa}}\left(u_j - u(L, t_j; 10^{\hat{\theta}^0})\right)$$

   for $j = 0, 1, \ldots, n-1$. Set $m = 0$.

3. Create a sample of size n by randomly sampling, with replacement, from the standardized residuals $\bar{r}_j$ to form a bootstrap sample $\{r_0^m, \ldots r_{n-1}^m\}$.

4. Create bootstrap sample points

$$u_j^m = u(L, t_j; 10^{\hat{\theta}^0}) + r_j^m, \quad j = 0, 1, \ldots, n-1.$$

5. Solve the OLS minimization problem (14) with the bootstrap-generated data $\{u_j^m\}$ to obtain a new estimate $\hat{\theta}^{m+1}$ which we store.

6. Increase the index m by 1 and repeat steps 3-5. This iterative process should be carried out for M times where M is large (we used $M = 1000$, as suggested in [5]). This will give M estimates $\{\hat{\theta}^m\}_{m=1}^M$.

Upon completing all M simulation runs, the following will give the mean and covariance matrix for the bootstrap estimator $\hat{\Theta}_{boot}$ of θ:

$$\hat{\theta}_{boot} = \frac{1}{M}\sum_{m=1}^{M}\hat{\theta}^m, \quad \hat{\Sigma}_{boot} = \frac{1}{M-1}\sum_{m=1}^{M}(\hat{\theta}^m - \hat{\theta}_{boot})(\hat{\theta}^m - \hat{\theta}_{boot})^T. \tag{19}$$

Then the standard errors for the bootstrap estimator $\hat{\Theta}_{boot}$ is given by

$$(SE_{boot})_i = \sqrt{\left(\hat{\Sigma}_{boot}\right)_{ii}}, \quad i = 1, 2, \ldots, \kappa,$$

where $\left(\hat{\Sigma}_{boot}\right)_{ii}$ is the (i,i)th entry of covariance matrix $\hat{\Sigma}_{boot}$. Hence, the endpoints of the confidence intervals for $(\hat{\Theta}_{boot})_i$ (the ith element of $\hat{\Theta}_{boot}$) are given by

$$(\hat{\theta}_{boot})_i \pm t_{1-\alpha/2}(SE_{boot})_i$$

for $i = 1, 2, \ldots, \kappa$.

Note that bootstrapping requires solving the inverse problem 1000 times. Even for a model that is solved in a short time (e.g., less than one minute), bootstrapping takes a significant time to compute (as we must solve the inverse problem many times and each inverse problem involves solving the model multiple times). Due to long computational times (e.g., one week for bootstrapping versus minutes for the asymptotic theory), we report here results only for a case using `lsqnonlin`, the trust-region-reflective option, to estimate E_0, E_1, and τ_1 in a one-relaxation-time model. It is worth noting that even though the bootstrapping algorithm can be implemented in parallel, this requires a considerable amount of computing resources (unavailable to most investigators) to achieve computational times comparable to that attained in using the asymptotic theory. For our purposes, the bootstrap results we provide are sufficient to indicate that the less conservative asymptotic error analysis yields a reasonable uncertainty measure in the inverse problem we investigate.

The bootstrapping results for a one-relaxation-time model obtained by using `lsqnonlin` TRR routine are summarized in Table 6. We see that confidence intervals for all parameters are wider

Table 6: TRR `lsqnonlin`: Parameter estimates, bootstrap standard errors (SE) and confidence intervals obtained at low, medium and high noise levels for one-relaxation-time model.

Noise level	Params	True Value	$\hat{\theta}_{boot}$	SE	95% Confidence Interval
Low noise	$\log_{10}(E_0)$	5.3424	5.3425	0.01547	(5.3120, 5.3730)
	$\log_{10}(E_1)$	1.6021	1.6025	0.5937	(0.4332, 2.7719)
	$\log_{10}(\tau_1)$	-1.3010	-1.2294	3.8697	(-8.8510, 6.3923)
Medium noise	$\log_{10}(E_0)$	5.3424	5.3434	0.03136	(5.2816, 5.4052)
	$\log_{10}(E_1)$	1.6021	1.5852	1.4590	(-1.2884, 4.4589)
	$\log_{10}(\tau_1)$	-1.3010	-1.2079	16.9971	(-34.6849, 32.2692)
High noise	$\log_{10}(E_0)$	5.3424	5.3434	0.061762	(5.2218, 5.4651)
	$\log_{10}(E_1)$	1.6021	1.6029	2.4545	(-3.2314, 6.4372)
	$\log_{10}(\tau_1)$	-1.3010	-0.1592	40.8381	(-80.5930, 80.2746)

than those obtained using the asymptotic error theory, especially for the cases of medium and high noise level (the same phenomenon was also observed for the results obtained for a zero-relaxed-time model by using `fmincon` routine, see [7] for details). However, this is expected. At the low noise level, we obtained fairly good results for E_0 and E_1 but the standard error for the relaxation time is larger in magnitude than the relaxation time value itself. This is even more prominent at higher noise levels – the results in the table indicate that on medium and high noise data sets, the estimation of τ_1 is not very robust. Note also that the estimation of E_1 begins to suffer as well, resulting in a higher standard error than its own value on the high noise data set. This is a further indication that we may have problems in the future estimating even the single relaxation time.

We depict histograms of the estimates in Figure 5. We see on a low noise data set that each

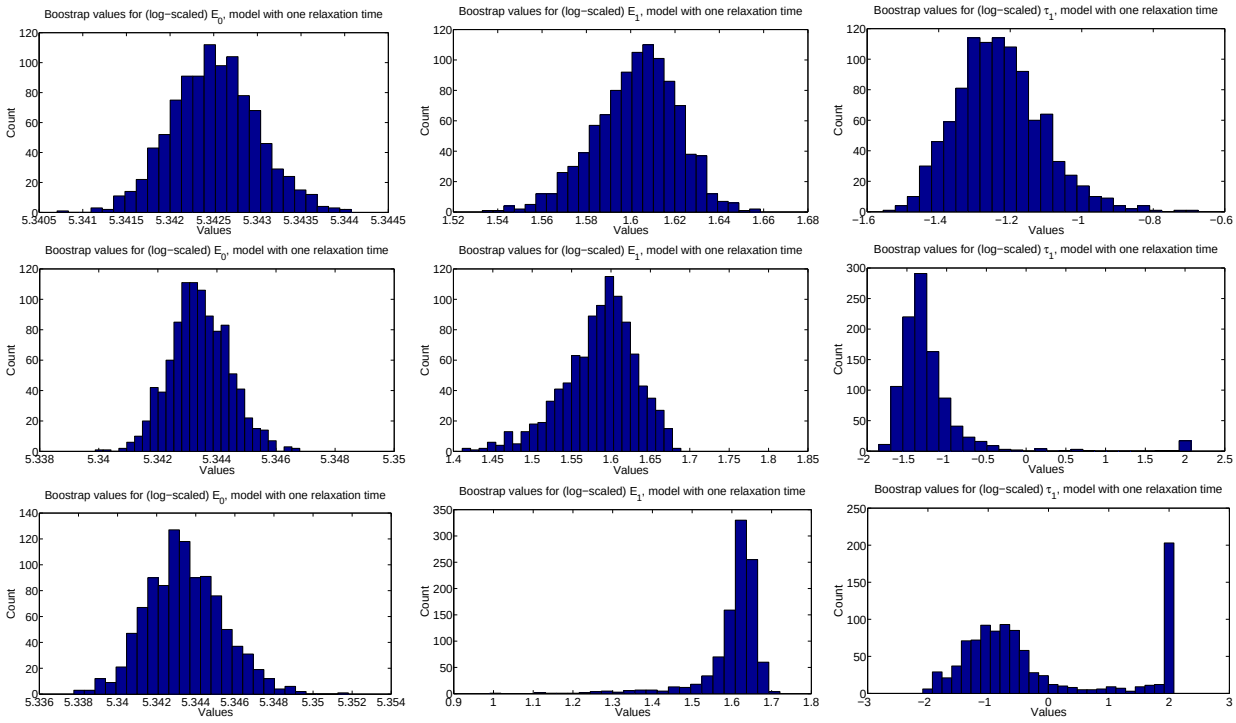

Figure 5: Histograms of bootstrap estimates $\hat{\theta}^m$ for a one-relaxation-time model obtained at low noise (upper row), medium noise (middle row) and high noise (bottom row) levels. (left column) Estimates for $\log_{10}(E_0)$; (middle column) Estimates for $\log_{10}(E_1)$; (right column) Estimates for $\log_{10}(\tau_1)$.

parameter estimator appears to be mostly normally distributed. This begins to break down for the case of middle noise level data set (shown in the middle row of Figure 5), where we begin to see some outliers at the $\log_{10}(\hat{\tau}_1) = 2$ level (which means the estimates were converging to our upper bound on that parameter) and also some more pronounced skewness in the count levels. Finally, on the high noise level (shown in bottom row of Figure 5) we have the distribution for the E_1 estimates skewed, and we also observe a clear proliferation of estimates of the first relaxation time approaching the value 2. This further supports the expectation of difficulty in estimating relaxation times, particularly when the noise level is high. It is worth noting here that for the zero-relaxation-time model each parameter estimator, obtained by `fmincon` routine, tends to have the shape of a normal distribution regardless of the noise level - see [7] for details.

4 Model comparison and hypothesis testing on amplitude

In this section, we develop a methodology for determining whether or not data came from a low-amplitude input traction. This simulates the problem of determining if the data came from a vessel experiencing a normal heartbeat or not. We will ultimately run the inverse problem without amplitude restrictions and use a scoring function to compare results with the score of the model solved at a low amplitude. A model comparison test will be implemented to determine if there is

statistical significance in the differences between the model solved with the unrestricted estimate and the model solved using the restricted amplitude value.

4.1 Setup

We first examine the sensitivity of the model with respect to the Van Bladel input amplitude parameter A, to insure that an estimation procedure is reasonable (if the model were insensitive to A then the results from the optimization routine would be suspect). The form of the sensitivity equation is nearly identical to that of the actual model, just with a lower amplitude. This is seen in Figure 6, which has a form similar to that of the model solution (depicted by the solid line in Figure 3). In both the low and high amplitude cases, the sensitivity with respect to amplitude is most

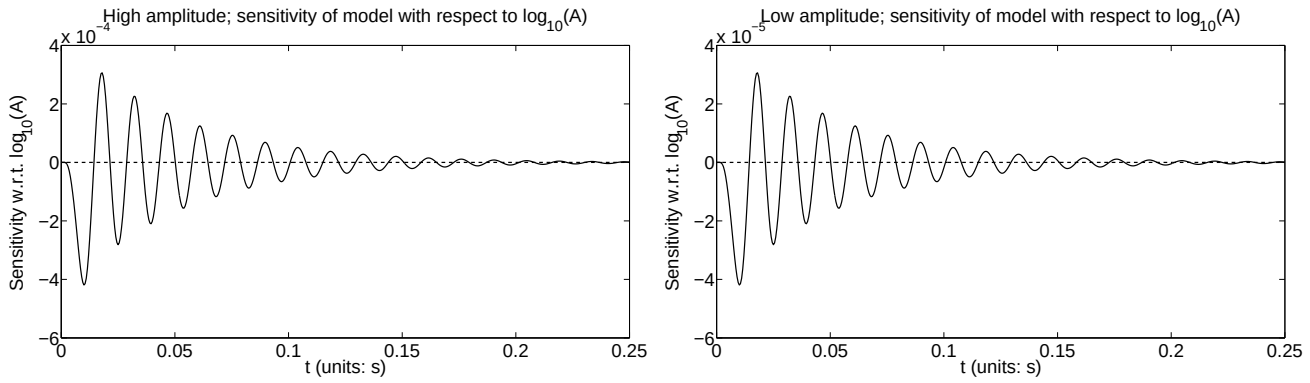

Figure 6: Sensitivity of model with respect to Van Bladel input parameter A around the baseline parameters (10). (left pane) High forcing function amplitude $A = 6 \times 10^3$. (right pane) Low forcing function amplitude $A = 6 \times 10^2$.

marked during early times and less so at later times; this makes perfect sense, as the amplitude is greater early on before being damped out. In the problem below, we will take data throughout the full time frame $t \in [0, 0.25]$ so with our sensitivity results we can be assured that the early data will drive estimation of the amplitude parameter.

4.2 Data generation

For the high amplitude data, we use the same low, medium, and high noise data sets as described in Section 3.2.1 and shown above in Figure 3. We form the low amplitude data by taking $A_{low} = A/10$ as our Van Bladel input amplitude parameter. Thus, the dynamics are roughly 10% the magnitude of the high amplitude data. This means the corresponding noise for the low noise, low amplitude data set will be generated with variance $\sigma^2 = 5 \times 10^{-7}$, medium noise with $\sigma^2 = 10 \times 10^{-7}$, and high noise with $\sigma^2 = 20 \times 10^{-7}$. The low amplitude input data set then is supposed to represent a normal heartbeat and the high amplitude data set then is meant to represent the input shear for a heartbeat in the presence of a stenosis in the vessel. Note that we are not yet exactly certain regarding the difference between these effects in an actual patient, so the data sets here are truly for a proof-of-concept investigation. The low amplitude data sets are depicted in Figure 7.

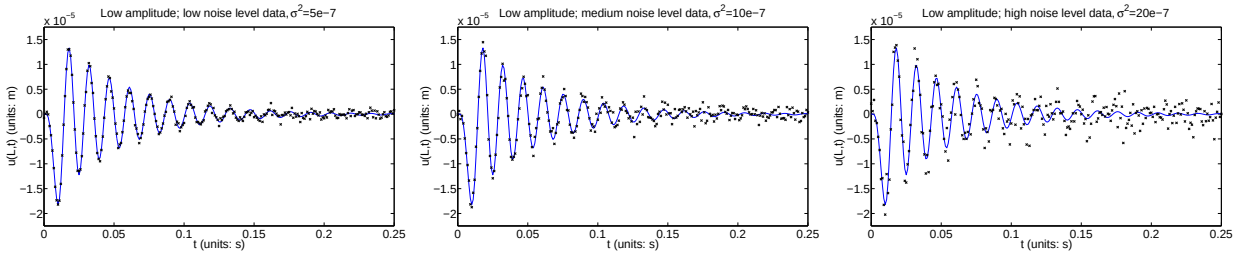

Figure 7: Simulated low amplitude noisy data around the true parameter values. (left pane) Low noise level, $\sigma_0^2 = 5 \times 10^{-7}$. (middle pane) Medium noise, $\sigma_0^2 = 10 \times 10^{-7}$. (right pane) High noise level, $\sigma_0^2 = 20 \times 10^{-7}$.

4.3 Hypothesis testing methodology

We can now begin to discuss the approach to model comparison and hypothesis testing that we will use by defining a model comparison test statistic. The work here follows the development in [16]. Our performance criterion for hypothesis testing will be

$$J(\vec{U}, \theta) = \sum_{j=0}^{n-1} [U_j - u(L, t_j; \theta)]^2.$$

For the purposes of this paper, we postulate that a normal (non-stenosed) vessel corresponds with a low amplitude input parameter $A \leq 6 \times 10^2$. Then, a stenosed vessel would have a high input amplitude parameter with $A > 6 \times 10^2$. The hypothesis test we use requires a set benchmark value for A, so we choose that benchmark to be $A_0 = 6 \times 10^2$. Then, we define the restricted parameter set

$$\mathbb{A}_H = \{A \in \mathbb{A} | A = A_0 = 6 \times 10^2\},$$

where $\mathbb{A} = [A_0, \infty)$ is the larger set of unrestricted admissible amplitudes.

Our null hypothesis H_0 is that the amplitude is a low amplitude, represented by $A \in \mathbb{A}_H = \{A_0\}$. The unrestricted amplitude model would then represent the amplitude parameter as $A = A_0 + \tilde{A}$ where $\tilde{A} \in [0, \infty)$. This framework will allow us to develop a test statistic to determine the confidence level of accepting or rejecting H_0 for a given data set. In other words, we will develop a test to determine if the data is statistically better represented by the benchmark A_0 than the unrestricted amplitude.

The first step is to determine the performance criterion at the benchmark amplitude $\hat{\theta}_H = 6 \times 10^2$, which we will denote $J(\vec{u}, \hat{\theta}_H)$ (Since the value $\hat{\theta}_H$ is fixed in our case, no optimization problem is needed to compute these values). We then run an optimization routine to determine an unrestricted input amplitude parameter estimate $\hat{\theta}$, which we then use to compute $J(\vec{u}, \hat{\theta})$. The value for $\hat{\theta}$ comes from solving the unrestricted optimization problem (14). As discussed in [16, 4], the model comparison statistic is defined as

$$\hat{V} = n \frac{J(\vec{U}, \hat{\Theta}_H) - J(\vec{U}, \hat{\Theta})}{J(\vec{U}, \hat{\Theta})}$$

with realization

$$\hat{v} = n\frac{J(\vec{u},\hat{\theta}_H) - J(\vec{u},\hat{\theta})}{J(\vec{u},\hat{\theta})}. \tag{20}$$

If our null hypothesis H_0 were true, the model comparison statistic $\hat{V}$ converges in distribution to V as $n \to \infty$ where $V \sim \chi^2(r)$ is a chi-square distribution with r degrees of freedom (r is the number of constraints in $\mathbb{A}_H$). For our problem, $r = 1$. Given the significance level α, we can obtain a threshold value ν such that the probability that V will take on a value greater than ν is α. In other words, $\text{Prob}(V > \nu) = \alpha$. In our context, if the test statistic $\hat{v} > \nu$ we reject H_0 as false with confidence level $(1-\alpha)100\%$. Otherwise we do not reject H_0 as false, at the specified confidence level. In Table 7 we include sample values from the $\chi^2(1)$ distribution for reference (table repeated from [16]).

Table 7: Sample $\chi^2(1)$ values.

α	ν	confidence
0.25	1.32	75%
0.1	2.71	90%
0.05	3.84	95%
0.01	6.63	99%
0.001	10.83	99.9%

We summarize in Table 8 the results of computing the OLS performance criterion for the low amplitude and high amplitude data each with the restricted/unrestricted parameters. Based on this table and Table 7, we see for both the low and medium noise cases with data generated with a low A value that we do not reject H_0 with high degrees of confidence. However, the case with high noise is somewhat less certain, though we would still likely not reject H_0 with a fairly high degree of confidence. The results are more stark in the cases where the data was generated from a high amplitude. Given that the magnitude of $\hat{v}$ is greater than 900 at all noise levels, we would reject H_0 as false on these data sets with confidence level more than 99.9%. Altogether, these results suggest robustness in our methodology for determining whether the data came from a normal vessel experiencing a heartbeat (low input amplitude) or from a abnormal (stenosed) response.

Table 8: Model comparison test results using (20) on low, medium, and high noise data sets generated with both high and low input amplitude parameter A values.

	$J(\vec{u},\hat{\theta})$	$J(\vec{u},\hat{\theta}_H)$	$\hat{v}$
Low A, low noise	6.3846e-11	6.3887e-11	0.1609
Low A, medium noise	2.6872e-10	2.6896e-10	0.2258
Low A, high noise	9.8836e-10	9.9658e-10	2.0878
High A, low noise	6.6812×10^{-9}	3.5229×10^{-7}	1.2984e+04
High A, medium noise	3.1016×10^{-8}	3.4730×10^{-7}	2.5596e+03
High A, high noise	9.9737×10^{-8}	4.6015×10^{-7}	907.0283

5 Conclusion

In this work we have carried out proof-of-concept investigations for estimating material parameters and created a model comparison test as a basis for distinguishing between data that comes from a normal or from a stenosed blood vessel. We found that the model was less sensitive to a second viscoelastic relaxation time than to the other parameters, and this was manifested as a difficulty in recovering two relaxation times. On the other hand, models with zero or one relaxation time allowed for more confidence in the estimation procedure (i.e., smaller standard errors). We compared asymptotic error theory with bootstrapping error theory, and found (as expected) that bootstrapping gives more conservative confidence intervals but not so much so that the asymptotic theory cannot be profitably used for uncertainty quantification in models with large computational costs rendering bootstrapping less desirable. In terms of the model comparison on the input amplitude parameter A, we were able to develop a successful methodology for statistically determining whether or not data came from a low amplitude input force. This will form the basis of a model comparison test we can use on experimental data sets.

In future efforts, we may need to examine the possibility of relative error instead of absolute error, which will necessitate a generalized least squares (GLS) cost function in our inverse problems due to changes in the error process. This will be coupled with a study of a statistical model for the measurement processes being used in the experiments at QMUL. The changes needed are discussed in [16]. If we require a GLS framework, we will use the proper model comparison framework for GLS problems given in [8]. The most important immediate efforts will be to apply the methods presented in this paper to the data from the QMUL experiments.

6 Acknowledgements

This research was supported in part (HTB,SH,ZRK) by Grant Number R01AI071915-09 from the National Institute of Allergy and Infectious Diseases, in part (HTB,SH,ZRK) by the Air Force Office of Scientific Research under grant number FA9550-09-1-0226, in part (ZRK) by the Department of Education with a GAANN Fellowship under grant number P200A070386, in part (CK,SS,JW) by the Engineering and Physical Sciences Research Council EP/H011072/1, and in part (MBr,SG,MBi) by the Engineering and Physical Sciences Research Council EP/H011285/1. The authors are grateful to several anonymous referees for their valuable comments and suggestions for improving this manuscript.

References

[1] H.T. Banks, *A Functional Analysis Framework for Modeling, Estimation and Control in Science and Engineering*, CRC Press/Taylor-Francis, Boca Raton London New York, 2012.

[2] H.T. Banks, J.H. Barnes, A. Eberhardt, H. Tran, and S. Wynne, Modeling and computation of propagating waves from coronary stenoses, *Computation and Applied Mathematics*, **21** (2002), 767–788.

[3] H. T. Banks and K. L. Bihari, Modeling and estimating uncertainty in parameter estimation, *Inverse Problems*, **17** (2001), 1–17.

[4] H.T. Banks and B.G. Fitzpatrick, Inverse problems for distributed systems: statistical tests and ANOVA, LCDS/CCS Rep. 88-16, July, 1988, Brown University; *Proc. International Symposium on Math. Approaches to Envir. and Ecol. Problems*, Springer Lecture Notes in Biomath., **81** (1989), 262–273.

[5] H.T. Banks, K. Holm and D. Robbins, Standard error computations for uncertainty quantication in inverse problems: Asymptotic theory vs. Bootstrapping, CRSC-TR09-13, North Carolina State University, June 2009; Revised May 2010; *Mathematical and Computer Modelling*, **52** (2010), 1610–1625.

[6] H.T. Banks, S. Hu and Z.R. Kenz, A brief review of elasticity and viscoelasticity for solids, *Advances in Applied Mathematics and Mechanics*, **3** (2011), 1–51.

[7] H.T. Banks, S. Hu, Z.R. Kenz, C. Kruse, S. Shaw, J.R. Whiteman, M.P. Brewin, S.E. Greenwald and M.J. Birch, Material parameter estimation and hypothesis testing on a 1d viscoelastic stenosis model: methodology, Technical Report CRSC-TR12-09, Center for Research in Scientific Computation, North Carolina State University, April, 2002.

[8] H.T. Banks, Z.R. Kenz and W.C. Thompson, An extension of RSS-based model comparison tests for weighted least squares , CRSC-TR12-18, N. C. State University, Raleigh, NC, August, 2012; *Intl. J. Pure and Appl. Math.,* **79** (2012), 155–183.

[9] H.T. Banks, Z.R. Kenz and W. Clayton Thompson, A review of selected techniques in inverse problem nonparametric probability distribution estimation, CRSC-TR12-13, N. C. State University, Raleigh, NC, May, 2012; *J. Inverse and Ill-posed Problems*, **20** (2012), 429–460; DOI 10.1515/jip-2012-0037.

[10] H.T. Banks and K. Kunisch, *Estimation Techniques for Distributed Parameter Systems*, Boston:Birkhaüser, 1989.

[11] H.T. Banks and N. Luke, Modeling of propagating shear waves in biotissue employing an internal variable approach to dissipation, *Communication in Computational Physics*, **3** (2008), 603–640.

[12] H.T. Banks, N. Medhin, and G. Pinter, Multiscale considerations in modeling of nonlinear elastomers, *International Journal for Computational Methods in Engineering Science and Mechanics*, **8** (2007), 53–62.

[13] H.T. Banks and G.A. Pinter, A probabilisitic multiscale approach to hysteresis in shear wave propagation in biotissue, *Multiscale Modeling and Simulation*, **3** (2005), 395–412.

[14] H.T. Banks, H. Tran and S. Wynne, A well-posedness result for a shear wave propagation model, *Intl. Series Num. Math.*, Vol. **143**, Birkhauser Verlag, Basel, 2002, 25–40.

[15] H.T. Banks and J.R. Samuels, Jr., Detection of cardiac occlusions using viscoelastic wave propagation, *Advances in Applied Mathematics and Mechanics*, **1** (2009), 1–28.

[16] H.T. Banks and H.T. Tran, *Mathematical and Experimental Modeling of Physical and Biological Processes*, CRC Press, Boca Raton, FL, 2009.

[17] K.P. Burnham and D.R. Anderson, *Model Selection and Inference: A Practical Information-Theoretical Approach*, Springer-Verlag, New York, 1998.

[18] B.D. Coleman and W. Noll, Foundations of linear viscoelasticity, *Reviews of Modern Physics*, **33** (1961), 239–249.

[19] P. Constantin and C. Foias, *Navier-Stokes Equations*, University of Chicago Press, Chicago, 1988.

[20] M. Fabrizio and A. Morro, *Mathematical Problems in Linear Viscoelasticity*, Studies in Applied Mathematics, Vol. 12, SIAM, Philadelphia, 1992.

[21] J.D. Ferry, *Viscoelastic Properties of Polymers*, John Wiley and Sons Inc., New York, 1970.

[22] W.N. Findley, J.S. Lai, and K. Onaran, *Creep and Relaxation of Nonlinear Viscoelastic Materials (With an Introduction to Linear Viscoelasticity)*, Dover Publications Inc., New York, 1989.

[23] A. Friedman, *Foundations of Modern Analysis*, Dover Publications, Inc., New York, 1982.

[24] Y.C. Fung, *Biomechanics: Mechanical Properties of Living Tissues*, Springer-Verlag, New York, 1993.

[25] J.M. Golen and G.A.C. Graham, *Boundary Value Problems in Linear Viscoelasticity*, Springer-Verlag, New York, 1988.

[26] Carola Kruse, Matthias Maischak, Simon Shaw, J.R. Whiteman, et al., High order space-time finite element schemes for acoustic and viscodynamic wave equations with temporal decoupling, BICOM Technical Report 12/3, 2012, in preparation.

[27] R. Lakes, *Viscoelastic Materials*, Cambridge Univ. Press, Cambridge, 2009.

[28] N. Luke, *Modeling Shear Wave Propagation in Biotissue: An Internal Variable Approach to Dissipation*, PhD Thesis, NCSU, Raleigh, 2006.

[29] S.E. Nissen, Application of intravascular ultrasound to characterize coronary artery disease and assess the progression or regression of atherosclerosis, *Am. J. Cardiol.*, **89** (2002), 24B–31B.

[30] S.E. Nissen and P. Yock, Intravascular ultrasound: novel pathophysiological insights and current clinical applications, *Circulation*, **103** (2001), 604–616.

[31] N.L. Owsley and A.J. Hull, Beamformed nearfield imaging of a simulated coronary artery containing a stenosis, *IEEE Trans. on Medical Imaging*, **17** (1998), 900–909.

[32] N.L. Owsley, A.J. Hull, M.H. Ahmed, and J. Kassal, A proof of concept experiment for the detection of occluded coronary arteries using array sensor technology, *Engr. in Medicine and Biol. Society*, IEEE 17th Annual Conf., **1** (1995), 145–146.

[33] Yu. V. Prohorov, Convergence of random processes and limit theorems in probability theory, *Theor. Prob. Appl.*, **1** (1956), 157–214.

[34] J.R. Samuels, Jr., *Inverse Problems and Post Analysis Techniques for a Stenosis-driven Acoustic Wave Propagation Model*, PhD Thesis, NCSU, Raleigh, 2008.

[35] J. Semmlow, J. Welkowtiz, W. Kostis, and J.W. Mackenzie, Coronary artery disease – correlates between diastolic auditory characteristics and coronary artery stenoses, *IEEE Trans. on Biomed. Engr.*, **30** (1983), 136–139.

[36] J. Verburg, Transmission of vibrations of the heart to the chestwall, *Adv. Cardiovasc. Phys.*, **5** (1983), 84–103.

[37] J. Verburg and E. van Vollenhoven, Phonocardiography: Physical and technical aspects and clinical uses, in *Noninvasive Physiological Measurements* (ed. P Rolfe), Academic Press, London (1979).

[38] J. Wloka, *Partial Differential Equations*, Cambridge University Press, Cambridge, 1987.

e-mail:
{htbanks, shu3, zrkenz}@ncsu.edu
{carola.kruse, simon.shaw, john.whiteman}@brunel.ac.uk
{s.e.greenwald, m.p.brewin,m.j.birch}@qmul.ac.uk