On the Advancement of Probabilistic Models of Decompression Sickness

by

Ethan Alexander Hada

Department of Mechanical Engineering and Materials Science
Duke University

Date:_____

Approved:

_____
Laurens E. Howle, Supervisor

_____
Richard Vann

_____
Wayne Gerth

_____
Brian Mann

_____
Stefano Curtarolo

Dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy in the Department of
Mechanical Engineering and Materials Science in the Graduate School
of Duke University

2016

ABSTRACT

On the Advancement of Probabilistic Models of Decompression Sickness

by

Ethan Alexander Hada

Department of Mechanical Engineering and Materials Science
Duke University

Date:_____

Approved:

_____
Laurens E. Howle, Supervisor

_____
Richard Vann

_____
Wayne Gerth

_____
Brian Mann

_____
Stefano Curtarolo

An abstract of a dissertation submitted in partial
fulfillment of the requirements for the degree
of Doctor of Philosophy in the Department of
Mechanical Engineering and Materials Science in the Graduate School of
Duke University

2016

Abstract

The work presented in this dissertation is focused on applying engineering methods to develop and explore probabilistic survival models for the prediction of decompression sickness in US NAVY divers. Mathematical modeling, computational model development, and numerical optimization techniques were employed to formulate and evaluate the predictive quality of models fitted to empirical data. In Chapters 1 and 2 we present general background information relevant to the development of probabilistic models applied to predicting the incidence of decompression sickness. The remainder of the dissertation introduces techniques developed in an effort to improve the predictive quality of probabilistic decompression models and to reduce the difficulty of model parameter optimization.

The first project explored seventeen variations of the hazard function using a well-perfused parallel compartment model. Models were parametrically optimized using the maximum likelihood technique. Model performance was evaluated using both classical statistical methods and model selection techniques based on information theory. Optimized model parameters were similar in magnitude to those of previously published probabilistic decompression models. Model performance results indicated that a novel hazard function definition that included both ambient pressure scaling and individually

fitted compartment exponent scaling terms outperformed alternative hazard function definitions.

We developed ten pharmacokinetic compartmental models that included explicit delay mechanics to determine if predictive quality could be improved through the inclusion of material transfer lags. A fitted discrete delay parameter augmented the inflow to the compartment systems from the environment. Based on the observation that symptoms are often reported after risk accumulation begins for many of our models, we hypothesized that the inclusion of delays might improve correlation between the model predictions and observed data. Model selection techniques identified two models as having the best overall performance, but comparison to the best performing model without delay and model selection using our best identified no delay pharmacokinetic model both indicated that the delay mechanism was not statistically justified and did not substantially improve model predictions.

Our final investigation explored parameter bounding techniques to identify parameter regions for which statistical model failure will not occur. When a model predicts a no probability of a diver experiencing decompression sickness for an exposure that is known to produce symptoms, statistical model failure occurs. Using a metric related to the instantaneous risk, we successfully identify regions where model failure will not occur and identify the boundaries of the region using a root bounding technique.

Several models are used to demonstrate the techniques, which may be employed to reduce the difficulty of model optimization for future investigations.

Contents

List of Tables

List of Figures

Acknowledgements

This effort was supported under NAVSEA contract N0002413C4104 under the supervision and mentorship of Dr. Laurens Howle of Duke University. Additional guidance and mentorship for this ongoing investigation can be attributed to F. Gregory Murphy and Ashleigh Swingler of Duke University, Dr. Wayne Gerth and Dr. David Doolette of the Naval Experimental Diving Unit, Dr. Richard Vann of Duke University Medical Center, and Dr. Michael Qin of the Naval Submarine Medical Research Laboratory.

1. Motivation

Decompression Sickness (DCS) is a risk for compressed air workers, divers, submariners, pilots, and astronauts. The modeling efforts of this project were focused on military divers specifically. Signs and symptoms of DCS are commonly associated with a reduction in the ambient environmental pressure of the subject. Symptoms of DCS are categorized as either Type I (Mild) or Type II (Serious) [1]. Type I symptoms are generally mild musculoskeletal injuries, especially pain in joints, and are much more common than Type II. Other Type I injuries, including cutaneous and constitutional effects, are less frequently reported. Type II injuries are categorized as more severe disorders and include neurologic and sensory effects, cardiopulmonary symptoms, and death [1]. DCS is believed to result from the formation of gaseous bubbles in blood and body tissues; a hypothesis that has guided both experimentation and modeling for most of the last century [2-7]. The primary motivation of this dissertation was exploration and development of probabilistic models and techniques used to predict the incidence of DCS.

Early modeling successes came in the form of deterministic gas content models which were initially developed by Haldane *et al.* [8]. The basic approach to the Haldane model was to divide the body into a series of independent, parallel, well-perfused tissue compartments that each represent the general mechanics of collections of connected or disconnected body organs and tissues but are not representative of specific body tissues. A compartment is a theoretical volume of material which is kinetically homogeneous with

respect to gas transfer and instantaneously applies material exchange throughout the entire volume. The idea of a compartment is motivated by the assumption that capillaries within the body tissue are spatially close such that diffusion within the compartment is instantaneous. We derived of the model differential equations explored in this dissertation assuming compartments were kinetically homogeneous.

Deterministic models have a binary calculation for the DCS outcome of a specific pressure exposure; either completely safe or completely unsafe. The DCS outcome of an exposure was determined by comparing the maximum supersaturation in each tissue compartment to a threshold value. Supersaturation refers to a state in which a greater amount of gas is dissolved in the tissue than can be tolerated by its equilibrium with the environment, as described by Henry's Law [9].

$$C_H = k_H P_H \tag{1}$$

where C_H is the solubility of gas in the tissue, k_H is the Henry gas law constant, and P_H is the partial pressure of the dissolved gas. If the supersaturation in any compartment exceeded its threshold, models indicate DCS will result. Thus, the Haldanian approach does not directly predict DCS, but instead models the factors which have been shown to be associated with symptoms (this is identical for all models explored in this dissertation). Many incremental improvements in deterministic models increased diver safety and expanded diving capabilities [3-7]. Deterministic models are generally simple, allowing them to be employed in submersible dive computers used today, but are limited in the

sense that they do not account for experimental observations which have demonstrated that not all divers will experience identical DCS outcomes for the same exposure (it should be mentioned that probabilistic models do not account for individual variation of divers, but only account for the statistical average across the experimental divers used for model evaluation).

Weathersby *et al.* [10] developed probabilistic DCS models by using survival analysis to address the binary outcome limitation of deterministic models. A general introduction to survival analysis can be found in Chapter 2. Probabilistic models explicitly model the probability of DCS for a given dive profile, allowing hyperbaric exposures to be planned with a calculated risk of DCS for divers. Development of probabilistic models was largely motivated by the observation that individual susceptibility to DCS can vary greatly for identical dive profiles.

Probabilistic DCS models are commonly parameterized through likelihood maximization; fitting model parameters to a training data set (discussed in Chapter 2). Likelihood maximization can be performed using a wide variety of optimization methods, and is a computationally expensive problem for two reasons. First, model likelihood hyper-surfaces are often locally discontinuous, which reduces the ability of gradient based optimization routines to accurately identify optimal parameter sets [11]. Second, the necessity of high resolution numerical integration of the risk function for accurate profile DCS probability calculation decreases overall optimization throughput (in this work we

use an integration step size of $1/10^{th}$ min for all models) [11]. Additionally, many of the best performing modern probabilistic models have a large number of fitted parameters, which increases the difficulty of optimization [12, 13]. By reducing either the number of fitted parameters or the difficulty of parametric optimization, researchers could investigate models more quickly, potentially allowing development of models with good predictive accuracy and low computational cost.

Parametrically simple models exist which are computationally inexpensive in both model development and execution time, but have been outperformed in the past by more complex models. Thalmann *et al*. [12] compared the linear-exponential (LE) model to the exponential-exponential (EE) model (described in Chapter 3). The LE model is a more functionally complex model than the EE model, and was determined to have superior predictive performance than the EE model. More recently, Howle *et al*. [11] re-examined the EE and LE models with a modification from the original model which included a minor mathematical error. Maximum LL parameter sets for EE were identified with a LL of -705, while the LE model obtained a maximum LL value of -697. Since the LE model requires a fitted crossover parameter, we believed that modifications to the more functionally simple EE model could potentially outperform the LE model through information theoretic model selection techniques.

We investigated two functional modifications to simple probabilistic models in the form of augmentation of the instantaneous hazard function (Chapter 3) and the

addition of explicit delay mechanics to simple pharmacokinetic models (Chapter 4). Our motivation for investigating these two techniques was to determine if the methods could improve predictive quality of two types of functionally simple probabilistic decompression models.

Parametric optimization of probabilistic models can be a very computationally expensive problem, especially for models with many fitted parameters. Determination of the parameter distributions in which to apply optimization techniques can be a challenging problem, which often motivates researchers to choose sufficiently large parametric search spaces that they feel confident good model parameters will be identified. In contrast, confinage of the parametric search space can reduce the difficulty of identifying the best predictive parameter set. By restricting parameter sets to only those which do not result in a predicted risk of zero on a profile with a recorded DCS injury, we can reduce the size of the search space considerably. In Chapter 5, we present a set of computational tools that simplify the exploration of the allowable model parameter space and help identify boundaries of model failure. Identification of the boundaries of the allowable parameter hyper-surface can improve optimization strategies through constrained optimization, increasing model optimization throughput.

Ultimately, the works contained in this dissertation are intended to demonstrate novel approaches developed in an effort to improve the predictive quality and ease of optimization and exploration for functionally simple probabilistic decompression models.

We hope that the techniques and approaches presented will aid in the ongoing efforts in development of real-time probabilistic computers and the overall improvement of both the capabilities and the safety of compressed air divers. Model predictive improvements are important for both recreational and operational divers, as the limited computational power of modern dive computers currently restricts the employment of real-time probabilistic models in submersible devices. The effort of this work is focused on the advancement of such probabilistic models through exploration of alternative hazard functions, simple functional changes to existing models, and the development of tools to simplify the difficult task of model parameterization.

2. Survival Analysis and the Hazard Function

Probabilistic models which predict the risk of DCS have primarily used survival analysis to estimate risk for a given dive profile [10]. Survival models are advantageous for modeling DCS because they operate with temporal data where the time an event occurs is assumed to be dependent on variable parameters. Additionally, outcome data is censored by either an event outcome or by the end of subject observation [14]. These kinds of models are widely applied in both prospective and retrospective pharmacological investigations in which subject exposure to a substance or condition of interest is related to a well-defined outcome. Probabilistic DCS models often simulate the kinetics of dissolved metabolic and non-metabolic gases, treating them as the exposure variable with DCS status as the model outcome [9].

Survival models depend on the formulation of a hazard function to describe the relationship between event probability and time. In this document, the term hazard function and risk function will be used interchangeably. The general form of the survival model is

$$P_0 = \exp\left(-\sum_i \int_0^{t_3} g_i h_i dt \right) \tag{2}$$

where P_0 is the probability of the exposure not resulting in DCS, t_3 is the final time of the exposure, $h_i(t)$ is a general hazard function, and g_i is a gain term that scales and non-dimensionalizes the hazard function in the i^{th} compartment and can be identified either

by parametric optimization or directly using the optimal gain technique developed by Howle [15]. Integration of the hazard function over a specified time interval provides a measure of the decompression stress accumulated within that interval. The summation of the decompression stress in all risk accumulating compartments is used by the exponential function to determine survival probability [16]. By the law of total probability, the sum of all possible outcomes must be unity, thus the probability of DCS, P_{DCS} is

$$P_{DCS} = 1 - P_0 \tag{3}$$

The hazard function is a mathematical abstraction used to relate the probability of a diver experiencing DCS on a given dive profile. Simple formulations of the hazard function may use a constant risk across the entire time domain, but more complex hazard definitions have been found to have superior predictive performance [4]. These functions generally use some combination of the ambient pressure and theoretical tissue gas content or properties of simulated bubbles along with estimated biological parameters to estimate DCS risk. An example of common hazard function is the supersaturation ratio, which was introduced by Weathersby *et al.* [4] and scales the accumulation of risk by the ambient pressure

$$h_i = \frac{P_{T_i} - P_{amb} + P_{FVG}}{P_{amb}} \tag{4}$$

where $P_{T_i}(t)$ is the inert gas tissue tension, $P_B(t)$ is the ambient pressure, and P_{FVG} accounts for gas pressure resulting from dissolved carbon dioxide and water vapor in the

blood ($P_{FVG} \approx$ 0.1917 Atm). Hazard functions scaled by ambient pressure have been commonly employed in the field [12, 17-19]. The hazard function defined in Eq.(4) can take on any value as there is no mathematical limit to the ambient pressure (although there is obviously a practical limit). Traditionally, the hazard function is positively clamped such that

$$h_i = \max\left(\frac{P_{T_i} - P_{amb} + P_{FVG}}{P_{amb}}, 0\right) \tag{5}$$

so that the hazard function can never take on negative values (this will be important in Chapter 5). In this dissertation, we will explore the effects on model predictive quality by comparing several forms of the hazard function in Chapter 3.

Probabilistic models used for prediction and prevention of DCS based on survival analysis is based in part on the method of maximum likelihood, which allows the optimization of model parameters to improve the correlation between model predicted estimates of DCS and observed incidents. Model optimization using the likelihood framework has several advantages, including allowing comparison of disparate non-nested models can be performed with information theoretic (IT) techniques, which has been very useful for model exploration [9, 20, 21]. Additionally, the likelihood framework allows model comparisons using classical statistical tools like the likelihood ratio test, which can be used to compare functionally similar models. All of the models we present in this dissertation were optimized using the method of maximum likelihood.

The likelihood function is a statistical function that defines the probability of a set of outcomes based on input parameters. In this context, we define the likelihood function in the same way that Weathersby et al. did to relate the predicted probabilities for DCS or no DCS in a single dive profile as

$$L = P_{DCS}{}^{\delta} P_0^{1-\delta} \tag{6}$$

where P_{DCS} and P_0 are defined as in equations (2) and (3). δ is a binary indicator for the injury outcome of the dive where $\delta = 1$ if the profile resulted in a DCS injury and $\delta = 0$ otherwise. Fractional weighting of the DCS outcome indicator δ to account for marginal DCS incidents has been explored in the past using values including 0.1 [12] and 0.5 [4], but recent a recent investigation by Weber [22] indicated that the marginal DCS incidents should be treated as non-events. Likelihood maximization is performed by parameterizing the model in such a way that the product of the likelihoods (shown in Eq.(7)) for the N profiles in the training data set is maximal.

$$\hat{L} = \prod_{n=1}^{N} L_n \tag{7}$$

A data set compromised of a large number of individual profiles will cause the likelihood value $\hat{L}$ to become very small, due to repeated multiplication individual profile likelihoods generated using Eq.(6) will be less than generally be smaller than 1 (

$L = 1$ would be a perfect prediction). To avoid numerical round off issues caused by small likelihood values, the logarithm of the likelihood (LL) is often maximized instead

$$LL = \sum_{n=1}^{N} \log\left(\hat{L}\right) \tag{8}$$

One important property to note is that model predictions that are significantly different from the observed data can cause the LL function to drop significantly. If the predicted risk is completely incorrect (i.e. $P_{DCS} = 0$ when $\delta = 1$), the LL value becomes negatively infinite (this is explored in detail in Chapter 5).

The likelihood function described by Eq.(6) uses the combination of the probability of DCS with the probability of no DCS over the entire length of the profile, and is referred to as an incidence-only likelihood equation. An alternative form of the likelihood definition that takes into account the time that the DCS event occurred is referred to as a time of onset equation, which replaces the $P\left(E\right)$ PDCS term from Eq.(3) with

$$P_{DCS} = \exp\left(-\sum_{i} \int_{0}^{t_1} g_i h_i dt \right) \left(1 - \exp\left(-\sum_{i} \int_{t_1}^{t_2} g_i h_i dt \right) \right) \tag{9}$$

where DCS symptoms were first identified in the predetermined time interval between t_1, the last known time without DCS symptoms, and t_2, the first known time when DCS symptoms existed (see Figure 1 [22]). Inclusion of the time of onset has been shown to improve the predictive correlation between models and observed decompression

incidence, but generally increases the difficult task of parameterizing models using likelihood maximization [19].

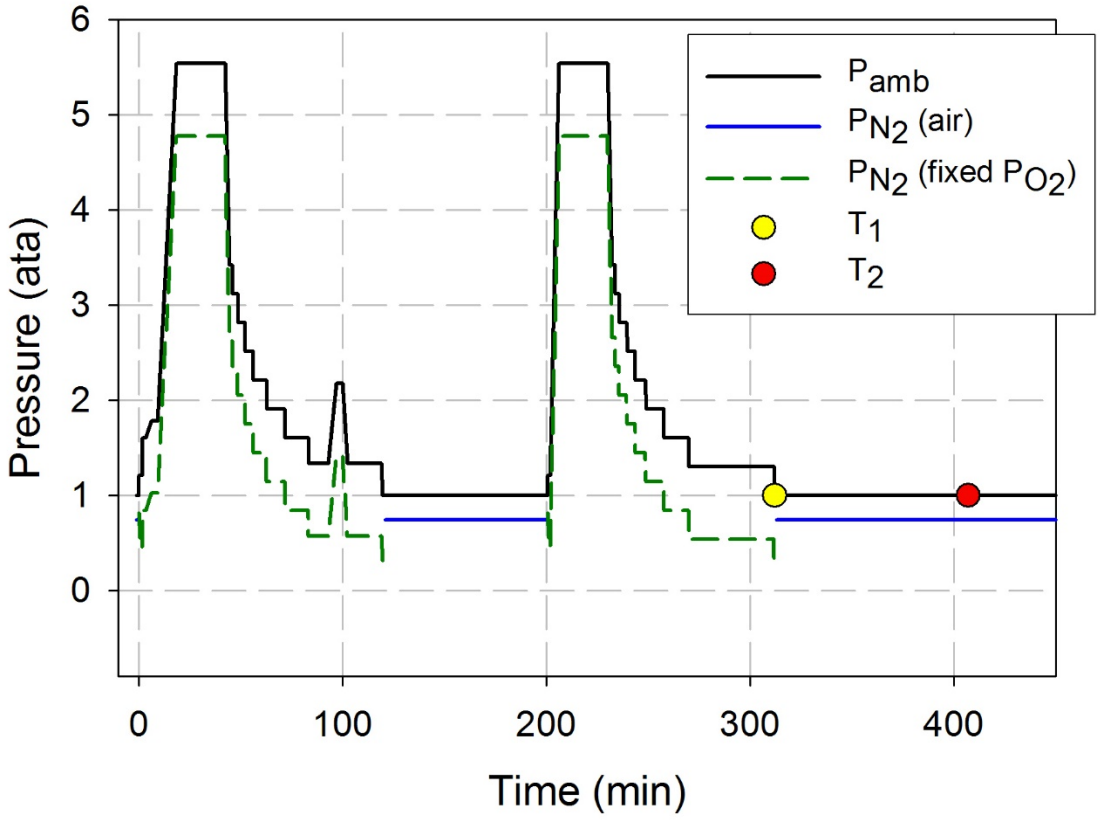

Figure 1: Diagram of a dive profile with two individual expeditions to depth and a brief surface interval. Change in atmospheric and gas pressures over time are shown along with a DCS hit between t_1 and t_2. Decompression stops are shown at various ambient pressures.

The form of the hazard function has a large impact on the risk accumulation properties of the model, so it is important that the best form is used in each application. Weathersby *et al.* [4] explored several hazard functional forms including the

supersaturation ratio (Eq.(4)), and found that ambient pressure scaling often improved model performance. The determination that risk should be scaled by ambient pressure was in agreement with the linear depth adjusted M-values method outlined by Workman in his 1957 report [6, 9]. Based on the performance of the supersaturation ratio hazard function, many subsequent probabilistic models used similar functional forms to describe instantaneous DCS risk [4, 12, 17-19]. It has been proposed that models using the supersaturation ratio may under-predict risk for profiles with a deep maximum depth, motivating exploration into alternative hazard functions [23]. Recently, the addition of an exponent on the numerator of the supersaturation ratio has been explored to allow for nonlinear scaling of risk with respect to ambient pressure. In Weathersby *et al*.[24], parametric optimization of the exponent parameter led to an indeterminate value, although a bounded range of exponential parameters was identified. Lovemann *et al.* [25] optimized the exponent parameter used in their investigation, but fewer than a third of the decompression incidents used for model optimization were man dives (226 out of 323 DCS cases were animal data). Lovemann *et al*. identified that the best using Akaike model selection (discussed in Chapter 3) included fitted exponential values. Variation of the hazard function has been shown to have significant effects on model predictions, so we believed an investigation of hazard functions could lead to improved predictive quality in our probabilistic decompression models.

Survival models applied to the prediction of DCS incidents have many advantages over deterministic models, motivating our interest in researching methods for advancing them. Probabilistic modeling has greater flexibility for comparing the predictions of different models by the use of information theoretical tools like the Akike Information Criterion and by the use of classical statistical methods. Parameterization of the models can be performed using likelihood maximization, allowing information in experimental data to be used for model predictions. Additionally, the flexibility of allowing organizations the ability to quantify risk is a significant advantage for mission planning. Probabilistic models possess many advantages for both divers and researchers and are the focus of the model explorations presented in this document.

3. An Information Theoretic Comparison of DCS Hazard Functions

3.1 Comparison of Hazard Function Performance with Identical Exponential Compartment Scaling Terms

We believe that through augmentation of the hazard function we may be able to get greater predictive quality in simple probabilistic models. If successful, we may achieve similar performance to more complex models. As stated earlier in Chapter 1, one of our motivations was to determine if computationally simplistic models could be modified to attain similar LL performance to more complex models. In Chapter 3, we explore seventeen hazard function variations on the probabilistic EE dive model developed by Thalmann *et al*. [12] to determine if any form of the hazard function can result in predictive performance similar to or better than more functionally complex probabilistic DCS models. In Section 3.1, we analyze the performance of 10 variations of instantaneous risk functions by maximizing the associated likelihood functions of a 3 compartment EE gas content model [11, 12]. Model performance is evaluated using classical statistical methods and information theoretic (IT) model selection techniques. In Section 3.2, we examine six model hazard function variations which allow nonlinear scaling exponent parameters to be individually fitted to each of the three compartments of the EE model. Using the results, we develop an AICc weighted model and evaluate its predictive performance. In this chapter, we identify one variation of the hazard function which appears to have the best predictive value in terms of the AICc weight.

3.1.1 Methods

The EE model is defined by its use of exponential functions to describe gas uptake and elimination in each of three parallel, perfusion-limited, uncoupled compartments. Figure 2 shows a graphical representation of the EE model, where P_{T_A} is the inert gas input from the environment, P_{T_i} is the inert gas tension in the i^{th} tissue compartment, T_i is the tissue time, and P_{T_A} is the inert gas wash out from the tissue compartments.

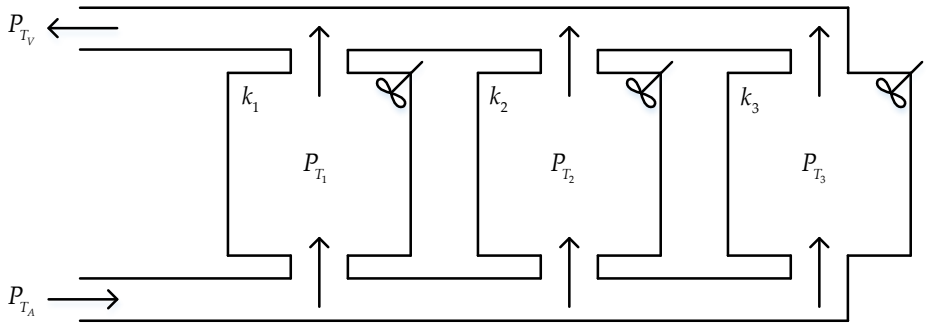

Figure 2: Diagram for the EE model with three parallel, independent compartments. Dissolved gas originates from the P_{T_A} side of the system and is washed out on the P_{T_V} side.

The instantaneous inert gas content for each compartment can be described by the differential system

$$\frac{dP_{T_i}}{dt} = k_i\left(P_{T_A} - P_{T_i}\right) \tag{10}$$

where k_i is the tissue rate (the inverse of the tissue time). Generally, the tissue rate is identified through maximum log-likelihood optimization on a training data set. The perfusion-limited assumption and the exponential-only kinetics of the EE model simplifies computational requirements, however it has been argued that the formation of bubbles in real body tissue slows gas kinetics substantially and exponential kinetics would inaccurately model gas exchange [9]. Several alternative gas kinetic models have been developed to allow for changing gas kinetics due to the formation of bubbles including the linear-exponential model [12]. We decided to explore variations of the hazard function using the EE model because of its mathematical simplicity and its low computation cost. These two features allowed us to explore the functions more rapidly and with a reasonable number of fully optimized parameters for each model to ensure good parametric quality (we obtained at least 32 solutions per model). Additionally, the EE model has a small number of fitted parameters, which reduced the possibility of model overfitting.

3.1.1.1 Hazard Functions

The mathematical form of the ten risk function definitions evaluated in this work are shown in Table 1 below. Half of the hazard functions included threshold terms and half did not. Threshold parameters adjust the level of supersaturation beyond which risk accumulates. Model names ending in a T (e.g. DST) indicate that a threshold term *Thr* was included; models without the T designation (e.g. DS) fixed the threshold term at 0 so it

would not affect model predicted instantaneous risk. Based on previously published results, models including a threshold parameter only included the term in the differential equations governing the third tissue compartment [11], which has a slower rate of gas transfer than the other two compartments.

The hazard functions for all models are shown in Table 1. The Depth Saturation (DS) and Depth Saturation Threshold (DST) models are treated as null models, simply implementing supersaturation as a hazard function [10]. Expanding on the functional definitions of DS and DST, the Depth Saturation Exponent (DSE) and Depth Saturation Exponent Threshold (DSET) hazard functions allow for nonlinear risk accumulation by applying a fitted exponent term to the supersaturation [24]. The Depth Saturation Ratio (DSR) and Depth Saturation Ratio Threshold (DSRT) models employ the supersaturation ratio, which as discussed in Chapter 2, scales the risk by the ambient pressure [4, 12, 17-19, 25]. Depth Saturation Ratio Exponent (DSRE) and Depth Saturation Ratio Exponent Threshold (DSRET) models apply a fitted exponent term to the DSR and DSRT hazard functions [23, 25]. Depth Saturation Ratio Exponent Numerator (DSREN) and Depth Saturation Ratio Exponent Numerator Threshold (DSRENT) only apply the fitted exponent term to the numerator of the hazard function. In all models using a fitted exponent, the exponent is identical for each of the three compartments. In Section 3.2, we examine the effect of allowing individually fitted exponents.

Table 1: Section 3.1 hazard functional forms. Models in this group which had fitted exponents used identical exponent values for each compartment.

Model	Functional Form
DS	$h_i = P_{T_i} - P_{amb} + P_{FVG}$
DST	$h_i = P_{T_i} - P_{amb} - Thr + P_{FVG}$
DSE	$h_i = \left(P_{T_i} - P_{amb} + P_{FVG}\right)^n$
DSET	$h_i = \left(P_{T_i} - P_{amb} - Thr + P_{FVG}\right)^n$
DSR	$h_i = \frac{P_{T_i} - P_{amb} + P_{FVG}}{P_{amb}}$
DSRT	$h_i = \frac{P_{T_i} - P_{amb} - Thr + P_{FVG}}{P_{amb}}$
DSRE	$h_i = \left(\frac{P_{T_i} - P_{amb} + P_{FVG}}{P_{amb}}\right)^n$
DSRET	$h_i = \left(\frac{P_{T_i} - P_{amb} - Thr + P_{FVG}}{P_{amb}}\right)^n$
DSREN	$h_i = \frac{\left(P_{T_i} - P_{amb} + P_{FVG}\right)^n}{P_{amb}}$
DSRENT	$h_i = \frac{\left(P_{T_i} - P_{amb} - Thr + P_{FVG}\right)^n}{P_{amb}}$
i : Compartment	P_{FVG} : Fixed venous gas
P_{T_i} : Tissue pressure	n : Fitted exponent
P_{amb} : Ambient pressure	Thr : Tissue threshold

3.1.1.2 Optimization Methods

Optimization of the LL function for each model was performed using the Nelder-Mead (NM) optimization routine in the Extreme Optimization software package (ver. 4.1) and a simulated annealing (SA) routine we developed based off of the algorithm described

in 'Numerical Recipes in C' [26]. NM is a gradient-free simplex optimization method that can be used to find a maximum point on a surface, in this case the maximum likelihood value. The approach cannot guarantee convergence to a global maximum, but it has been shown to provide a good balance between convergence speed and average solution quality [11, 22, 26, 27]. Gradient free optimization methods have previously been shown to perform well on the log likelihood function, which can be locally discontinuous (explored further in Chapter 5). Parameters for the NM algorithm were considered to be converged when either the difference between the best and worst LL in the region were within a set tolerance (10^{-4} LL) or when the size of the uncertainty of the optimal parameter set was within a set tolerance (10^{-4} for all parameters). We initially employed only the NM algorithm for model optimization, however, examination of the Hessian matrix revealed many of the optimized parameter sets to be suboptimal.

To improve parameter LL quality we employed SA, which has more robust, but slower, convergence properties [17, 28-30]. SA is a modified version of the Metropolis-Hastings Algorithm described in Kirkpatrick *et al.* [30]. SA has been shown to perform well under a variety of problems, including parameter optimization of DCS predictive models [15, 17]. Convergence for the SA algorithm was determined when no adjacent parameter set was more than 10^{-4} LL units better than the optimized parameter set. All optimized model parameters were drawn from normal distributions. The mean and variance for each distribution were set by examination of 256 randomly distributed

parameter sets for each model. Of those 256 parameter sets, the top 32 (by LL ranking) were used to set each of the variables for the normal distribution for each parameter.

3.1.1.3 Data

The BIG292 data set, comprised of 3322 individual man dives in 1038 dive profiles with 190 DCS incidents, was used as training data for model optimization [31]. It is important to note that of the 1038 dive profiles, 110 were recorded as marginal DCS events and have often been published with δ outcome indicator values between 0 and 1 (as shown in Eq.(6)). In the current study, marginal DCS events were classified as non-events, as earlier investigations have demonstrated that the inclusion of marginal DCS incidents does not significantly affect the quality of model predictions [31].

For model analysis, profiles were grouped by dive type as well as maximum depth, as it has been suggested that the effects of dive depth on DCS risk may be nonlinear [25]. Dive data was organized into four groups by depth, shown in Figure 3. BIG292 contains six defined dive types (Saturation, Sub-Saturation, Single Air, Single Non-Air, Repetitive Air, and Repetitive Non-Air), and model performance was also investigated using the separation of profiles by dive type.

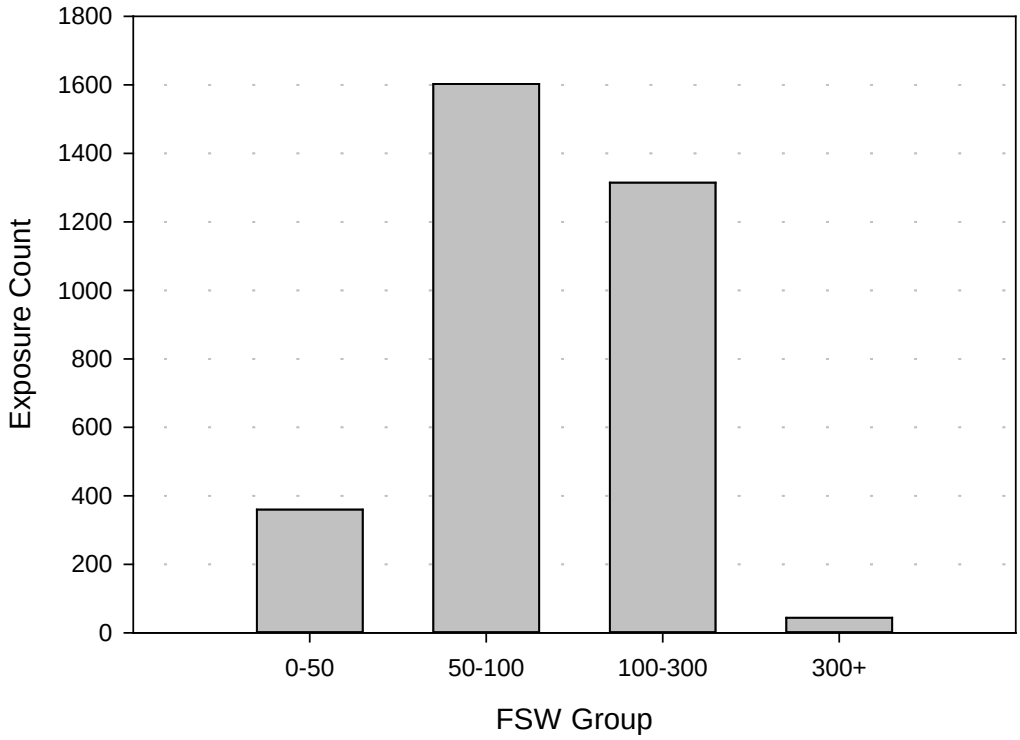

Figure 3: Maximum Depth for BIG292. Depth groups were set as 0-50, 50-100, 100-300, and 300+ feet of sea water (FSW)

3.1.1.4 Model Comparison

Model performance was evaluated using the Pearson Xi-square test on the entire data set, by dive type, and by the depth groups defined above [16]. Traditional likelihood ratio testing is not an appropriate metric for model performance comparison because the hazard functions are non-nested. Parameter confidence intervals were generated using the standard error approximation

$$s.e.\left(\hat{\beta}_i\right) = \mathrm{cov}\left(\hat{\beta}_i, \hat{\beta}_i\right)^{\frac{1}{2}} \tag{11}$$

where $\hat{\beta}_i$ is the i^{th} parameter of the maximum likelihood model parameter set $\hat{\beta}$ [16, 24]. To avoid issues related to the non-nested structure of the models, model selection was performed using the Akaike Information Criterion (AIC),

$$AIC = 2k - \ln(L) \tag{12}$$

which is an information theoretic tool that can be used to evaluate performance of a model performs with *k* fitted parameters and a maximized likelihood value *L* [32, 33]. Disparate model structures can be compared through the calculation of the AIC weight, generated by taking the difference between the AIC of each candidate model and the lowest AIC value such that

$$\Delta_i\left(AIC\right) = AIC_i - \min AIC \tag{13}$$

Each AIC difference $\Delta_i\left(AIC\right)$ is used to generate the relative likelihood of each model, defined as

$$\psi_i = \exp\left(-\frac{1}{2}\Delta_i\left(AIC\right)\right) \tag{14}$$

which is used to generate Akaike weight $w_i\left(AIC\right)$ by normalizing each relative likelihood ψ_i by the sum of the of relative likelihoods from all M models explored

[22]

$$w_i\left(AIC\right)=\frac{\exp\left(-\frac{1}{2}\Delta_i\left(AIC\right)\right)}{\sum_{m=1}^{M}\exp\left(-\frac{1}{2}\Delta_m\left(AIC\right)\right)} \tag{15}$$

where $\sum_{m=1}^{M}\exp\left(-\frac{1}{2}\Delta_m\left(AIC\right)\right)$ is the sum of the Akaike weights for all the M models in the comparison. AIC weights sum to 1 and can therefore be thought of as percentages of the total evidence ratio (we will use AIC weight and percent interchangeably) [34]. Due to the low ratio of decompression incidents to the number of fitted parameters in all models, we employed AICc, a modified form of AIC which accounts for finite data samples, defined as

$$AICc=2k-\ln(L)+\frac{2k(k+1)}{s-k-1} \tag{16}$$

where s is the total number of observations in the data set. As s becomes large, the correction term of the AICc diminishes and so the value of AICc approaches the value of AIC. AICc can be to generate the AIC differences $\Delta_i\left(AIC\right)$, relative likelihoods ψ_i and AIC weights $w_i\left(AIC\right)$.

3.1.2 Results & Discussion

3.1.2.1 Optimization Results

Optimized model parameters and confidence intervals for the 10 models and are displayed in Table 2. Results from the NM algorithm were evaluated for optimality by determining the invertibility of the Hessian matrix of the optimized parameter set. For

several of the parameter sets, the Hessian was noninvertible, indicating the solution was not at a local maxima of the LL hyper-surface. This was not entirely surprising as the NM algorithm has been found to have good average functional improvement but may have poor convergence properties [26, 35]. Interestingly, the scale of the exponential parameters tended to be near one, with the 95% parameter confidence interval encompassing $n = 1$ for all models except DSREN. This implies that only a small nonlinearity was useful for maximizing LL.

Table 2: Optimized LL hazard function model parameters

Model	DS	DST	DSE	DSET	DSR	DSRT	DSET	DSRE	DSRET	DSREN	DSRENT
T_1 (min)	1.06 ± 1.15	1.01 ± 0.90	1.13 ± 1.05	1.27 ± 1.84	1.67 ± 1.70	1.70 ± 1.53	1.27 ± 1.84	1.78 ± 1.58	1.74 ± 1.28	1.65 ± 1.62	3.02 ± 4.16
T_2 (min)	233.43 ± 66.41	222.85 ± 51.29	243.30 ± 70.48	224.29 ± 62.36	222.71 ± 73.13	215.06 ± 48.80	224.29 ± 62.36	212.76 ± 113.78	217.24 ± 58.08	208.23 ± 116.33	228.78 ± 55.97
T_3 (min)	483.94 ± 280.99	519.92 ± 141.80	512.01 ± 367.14	519.50 ± 176.53	495.31 ± 291.12	547.13 ± 79.62	519.50 ± 176.53	482.44 ± 196.90	548.54 ± 88.02	461.57 ± 159.82	565.97 ± 101.44
g_1 (min^{-1})	3.9 E-3 ± 6.71 E-3	4.2 E-3 ± 6.18 E-3	4.33 E-3 ± 6.86 E-3	2.94 E-3 ± 6. E-3	3.09 E-3 ± 4.5 E-3	3.02 E-3 ± 4.01 E-3	2.94 E-3 ± 6. E-3	3.77 E-3 ± 5.25 E-3	2.78 E-3 ± 3.07 E-3	4.08 E-3 ± 6.51 E-3	1.55 E-3 ± 2.52 E-3
g_2 (min^{-1})	5.03 E-4 ± 3.05 E-4	4.74 E-4 ± 2.04 E-4	7.7 E-4 ± 6.49 E-4	5.04 E-4 ± 4.48 E-4	4.85 E-4 ± 3.68 E-4	4.76 E-4 ± 1.58 E-4	5.04 E-4 ± 4.48 E-4	6.76 E-4 ± 7.69 E-4	4.48 E-4 ± 3.17 E-4	7.37 E-4 ± 9.3 E-4	5.94 E-4 ± 4.69 E-4
g_3 (min^{-})	1.74 E-4 ± 2.88 E-4	3.67 E-4 ± 3.11 E-4	2.66 E-4 ± 5.69 E-4	3.64 E-4 ± 3.44 E-4	2.43 E-4 ± 3.55 E-4	6.29 E-4 ± 3.33 E-4	3.64 E-4 ± 3.44 E-4	6.91 E-4 ± 1.08 E-3	5.99 E-4 ± 4.1 E-4	9.37 E-4 ± 1.19 E-3	6.83 E-4 ± 4.61 E-4
Thr (ATM)	-	77.2 E-3 ± 4.6 E-2	-	71.5 E-3 ± 4.02 E-2	-	101.1 E-3 ± 2.46 E-2	71.5 E-3 ± 4.02 E-2	-	103.9 E-3 ± 2.42 E-2	-	97.4 E-3 ± 2.97 E-2
n	-	-	1.28 ± 0.27	1.03 ± 0.36	-	-	1.03 ± 0.36	1.39 ± 0.36	0.973 ± 0.33	1.55 ± 0.31	1.12 ± 0.37
P_{DCS}	189.85 ± 25.99	187.7 ± 25.6	189.46 ± 25.86	189.35 ± 25.82	190.08 ± 26.09	190.24 ± 25.97	189.35 ± 25.82	190.02 ± 26.01	186.96 ± 25.58	189.43 ± 25.88	189.42 ± 25.85
LL	-971.02	-967.16	-968.98	-967.18	-976.87	-967.7	-967.18	-974.26	-967.72	-970.57	-967.78

We found that the addition of SA optimization upon the NM optimized solution was, in several cases, advantageous for improving LL model fit and yielded parameter sets which were locally optimal. Improvements and AICc values can be seen in Table 3. We remind the reader that the data set used in parametric optimization considered marginal DCS outcomes as non-events, which results in the LL values reported in Table 3 being not directly comparable to other published models using the data set. Classification of marginal events as non-events will artificially improve the reported LL compared to other models. In the sections below, model selection using AICc weighting and model performance analysis using Xi square testing is presented. By AICc weighting, the DST model was the best individual performing model, while the best model according to Xi square results was inconsistent and included DSE for the full group, and DS, DST, DSE, and DSET for the depth groups 0-50 and 50-100 feet of sea water (FSW).

Table 3: Hazard Function AICc Values and SA LL Improvements

Model	Best	+Δ LL NM-SA
DS	-976.28	0.41
DST	-972.90	0.46
DSE	-974.67	0.49
DSET	-971.29	1.63
DSR	-981.59	0.07
DSRT	-973.00	0.01
DSRE	-979.45	0.61
DSRET	-973.03	1.61
DSREN	-975.97	0.74
DSRENT	-969.06	1.69

3.1.2.2 Information Theory

Using the maximum LL identified, AICc weights were calculated for each model. In the case where one model has most of the informative weight (>= 0.9), inference can be made using only that model, as it is the most parsimonious with the training data [36]. Alternatively, if a small group (preferably fewer than 2 in 10, a standard metric used in the field) of model AICc weights combine to form greater than 90% of the total, multimodel inference (MMI) could be used to generate predictions by combining the PDCS of several models (MMI is covered in Section 3.2) [36]. Results from the AICc weighting were inconclusive, indicating no clear model outperformed competitors shown in Table 4 and graphically depicted in Figure 4. Additionally, no combination of two models accounts for more than 60% of the overall evidence. DST had the largest AICc weight, but with a value of less than 40% of the total, model selection based on AICc would be of limited certainty. Including a threshold parameter increased the evidence

ratio in each model, which indicates that the additional parameter in each model improves its explanatory power for the outcomes of BIG292.

Table 4: AICc values and weights for models explored in Section 3.1. DST and DSET had the greatest weight, but no model overall was selected by AICc weighting.

Model	AICc	AICc Weight
DS	-971.02	0.022
DST	-967.16	0.384
DSE	-968.98	0.062
DSET	-967.18	0.137
DSR	-976.86	>0.000
DSRT	-967.70	0.225
DSRE	-974.26	>0.000
DSRET	-967.72	0.081
DSREN	-970.57	0.013
DSRENT	-967.78	0.076

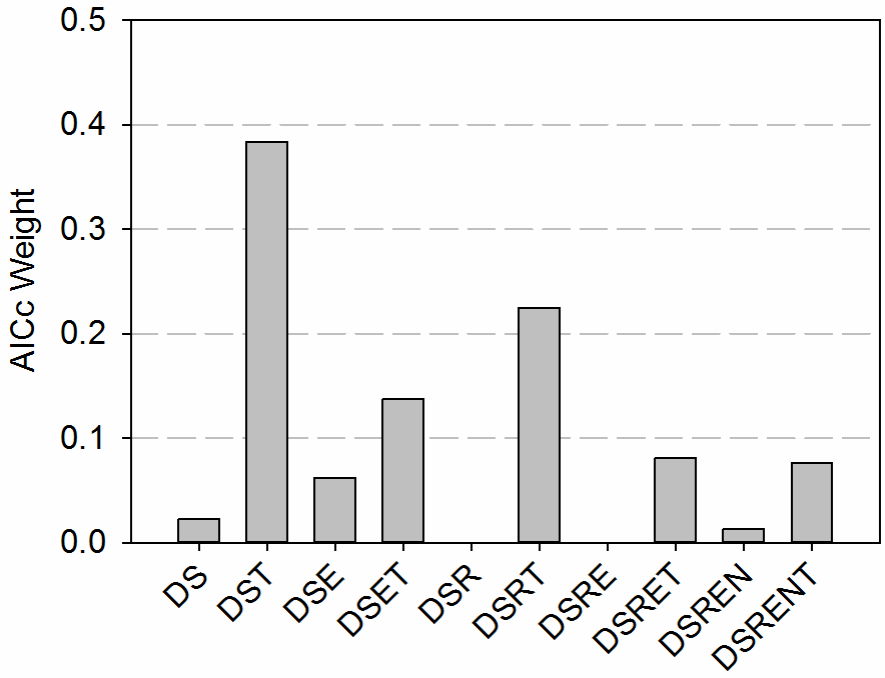

Figure 4: AICc weighting of the models. DST accounted for the largest percentage but did not account for more than 90% of the overall evidence. The results indicate that no single model could be selected as the overall likeliest model.

Given the small number of models in the explored set, we were surprised to find that no one model or small combination of models had the majority of the AICc weight. A smaller group of models will emphasize differences in AIC performance by dividing the total evidence ratio in larger proportions. The combination of AICc weights from all threshold models was greater than 90%, but the combination of 5 models out of 10 is unsatisfactory from a model selection perspective, as more than 2 in 10 models is often considered inconclusive for model selection [36]. The higher AICc weights associated with

all threshold models compared to their respective bases indicates that the additional fitted parameters are justified. This was not unexpected as the addition of threshold parameters has been found to be statistically justified in the past [12]. Because AICc model selection failed to indicate a best performing model, we elected to include additional model parametric freedom by allowing for individual fitting for compartmental exponent parameters (discussed in Section 3.2).

3.1.2.3 Group Xi Square Results

Xi square testing results for each model compared to the entire BIG292 data set are displayed in Figure 5. By group Xi square testing, no model predictions grouped by dive type were consistent with observations. The predicted risk from the DSE model was most closely associated with the observed incidents with a Pearson statistic of 0.4. DSR deviated from observations the most with a Pearson statistic of 0.05. Interestingly, the relative error of total predicted DCS from all models was no greater than 0.02% compared to the 190 incidents of the training data, but no model performed well by group Xi square testing. Poor performance by every model resulted from inaccurate risk distribution by dive category, which could indicate a limitation of the EE model type to properly fit the dive data. The DSR hazard function model deviated the most from the observed data in BIG292, suggesting that other hazard functional forms provide superior performance with respect to the distribution of DCS risk by dive type.

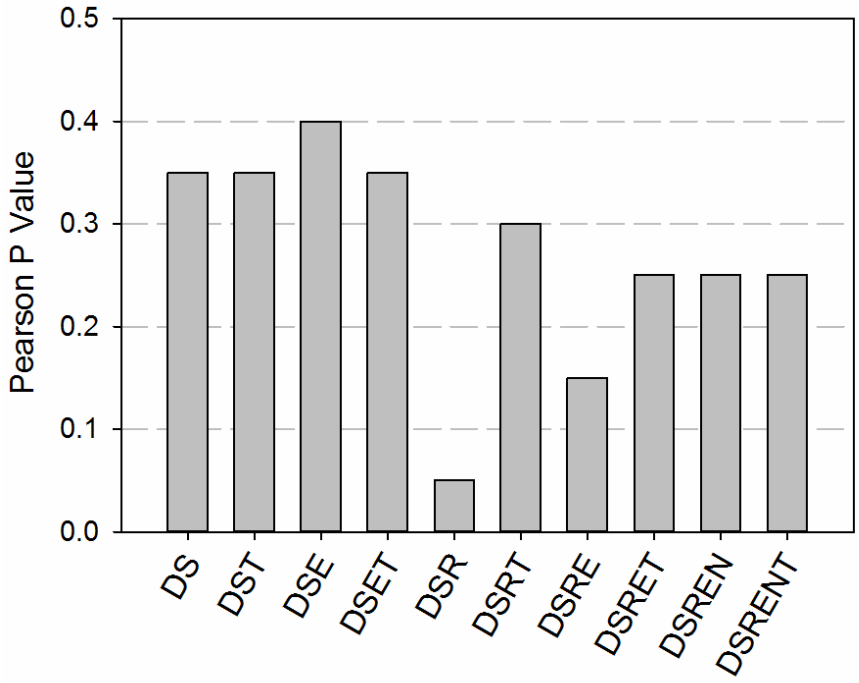

Figure 5: Group Xi Square p-value results for all models tested across all dive profiles in BIG292. DSE had the closest correlation with BIG292, while DSR was the least correlated.

3.1.2.4 Comparison by Dive Type

Figure 6 graphically displays the Xi square results by individual dive type with the number of degrees of freedom set to the number of experimental trials. Total predicted DCS incidents within each dive type was compared to the total observed incident count in each data set. In Single Air dives, no model Pearson test statistic exceeded 0.9, with DSR achieving the highest value of 0.8. For Repetitive Air profiles, all models except DSR exceeded 0.9. For Single Non Air, no model exceeded a Pearson statistic of 0.4. For

Repetitive Non Air, only DST and DSRET exceeded 0.9. For Saturation, all models except for DS, DSR, and DSRE exceeded 0.99. For Sub-Saturation, no model exceeded 0.55.

Hazard function performance was found to be inconsistent based on the results from Xi square tests. While DSR achieved the largest Pearson test statistic on Single Air dives, for all other dive types it was the weakest performer (with the exception of Sub-Saturation, for which no model performed well). Some models predicted well for some dive types but poorly on others, e.g. DSR's prediction was most closely associated with Single Air dives but the least correlated with Saturation dive types.

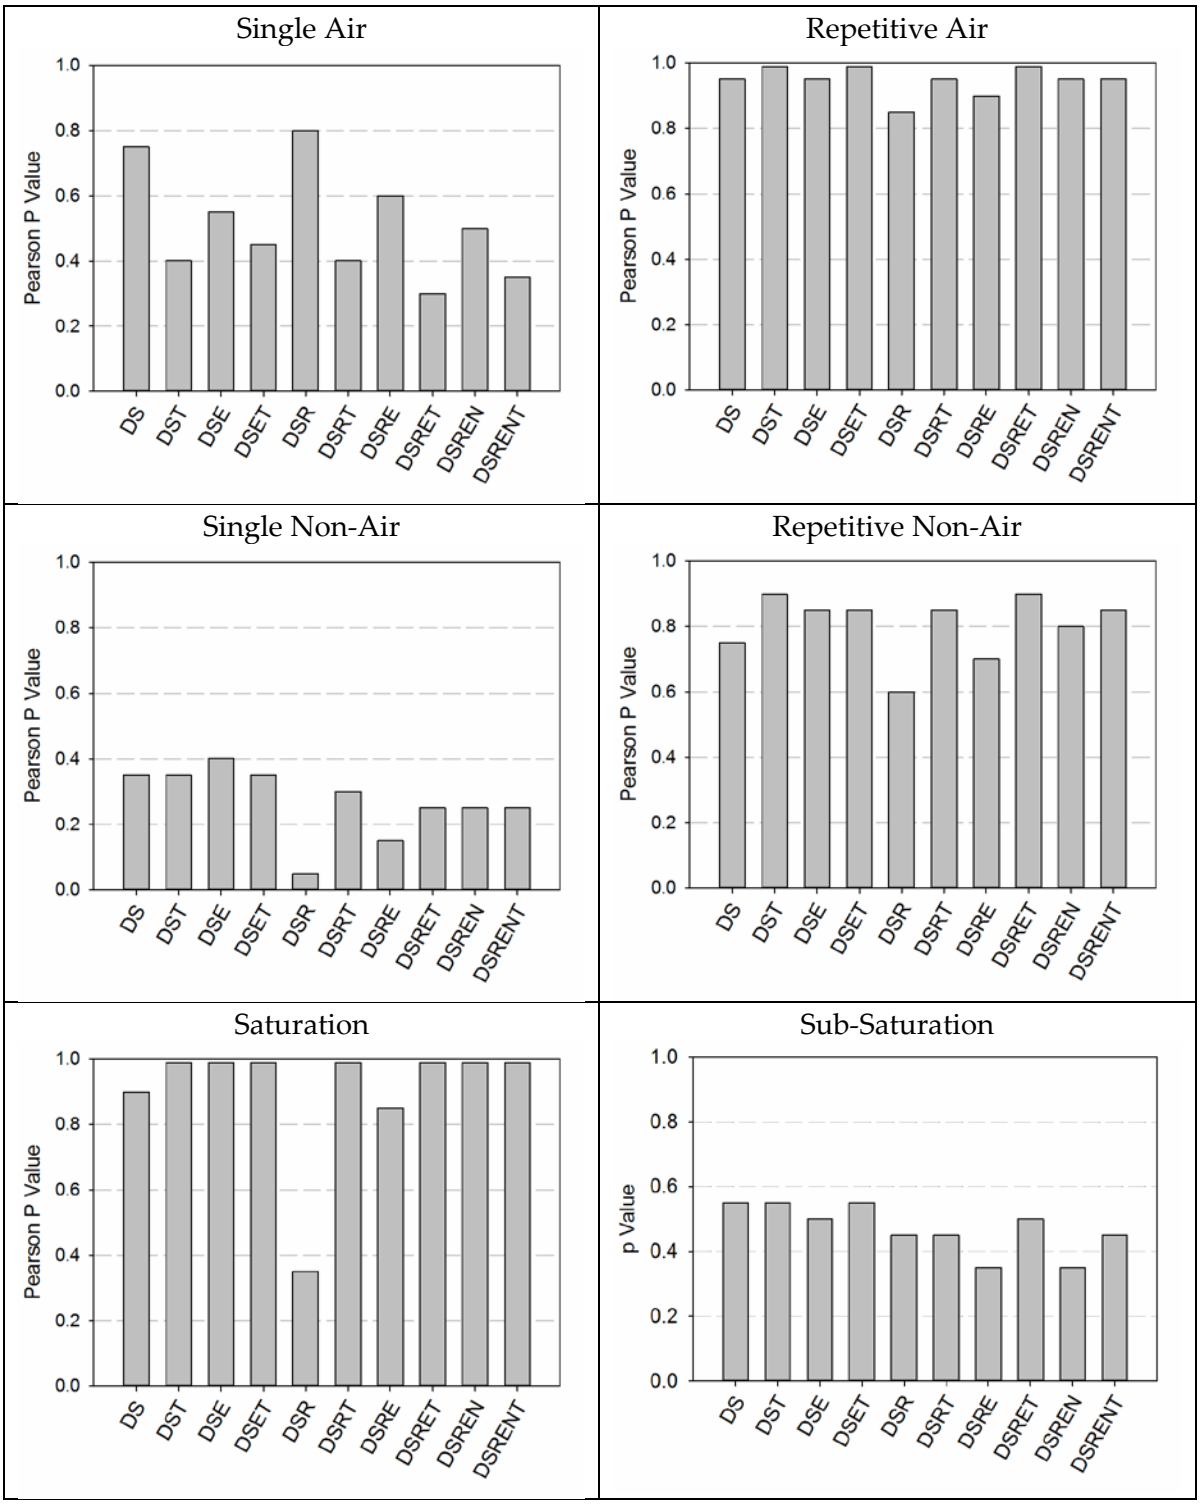

Figure 6: Xi square tables for models explored in Section 3.1 by dive type

3.1.2.5 Comparison by Dive Depth

Model-predicted DCS compared to the observed cases segmented by depth can be found in Table 5**.** All models over predicted DCS on profiles with maximal depths from 0-50 FSW and 50-100 FSW compared to observed DCS cases. All models under-predicted DCS for profiles with maximal depths of greater than 100 FSW. In depth categories for 0-50 and 50-100 FSW, the DSR and DSRT models were outperformed by several alternative model forms based on the Pearson test statistic including DS, DST, DSE, and DSET. Overall, no functional form performed best on all depth categories.

Table 5: Model predicted DCS count by maximum profile depth. All model variants overpredicted the number of incidents for profiles with maximum depths shallower than 100 FSW and underpredicted the number of cases for profiles deeper than 100 FSW.

Depth / Model (DCS)	0-50 FSW (18)	50-100 FSW (88)	100-300 FSW (82)	300+ FSW (2)	Total (190)
DS	26.43	110.46	52.4	0.56	189.85
DST	26.69	108.24	52.18	0.59	187.7
DSE	27.21	109.25	52.29	0.71	189.46
DSET	26.95	109.21	52.73	0.46	189.35
DSR	28.58	111.15	49.86	0.5	190.08
DSRT	30.49	109.99	49.27	0.49	190.24
DSRE	31.29	110.22	47.86	0.65	190.02
DSRET	29.89	108.15	48.47	0.45	186.96
DSREN	31.02	109.39	48.28	0.74	189.43
DSRENT	30.44	109.35	49.36	0.27	189.42
Model Average	**28.9**	**109.5**	**50.3**	**0.54**	**189.25**

No Xi square test statistic for comparison between DCS incidents in dive profiles segmented by maximum dive depth and model-predicted DCS exceeded a value of 0.95 in any depth category (see Figure 7). Model performance was similar in all four depth categories, with a maximal disparity of Pearson test statistic of 0.4, with the exception of the 100-300 FSW range which had poor performance from every model. Interestingly, the models that did not scale the supersaturation by the ambient pressure outperformed the models which did for the 0-50 FSW depth group. This could suggest that for shallower maximal depth pressure exposures scaling by the ambient pressure could lead to slightly worse overprediced risk, as the models without ambient pressure scaling predicted lower risk for profiles in this depth category (closer to observed incidence).

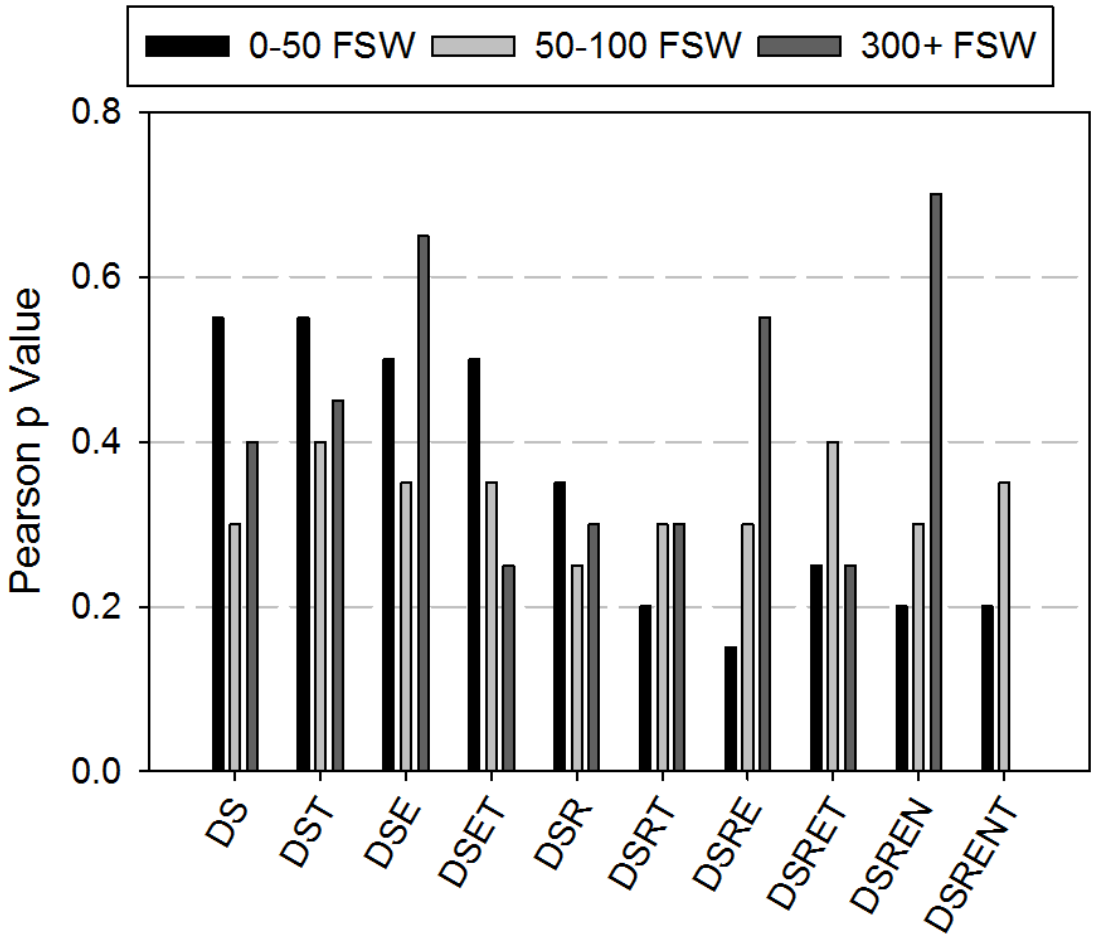

Figure 7: The maximal Pearson Test Statistic for each depth grouping is shown. No model exceeded 0.05 significance level in 100-300 FSW range and it was not included in this graph. Performance across models varied greatly between depth groups, but models without ambient pressure scaling outperformed models with ambient pressure scaling for the 0-50 FSW depth group.

Comparison of Hazard Function Performance with Independent Exponential Compartment Scaling Terms

Our initial exploration of hazard functional forms using the EE model is expanded in this section by allowing the independence of the compartment exponent parameters associated with the DSE, DSET, DSRE, DSRET, DSREN, and DSRENT hazard functions. Allowing independent exponents on each of the tissue compartments may improve performance enough to justify their inclusion, as we found that the addition of both threshold and exponent parameters aided in model performance for some dive type and depth groups in Section 3.1.

AICc weighting did not identify a small subset of models with combined weights exceeding the 90% significance threshold for the models analyzed in Section 3.1. With the addition of the presently explored models, we identified two models with AICc weights that sum to more than 90% of the total. We then explored multimodel inference (MMI) as a means of increasing model robustness. In many cases where a single model is not selected by AICc weighting, combining the predicted DCS from several models can improve prediction quality by combining the data explanatory power of more than one model [36]. Using AICc weighting, an averaged prediction for probability of a decompression injury can be calculated that reflects the estimates of the best models.

3.2.1 Methods

Hazard functions allowing individually fitted exponents were designated DSE111, DSE111T, DSRE111, DSRE111T, DSRE111N, and DSRE111NT, shown in Table 6. Models were optimized using NM and SA and were analyzed using AICc weighting and Xi square testing. Additionally, MMI was used to generate a PDCS estimation comprised of models with a combined AICc weight of at least 0.9.

Table 6: Section 3.2 hazard functions with individually fitted compartmental exponents.

Model	Functional Form
DSE111	$h_i = \left(P_{T_i} - P_{amb} + P_{FVG}\right)^{n_i}$
DSE111T	$h_i = \left(P_{T_i} - P_{amb} - Thr + P_{FVG}\right)^{n_i}$
DSRE111	$h_i = \left(\frac{P_{T_i} - P_{amb} + P_{FVG}}{P_{amb}}\right)^{n_i}$
DSRE111T	$h_i = \left(\frac{P_{T_i} - P_{amb} - Thr + P_{FVG}}{P_{amb}}\right)^{n_i}$
DSRE111N	$h_i = \frac{\left(P_{T_i} - P_{amb} + P_{FVG}\right)^{n_i}}{P_{amb}}$
DSRE111NT	$h_i = \frac{\left(P_{T_i} - P_{amb} - Thr + P_{FVG}\right)^{n_i}}{P_{amb}}$
i : Compartment	P_{FVG} : Fixed venous gas
P_{T_i} : Tissue pressure	n_i : Fitted exponent
P_{amb} : Ambient pressure	Thr : Tissue threshold

3.2.2 Results & Discussion

3.2.2.1 Optimization Results

Optimized model parameters with confidence intervals for the 6 optimized models are shown in Table 7. Parameter confidence intervals, as well as optimized parameters for most models were numerically similar to the maximum LL model parameters identified in Section 3.1. SA improved LL much more substantially for the models with individually fitted exponents (shown in Table 8), averaging an LL improvement of 4.86 compared to an average LL improvement of only 0.77 for the models explored in Section 3.1. Of notable difference between the free exponent models and those explored in Section 3.1 was the difference in the scale of the exponent parameters in this exploration. In 3.1, optimized exponent values were found to be relatively near 1, whereas in the free exponent models the values varied substantially. Compartments 1 and 3 overlapped in the mean of the exponent values, but compartment 2 had very small exponent values, averaging three orders of magnitude smaller than the values found for compartments 1 and 3. This small exponential value causes a nearly flat risk accumulation for any positive supersaturation value within compartment 2, which could reduce model statistical failure (explored in Chapter 5).

Table 7: Optimized LL hazard function model parameters Independent Exponent

Model	DSE111	DSE111T	DSRE111	DSRE111T	DSRE111N	DSRE111NT
T_1 (min)	1.49 ± 1.17	1.54 ± 1.19	2.49 ± 2.68	2.16 ± 2.00	1.62 ± 1.25	1.80 ± 1.46
T_2 (min)	130.55 ± 8.20	128.16 ± 4.64	129.40 ± 25.06	130.90 ± 18.43	131.05 ± 1.23	130.49 ± 19.37
T_3 (min)	403.25 ± 62.51	410.54 ± 1.76	493.90 ± 67.06	493.46 ± 64.41	467.09 ± 8.08	465.05 ± 67.62
g_1 (min^{-1})	4.12 E-3 ± 5.75 E-3	3.92 E-3 ± 5.44 E-3	2.92 E-3 ± 4.95 E-3	3.72 E-3 ± 5.75 E-3	5.12 E-3 ± 6.96 E-3	4.54 E-3 ± 6.79 E-3
g_2 (min^{-1})	5.14 E-5 ± 6.1 E-5	5.27 E-5 ± 7.2 E-5	6.34 E-5 ± 7.6 E-5	7.13 E-5 ± 6.7 E-5	8.52 E-5 ± 1. E-4	7.43 E-5 ± 7.5 E-5
g_3 (min^{-1})	1.33 E-3 ± 6.34 E-4	1.26 E-3 ± 6.39 E-4	4.01 E-3 ± 3.22 E-3	2.87 E-3 ± 2.68 E-3	2.99 E-3 ± 1.62 E-3	2.73 E-3 ± 1.59 E-3
Thr_3 (ATM)	-	19.2 E-3 ± 2.45 E-2	-	46.4 E-3 ± 5.68 E-2	-	13.3 E-3 ± 7.94 E-2
Exp_1	2.06 ± 1.72	2.12 ± 1.73	1.81 ± 1.65	1.77 ± 1.49	2.02 ± 1.74	2.15 ± 1.73
Exp_2	31.1 E-6 ± 0.44	4.7 E-3 ± 0.55	5.4 E-3 ± 0.46	63.3 E-3 ± 0.38	88.7 E-3±0.48	54.5 E-3 ± 0.41
Exp_3	1.76 ± 0.37	1.6 ± 0.41	2.42 ± 0.61	1.88 ± 0.81	2.34 ± 0.47	2.16 ± 0.85
P_{DCS}	188.6 ± 25.69	189.47 ± 5.76	189.99 ± 25.94	189.12 ± 25.78	192.58 ± 6.25	189.98 ± 25.88
LL	-961.2	-961.16	-960.56	-959.52	-957.59	-957.54

Table 8: LL improvement using SA after initial optimization using NM. Several models improved substantially with LL reductions of more than 7 units.

Model	Best	+Δ LL NM-SA
DSE111	-968.18	1.22
DSE111T	-967.92	3.89
DSRE111	-965.72	7.66
DSRE111T	-964.66	7.37
DSRE111N	-963.29	1.55
DSRE111NT	-963.42	7.50

3.2.2.2 AICc Model Selection

AICc weights were calculated using both the current 6 models and the 10 models from Section 3.1, the results of which are shown in Table 9 and graphically in Figure 8. Calculated AICc results revealed that only DSRE111N and DREN111NT were needed to obtain more than 90% of the evidence weight for the AICc method. This shows that the addition of unique exponents on each tissue is a statistically justified modification. It is interesting that only one base model type was selected by the AICc weighting method for use in MMI, as there was no indication of a favorite model in the base investigation. Due to the close formulation of the two selected models, we checked the AICc weight distribution by removal of the DSRE111NT model. In this instance, DSRE111N accounted for 88% of the total evidence ratio. In the inverse case, removal of DSRE111N resulted in DSRE111NT accounting for only 78% of the total evidence ratio. This indicates that the EE model optimized using a hazard function of the general formulation of DSRE111N has superior predictive performance than did the other models in this investigation.

Table 9: AIC weights for all models explored in Chapter 3. DSRE111N and DSRE111NT both significantly outperformed the other models in this exploration.

Model	Best LL	AICc Weight
DS	-971.02	>0.000
DST	-967.16	>0.000
DSR	-968.98	>0.000
DSRT	-967.18	>0.000
DSE	-976.87	>0.000
DSET	-967.70	>0.000
DSRE	-974.26	>0.000
DSRET	-967.72	>0.000
DSREN	-970.57	>0.000
DSRENT	-967.78	0.004
DSE111	-961.20	0.018
DSE111T	-961.16	0.007
DSRE111	-960.56	0.033
DSRE111T	-959.52	0.034
DSRE111N	-957.59	0.651
DSRE111NT	-957.54	0.251

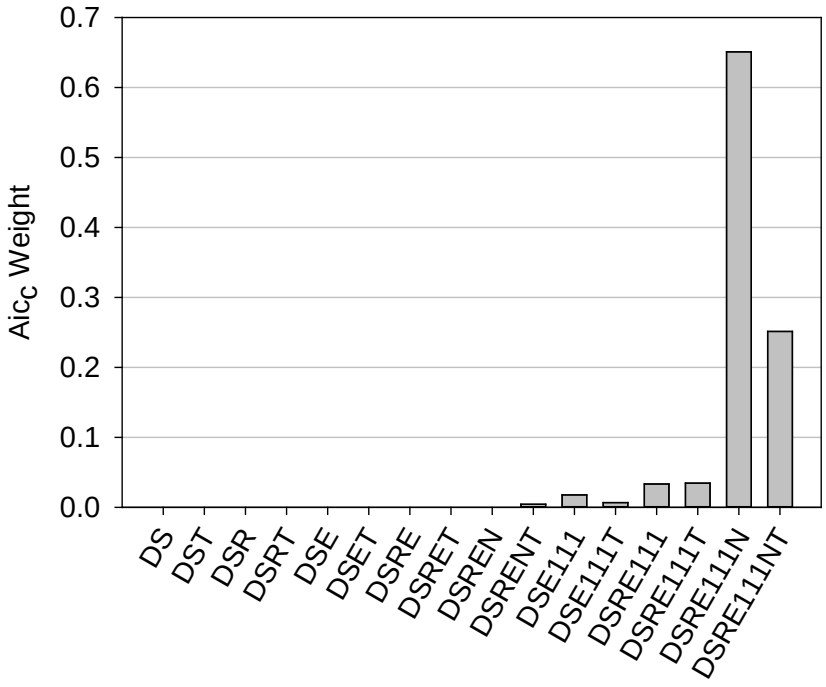

Figure 8: AICc weights for the initial 16 models tested. The two best models, DSRE111N, and DSRE111NT made up more than 90% of the AICc weight for the group. DSRE111N is the performance leader by a wide margin, but alone does not achieve the +90% of total evidence weight recommended for confident inference.

The two models selected by AICc weighting, DSRE111N and DSRE111NT, were used to create multimodel weighted average PDCS that distributed risk similarly to many of the other models by both dive type and depth. Using only the two selected models, the AICc weights were normalized (shown in Table 10) and used for predictive inference by combining the weighted risk predictions for each profile such that

$$P_{DCS}^{AIC} = 0.722\left(P_{DCS}^{\varepsilon}\right) + 0.278\left(P_{DCS}^{\nu}\right) \tag{17}$$

where P_{DCS}^{AIC} is the predicted DCS risk for each profile, P_{DCS}^{ε} is the PDCS for DSRE111N, and P_{DCS}^{ν} is the PDCS for DSRE111NT.

Table 10: Normalized AICc weights for DSRE111N and DSRE111NT used for the AICc weighted multimodel. DSRE111N accounted for over twice as much of the evidence ratio as DSRE111NT, but not enough to justify model selection by itself.

Models	Initial AICc Weight	Normalized Multipliers
DSRE111N	0.651	0.722
DSRE111NT	0.251	0.278

Xi square results for the AICc multimodel were comparable to other models in the base and free exponent groups. The AICc multimodel did not outperform the two models from which it was formed in any Xi square comparison, which was not surprising as the two models composing the multimodel had similar performance for all Xi square tests.

3.2.2.3 AICc Multimodel Performance

The predicted risk of all test models was compared to the DSR model as a baseline because of its simple form for scaling risk accumulation by the ambient pressure, a common technique used in probabilistic decompression models. While most risk predictions for profiles in the BIG292 data set were similar to the predictions of DSR, profiles deeper than 300 FSW and Saturation dive types had noticeably larger PDCS compared to DSR. Figure 9 illustrates these features in the forms of the Saturation and 300+ FSW data groups using the AICc multimodel for comparison. Models that predicted higher PDCS for deep and saturation dives performed better in Pearson testing and in AICc selection compared to the DSR model. DSRT, DSRE, DSRET, and DSRE111 did not

differ significantly from DSR in both 300+ FSW and Saturation dive profiles. The difference in predicted risk between the AICc multimodel and DSR implies that DSR under-predicts the risk of saturation and 300+ FSW dive profiles, since the AICc multimodel outperformed the DSR model in both categories.

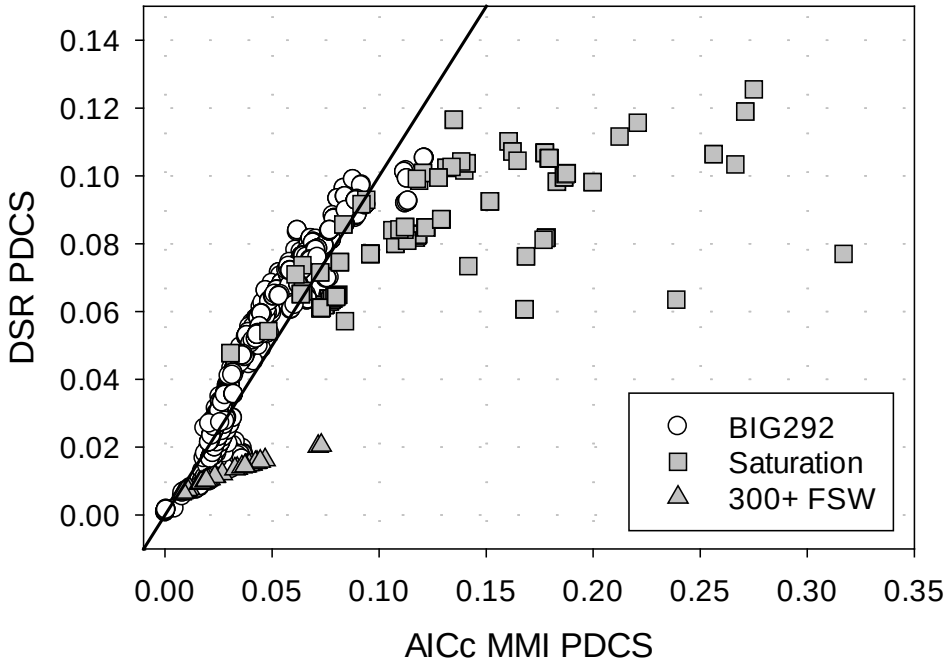

Figure 9: PDCS comparison for AICc MMI model vs DSR model. The AICc MMI predicted greater risk on 300+ FSW and Saturation profiles.

The risk accumulated by each compartment for each profile is shown for both DSR and the AICc multimodel in the ternary plot in Figure 10. The data points were generated by normalizing each individual integrated compartment risk by the total accumulated risk for each profile in the training data set, demonstrated functionally as

$$C_r^i = \frac{P_{DCS}^i}{\sum_{n=1}^{N} P_{DCS}^n} \tag{18}$$

where C_r^i is the normalized risk accumulated by compartment i, P_{DCS}^i is the total risk accumulated by compartment i, and N is the total number of compartments in the system. As shown in the figure, the DSR model accumulated risk across two or three compartments. The profiles which had risk accumulation from more than one compartment were primarily Single Air exposures. The accumulation of risks across compartments in the DSR model suggests a greater deal of overlap of the compartments than in the AICc multimodel. Of interest is the accumulation of risk by only one compartment for many of the AICc predictions. These profiles, indicated in the bottom left of the figure, were all submarine escape profiles. These saturation-drop out profiles have different risk accumulation properties than the majority of the training data set due to their short decompression time.

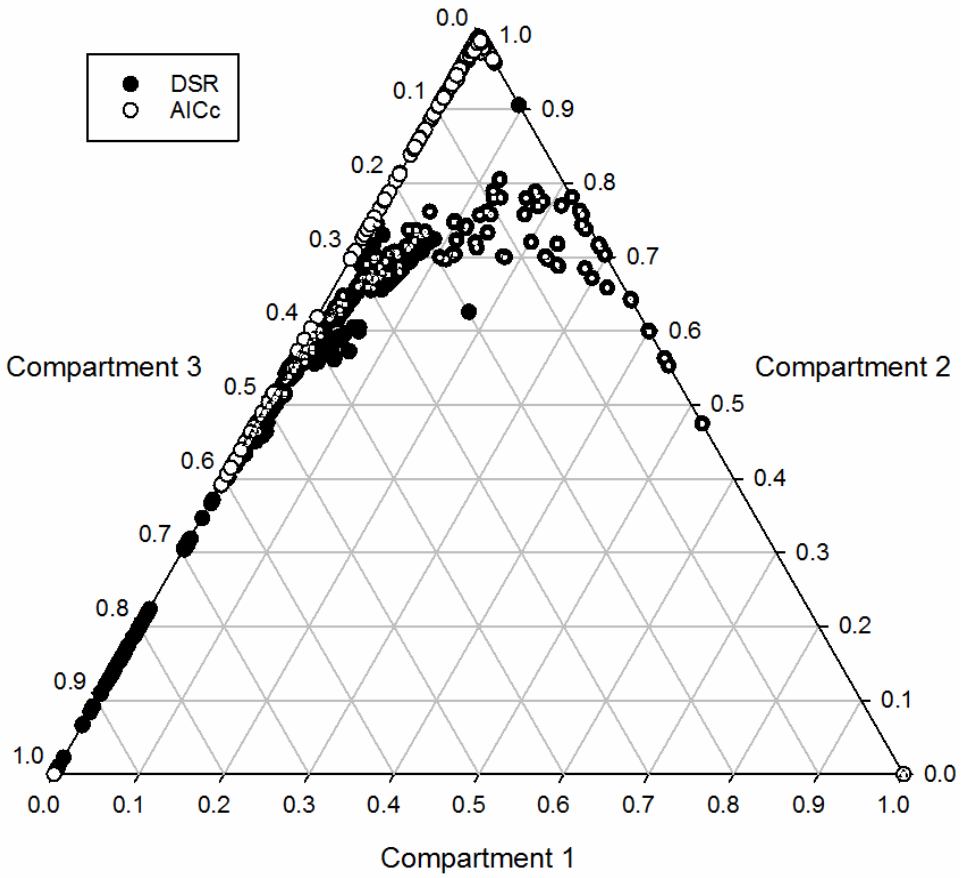

Figure 10: DSR vs AICc Model ternary plot. Risk spread between compartments two and three on the DSR model is largely associated with the Single Air dive category (indicated by a white dot).

3.2.2.4 Group Xi Square Results

Group Xi Square testing results can be seen in Figure 11, which includes the models from Section 3.1. The greatest overall Pearson testing statistic of 0.4 was achieved by DSE111, matching the value from the DSE model. The addition of a threshold term improved Xi square results in all models with the exception of DSE111. No hazard

function variant achieved a significant correlation with the training data, likely indicating a limitation of the underlying mechanics of the EE model for achieving predictive correlation with the entire data set. It was unexpected that the DSRE111N and DSRE111NT models did not outperform the other models in the set because they were selected by AICc weighting.

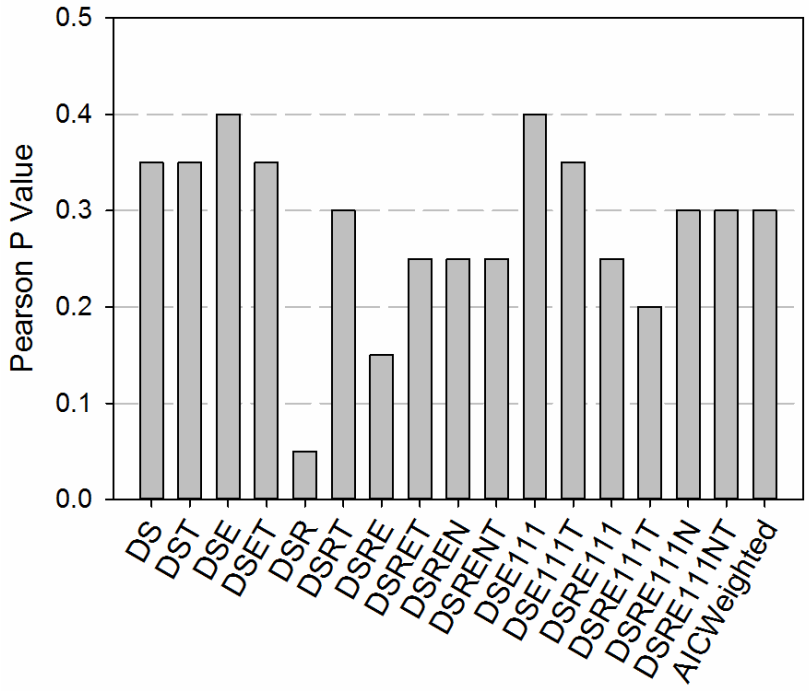

Figure 11: Average Chi Square on BIG292 data set across all models explored in Chapter 3. All models predictions differed from the observed DCS incident counts, however DSR differed the most significantly. This was largely due to poor predictive accuracy with respect to the Saturation profile type.

3.2.2.5 Comparison by Dive Type

Figure 12 graphically depicts the Xi square results across the free exponent models. For Repetitive Air, Single Non-Air, Repetitive Non-Air, and Saturation dive types all models performed well and achieved Xi square test results of at least 0.9. No model performed well on Single Air or Sub-Saturation profiles, whereas all models predicted well in the other four dive types, which matched performance in the model base. Overall, the results for comparison by dive type indicated similar performance between all tested models. Notably for Single Non-Air, Repetitive Air and Repetitive Non-Air all the models with decoupled exponents performed well, as compared to models with identical exponents in the base group.

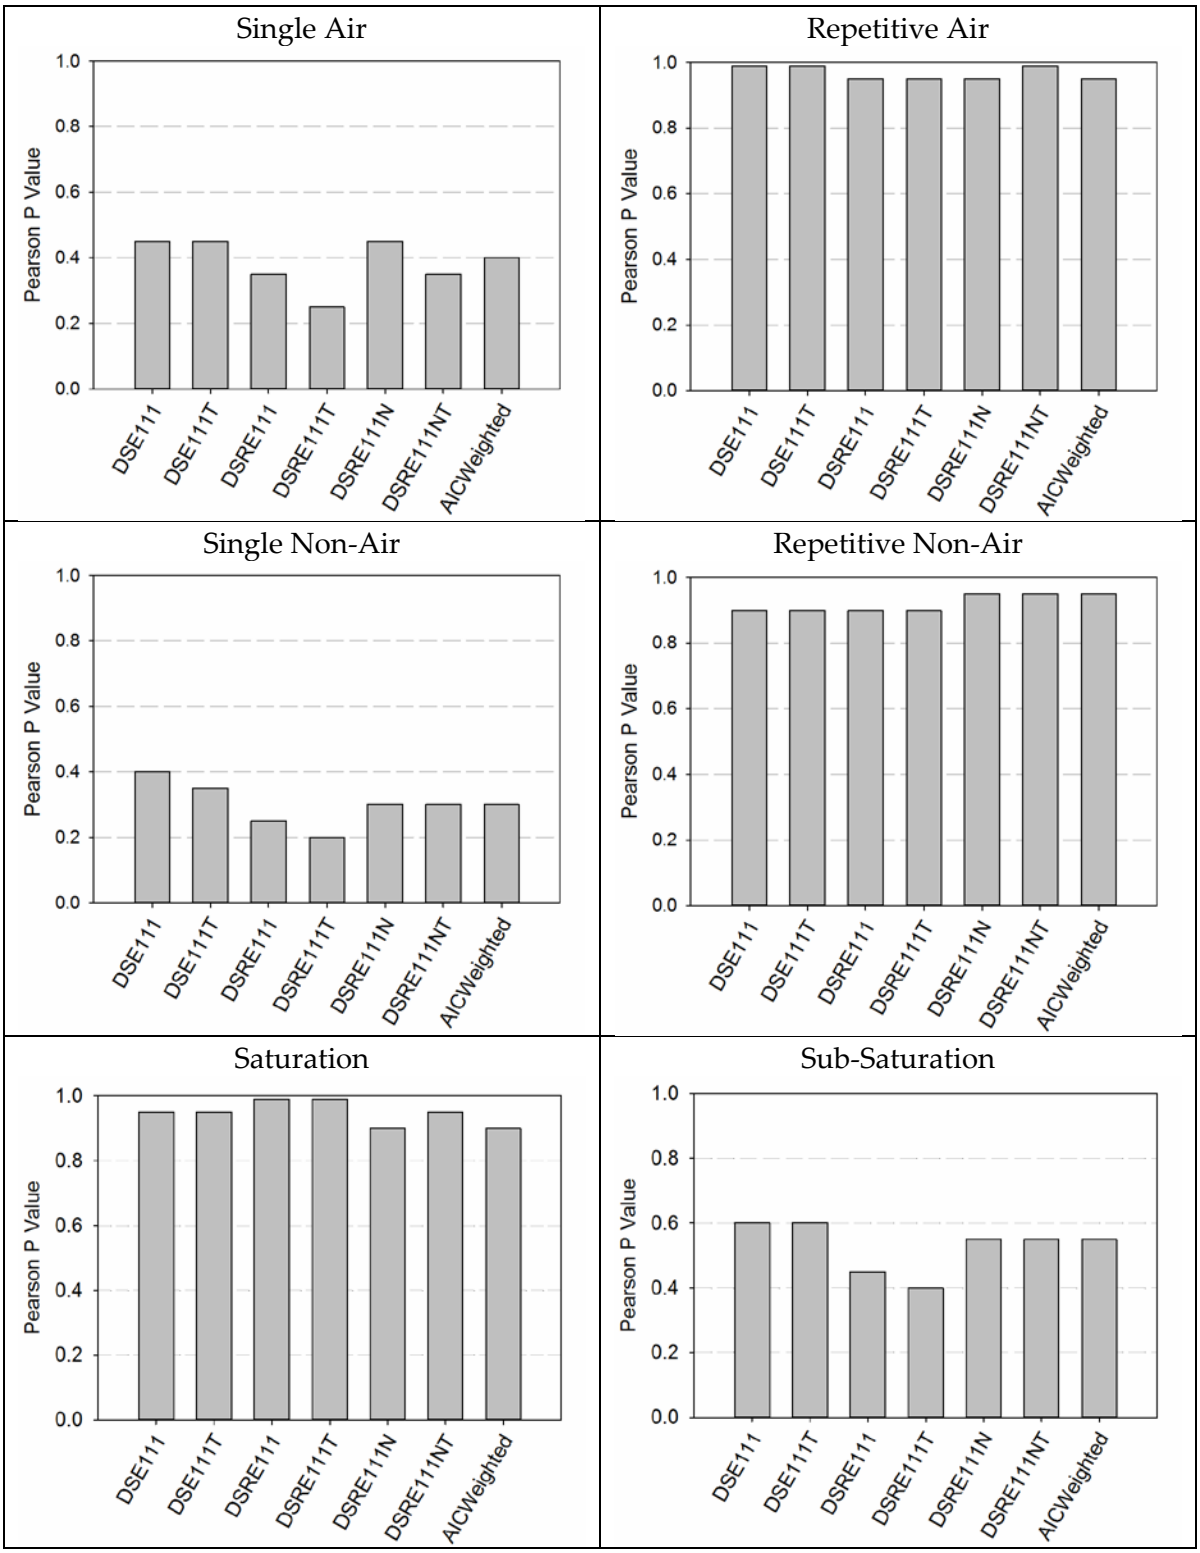

Figure 12: Xi square statistics by dive type for the models of Section 3.2.

3.2.2.6 Comparison by Dive Depth

In Table 11 we present the predicted number of DCS occurrences segmented by maximum profile depth for the free exponent models. Model predictions for each were similar to the corresponding fixed exponent models presented in Section 3.1, with the only significant difference being the predictions in the 300+ FSW profile group. The average PDCS for the models explored in Section 3.1 was 0.6, while for the free exponent model PDCS average was 0.94. This indicates that the addition of free exponent parameters on the models was not advantageous for improvement of average PDCS predictive quality. Deeper profile sets (100-300 and 300+) were found to have a greater predicted risk on average, which indicates that the expected effect of nonlinear scaling of the hazard function slightly increased predicted risk with depth.

Table 11: Model predicted DCS count by maximum profile depth. Similarly to the models in Section 3.1, all models over predicted DCS incidents in profiles with maximum depths less than 100 FSW and under predicted DCS for profiles deeper than 100 FSW.

Model / Depth (DCS)	0-50 FSW (18)	50-100 FSW (88)	100-300 (82)	300+ FSW (2)	Total (190)
DSE111	26.85	106.75	54.25	1.02	188.87
DSE111T	27.11	107.11	54.25	0.99	189.47
DSRE111	34.94	107.05	47.49	0.52	189.99
DSRE111T	31.57	107.77	48.53	0.67	188.53
DSREN111	31.07	108.4	51.95	1.16	192.58
DSREN111T	30.83	106.92	51.17	1.06	189.98
AICWeighted	31.01	107.99	51.74	1.13	191.86
Model Average	**30.48**	**107.43**	**51.34**	**0.94**	**190.18**

Model performance on dive data segmented by depth revealed that only profiles in the 300+ FSW category crossed the Pearson statistic 0.9 significance threshold (see Figure 13). All models differed from the observed DCS cases in the 100-300 FSW category significantly, which was similar to the predictions from the models explored in Section 3.1. It is important to note that the low number of DCS incidents in profiles in the 300+ FSW dive may positively bias model performance in this depth category. Similarly to the results in Section 3.1, we see that the results in the 0-50 FSW depth group are superior for the DSE111 and DSE111T models, which do not scale the supersaturation by the ambient pressure as do the other models in this section.

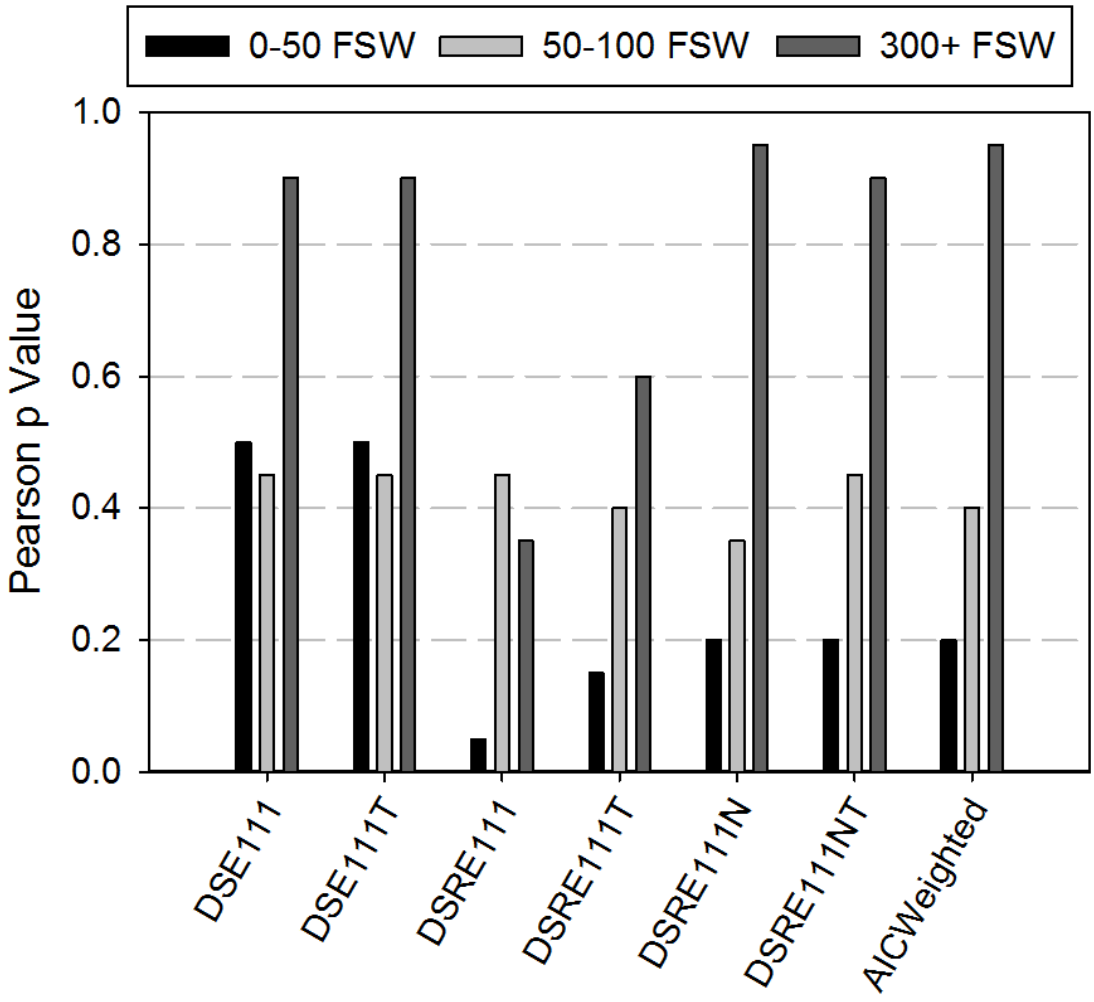

Figure 13: Xi Square results for the decoupled exponent hazard function model optimizations on the depth groups 0-50, 50-100, and 300+ FSW. The 100-300 FSW group is not shown because no model achieved a Pearson statistic of at least 0.05. Similarly to the results in Section 3.1, the models without ambient pressure scaling outperformed models with ambient scaling on the 0-50 FSW depth group.

3.3 Conclusion

In this chapter, we presented the results from parametric optimization of 16 variations of the EE model, each using a unique form of the hazard function. Predictive quality was determined using AICc weighting, as well as Xi square testing on profile subgroups by dive type and maximum profile depth.

AICc weighting indicated that two hazard function variations, specifically DSRE111N and DSRE111NT, were necessary to explain over 90% of the total evidence ratio. Removal of the DSRE111NT model from the set allowed the DSRE111N model to account for 88% of the overall evidence ratio individually. These results indicate that use of a hazard function of the general form of DSRE111N provides superior predictive quality over the alternative hazard functions explored in this chapter.

An AICc weighted multimodel was generated using the DCS incidence predictions of DSRE111N and DSRE111NT to determine if multimodel inference could improve predictive quality. Xi square testing indicated that the AICc multimodel predictions did not differ significantly from those of its source models. This is was likely due to the fact that the two models from which it was composed had nearly identical hazard function definitions, differing only by the inclusion of a threshold term on the third compartment which had a relatively small magnitude (on the order of 0.1 atmospheres). Predictions of the AICc multimodel were compared to the DSR model and indicated that the AICc multimodel separated risk predictions between compartments more than DSR,

which had risk accumulation in two or three compartments for many of the profiles in the Single Air dive type. Due to the fact that DSR had the closest overall predictive correlation to the DCS incidences in the Single Air dive type, this indicates that having risk accumulated across several compartments improves predictive performance for the EE model on Single Air profiles. The AICc multimodel had better predictive correlation with Saturation dive profiles, and often predicted substantially more risk per profile than did DSR, as was shown in Figure 9. Overall, the AICc multimodel did not uniformly improve on predictive quality compared to the other models in this exploration.

Interestingly, the scale of the optimized exponent parameters was similar to the parameter ranges of earlier explored hazard functions with nonlinear exponential risk scaling for both the identical compartment exponent models explored in Section 3.1 and the independent exponent models from Section 3.2. For the independent exponent models, exponent terms for compartments 1 and 3 were similar in magnitude (around 2), while the exponent for compartment 2 was always much smaller (around 0.0361). Also of note was the fact that parametric uncertainty was similar in our findings to that of previous explorations. Parameters in our explored models were similar to previously explored EE or Linear Exponential models [37].

Overall, the results of our investigation demonstrated that selection of the best hazard function can significantly improve model predictive quality. The addition of individually fitted compartment hazard exponent parameters and scaling by ambient

pressure improved model predictive quality in general on all dive types and for profiles with maximum depths of greater than 50 FSW. The hazard function DSRE111N utilizes both of these features and should be considered for use in future model explorations.

4. Modeling Using Delay Differential Equations

Pharmacokinetic (PK) models used to predict the occurrence of DCS often divide the body into a series of uncoupled tissue compartments and simulate the transfer of metabolic and non-metabolic gases between each compartment and the pulmonary system [38, 39]. Gas transfer between compartments in such models is often assumed to be instantaneous. In biological systems, the transfer of materials between body tissues is not instantaneous, and models can be made more realistic by incorporating mechanisms for to account material transport between compartments [40]. We can mathematically account for time-limited material transfer mechanics through the use of delay differential equations (DDE). In this chapter, we present the results of model optimization and selection on ten of our previously developed PK models which have been modified to include explicit delay mechanics.

Depending on model layout, a diffusion compartment may act as a storage vessel of dissolved gases by accumulating gas when ambient pressures are high. Other PK models have explored serially linked compartments which can have a longer gas uptake time for downstream compartments due to gas washout from upstream compartments [2, 41, 42]. Models that include diffusion or serially linked compartments and shift the accumulation of risk away from the time of first ascent have an implicit delay mechanism due to their structure. Implicit delay can improve the correlation between model

predictions and observed DCS incidents (which frequently occur after surfacing) by accumulating risk within the reported occurrence window [43].

Compartmental systems with lags have been investigated. In this work, only discrete time lags are explored. Continuous time lag functions exist which allow for the variation of the delay in transfer mechanics and could be explored in the future. In this work, we focus exclusively on discrete time lags. The addition of discrete delay mechanics into linear compartmental systems has been shown to have stable solutions and so we believed parametric optimization would identify the best time delay parameters for each model [40].

Lags are often advantageous for 'incomplete models' when the compartmental structure is missing some upstream mechanics. The addition of lags can allow for the incomplete model structures to account for model subsystems which are not present in the model differential equations, explained in detail by Jacquez [44]. This is advantageous for our purposes because we can potentially simulate more complicated systems using fewer fitted parameters. Jacquez [45] concluded that any compartmental systems that uses discrete time lags can be approximated by using the method of linear chain subsystems. The advantage of using lags in the context of predicting DCS incidence is primarily due to the fact that the inclusion of delay mechanics can account for more complex upstream mechanisms, collapsing a potentially large number of parameters into one term (thereby increasing the chances for good model performance by AIC selection).

Part of our motivation for investigating DDE models was the observation that the occurrence densify function (ODF) for DCS incidents has a bimodal structure, as shown in Figure 14 and has been previously reported by Gerth [16]. The ODF is maximal at 0, which represents the time of final surfacing, but a second peak occurs roughly 120 minutes post surfacing. We believed that the use of delay mechanics might allow our pharmacokinetic models to more closely match the prediction times associated with the 120 minute peak.

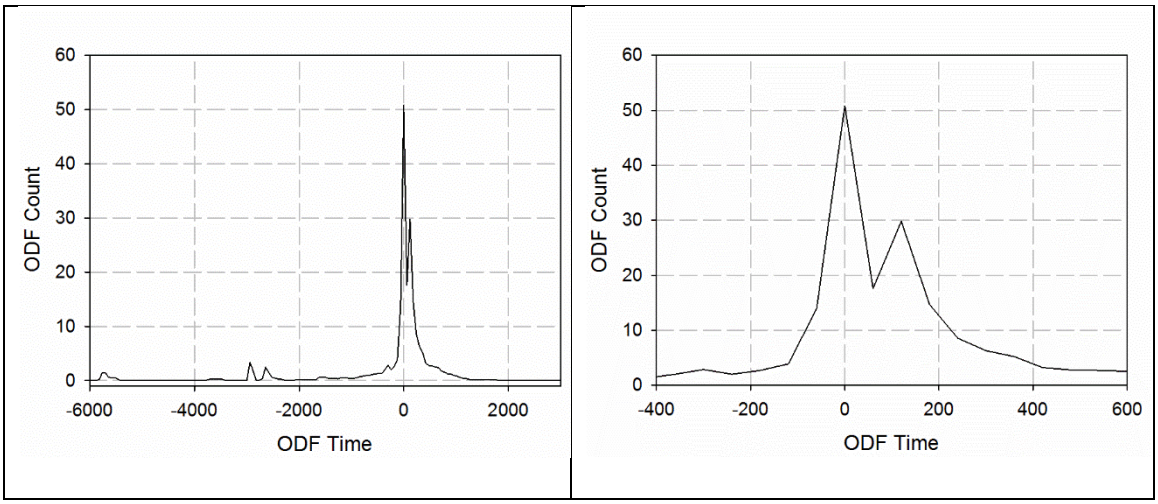

Figure 14: ODF of the Parker *et al.* data set used for model training in this chapter. On the left, the ODF window from 600 minutes before surfacing to 3000 minutes after surfacing is shown. On the right, the same data is shown but with a much tighter selection to display the bimodal peak of the ODF.

We investigated the distribution times of the reported DCS incidents to determine how long (in minutes) an explicit delay term should be for our models to capture the behavior of the ODF shown in Figure 14. Figure 15 displays a histogram of the T_1

distribution times centered around the time of last surfacing (0 on the abscissa corresponds to symptoms being first reported immediately on surfacing while a positive value indicates symptoms reported after surfacing). As the figure indicates, 58% of the total reported DCS incidents had T_1 times reported after surfacing (128 out of 223 total DCS incidents in the data set). The average T_1 time for the profiles which reported first symptoms after surfacing was 85 minutes after the final surfacing (last surfacing time for repetitive exposures). In Figure 16, we present a histogram of the profiles in the last bin of Figure 15. The profiles that had first reported DCS symptoms after the time of last surfacing appear to be two segmented into a 0-40 minute group (which contained 44 profiles) and a 120+ minute group (which contained 68 profiles). The observation that many profiles had their first reported symptoms after surfacing partially motivated our investigation for adding explicit delay mechanisms into our previously developed pharmacokinetic models.

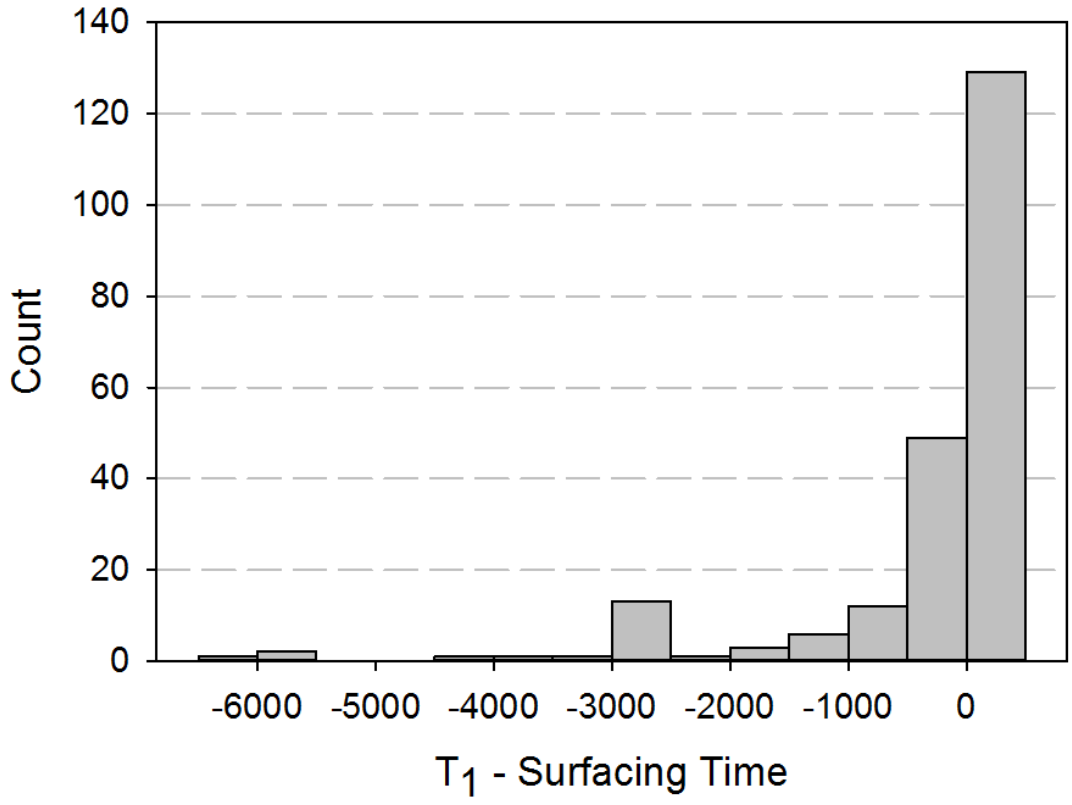

Figure 15: Histogram of the T_1 – Surfacing time for the Parker *et al.* data set. 128 of the 219 total profiles in the set had T_1 times that occurred after surfacing while 90 profiles had times T_1 observations prior to surfacing.

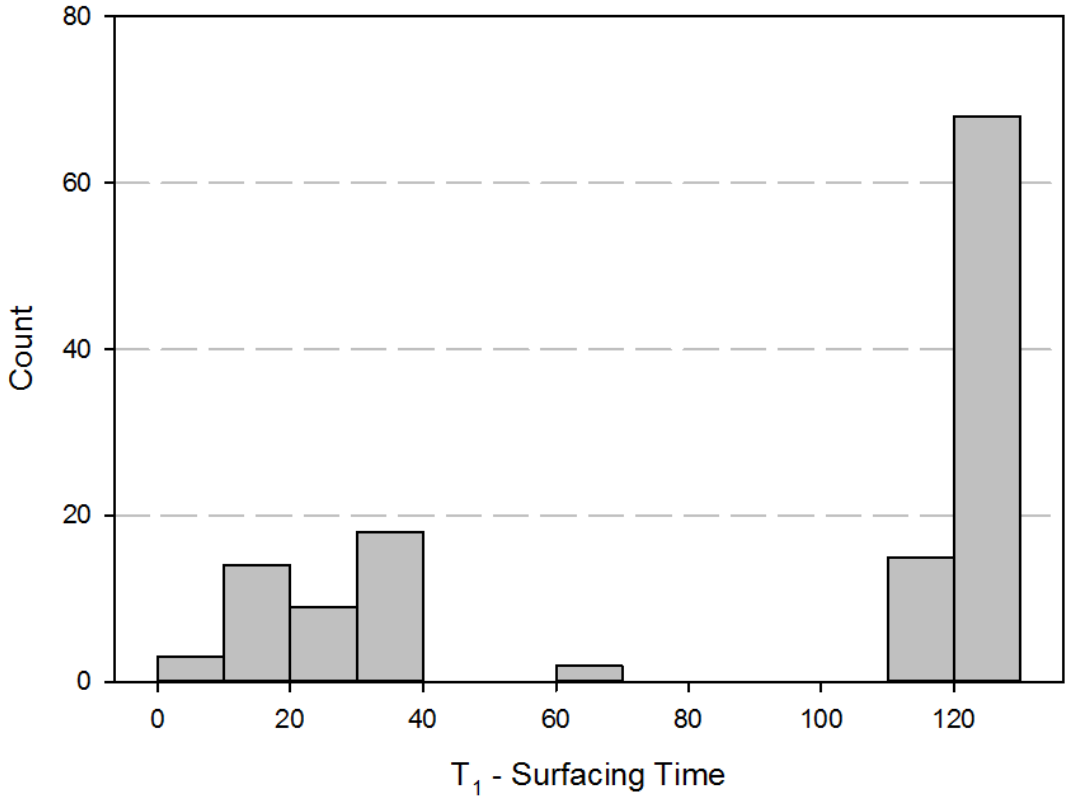

Figure 16: Histogram of the T_1 times for profiles with DCS incidents for which the first observed symptoms occurred after surfacing. 44 profiles were categorized in bins between 0 and 40 minutes, while 68 profiles fell in the 120+ minute bin.

PK models without delay assume instantaneous transfer of material between compartments. We can add explicit delay mechanisms to compartmental systems through the addition of a constant lag time τ, which has the effect of accounting for the transit time of material between compartments. As an example, consider the two compartment serially system PKS2T, depicted in Figure 17, which has no explicit delay mechanism.

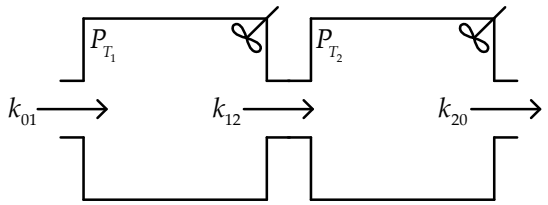

Figure 17: PKS2T is a serially linked, two compartment pharmacokinetic model described by Eq.(19). Material is received from the environment (q_0, not shown) and travels through compartments 1 and 2 before returning to the environment.

The differential equations describing the gas content of the two compartments in the system are

$$\begin{aligned} \frac{dP_{T_1}}{dt} &= k_{01}P_{T_A} - k_{12}P_{T_1} \\ \frac{dP_{T_2}}{dt} &= k_{12}P_{T_1} - k_{20}P_{T_2} \end{aligned} \tag{19}$$

where P_{T_A} is the gas input from the environment, P_{T_1} and P_{T_2} are the dissolved gas tensions of compartments 1 and 2, and the k_{ij} terms are the transfer rate coefficients from compartment i to compartment j. Note that the k_{ij} terms are simply the inverse of the tissue times reported in Chapter 3. For compartmental models like PKS2T that do not have explicit delay, risk accumulation begins immediately when the supersaturation value in any risk bearing compartment is positive.

A simple mathematical technique can introduce explicit delay through the use of DDEs in the gas transfer between compartments, allowing delayed risk accumulation with respect to reductions in ambient pressure [46]. If we add a delay mechanism between

the compartments, as shown in Figure 18, the differential system described by Eq.(19) changes such that

$$\begin{aligned}\frac{dP_{T_1}}{dt} &= k_{01}P_{T_A} - k_{12}P_{T_1} \\ \frac{dP_{T_2}}{dt} &= k_{12}P_{T_1}\left(t-\tau\right) - k_{20}P_{T_2}\end{aligned} \tag{20}$$

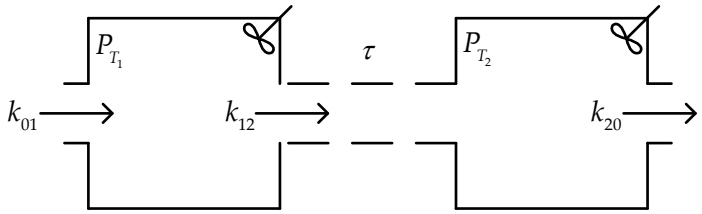

Figure 18: Displayed is a modified version of the PKS2T model which has a transfer delay time between compartments 1 and 2 of τ.

where t is the time and τ is the compartmental transfer delay. Material is lost from compartments 1 and 2 identically as in Eq.(19), but material gained in compartment 2 depends on the state of compartment 1 at time $t-\tau$. Delay mechanisms can be applied between the input to a system, or between compartments in a system. In this chapter, we focused only on delays to the input for the compartmental system to limit the search space to a manageable size, as examination of all permutations of inter-compartmental delay variants would have been too computationally intensive to finish exploration in a reasonable amount of time. As a demonstration, we apply the delay on the input to compartment 1 of Eq.(19), constructing the DDE describing PKDDES2T

$$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right) - P_{T_1}\right)$$
$$\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1} - P_{T_2}\right) \tag{21}$$

DDEPKS2T (shown below graphically in Section 4.1) accumulates risk only in the second compartment. Using explicit delay on the input compartment, both the beginning of risk accumulation and the position of peak risk can be moved away from the time of first decompression. In Figure 19, the effect on the instantaneous risk from a delay of 200 minutes on the PKS2T model is shown.

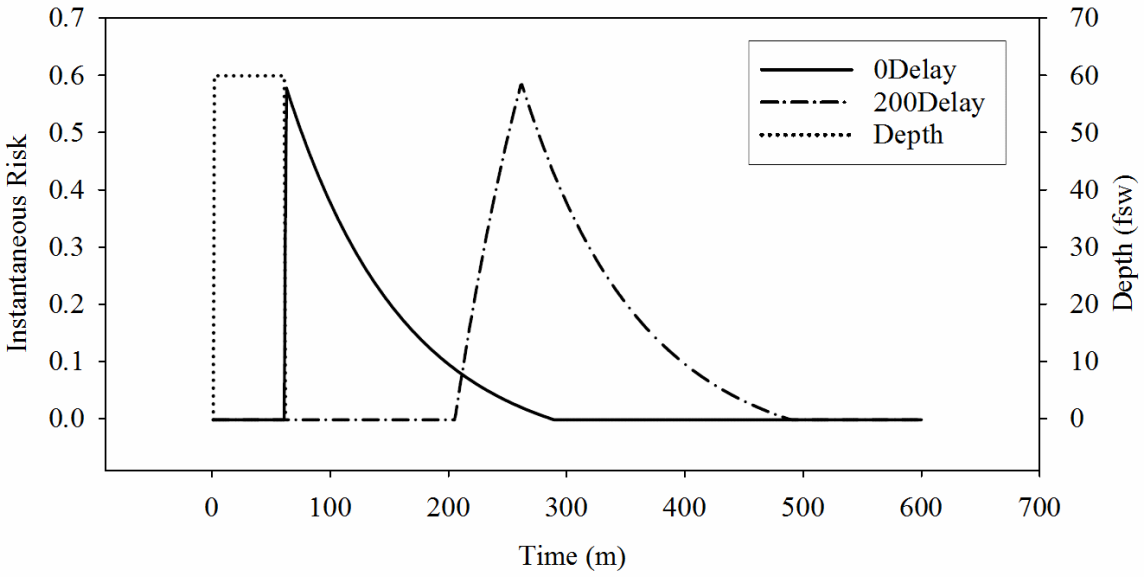

Figure 19: PKDDES2T with delay parameter τ set to 0 and 200 minutes on a hypothetical 60 minute exposure to 60 FSW. Both the peak risk and the beginning of risk accumulation are adjusted using the delay parameter.

With a delay of 0 minutes the system behaves identically to the differential system described by Eq.(19), with the time of maximum instantaneous risk occurring right at the

moment of first decompression (60 minutes). With a delay of 200 minutes, risk accumulation begins 200 minutes after the beginning of the pressure exposure, with the peak risk occurring 200 minutes after the time of first decompression. If the [T_1, T_2] interval for a DCS incident occurs after surfacing, it is simple to see that having a larger fraction of the risk accumulation occur within that window would improve the correlation between model predictions and the observed incident.

4.1 Methods

Ten models were derived by applying explicit delay to the arterial input of previously explored PK models [2, 38, 39, 47]. Models ranged from one compartment system (PKDDEPLB) to four total compartments (PKDDECS2T, PKDDECS2T3). Risk was accumulated in one (PKDDEPLB) to three (PKDDECS2T3) compartments depending on the model. The saturation ratio hazard function (DSR from Chapter 2) with time of onset risk accumulation was used in every model [10, 23]. In models that included diffusion mechanics, equal bidirectional gas diffusion between connected compartments was assumed. Model differential equations and diagrams are provided in Table 12.

Table 12: Diagrams and differential equations for the delayed input pharmacokinetic models explored in Chapter 4. An X within the compartmental diagram indicates that the compartment accumulated risk. In the case of PKCS2T and PKCS2T3, compartment 1 accumulated risk in PKCS2T but not in PKCS2T3 and it is marked with an 'o.'

Model Diagram	Model Differential Equations
PKDDEPLB 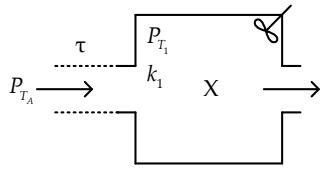	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right) - P_{T_1}\right)$
PKDDES2T 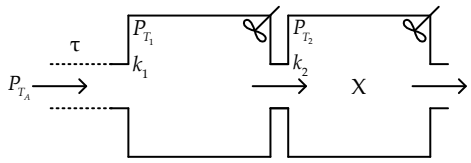	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right) - P_{T_1}\right)$ $\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1} - P_{T_2}\right)$
PKDDES3T 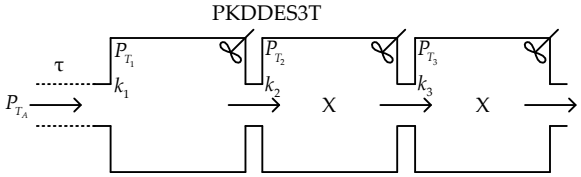	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right) - P_{T_1}\right)$ $\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1} - P_{T_2}\right)$ $\frac{dP_{T_3}}{dt} = k_3\left(P_{T_2} - P_{T_3}\right)$
PKDDEPDB 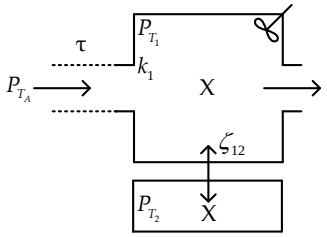	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right) - P_{T_1}\right) - \zeta_{12}\left(P_{T_1} - P_{T_2}\right)$ $\frac{dP_{T_2}}{dt} = \zeta_{12}\left(P_{T_1} - P_{T_2}\right)$
PKDDES2LP 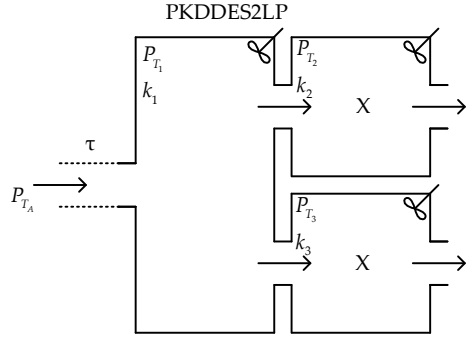	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right) - P_{T_1}\right)$ $\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1} - P_{T_2}\right)$ $\frac{dP_{T_3}}{dt} = k_3\left(P_{T_1} - P_{T_3}\right)$

PKDDES2LPD 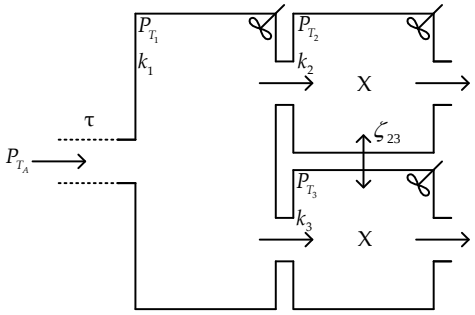 	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right)-P_{T_1}\right)$ $\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1}-P_{T_2}\right)-\zeta_{23}\left(P_{T_2}-P_{T_3}\right)$ $\frac{dP_{T_3}}{dt} = k_3\left(P_{T_1}-P_{T_3}\right)+\zeta_{23}\left(P_{T_2}-P_{T_3}\right)$
PKDDEPLCCD 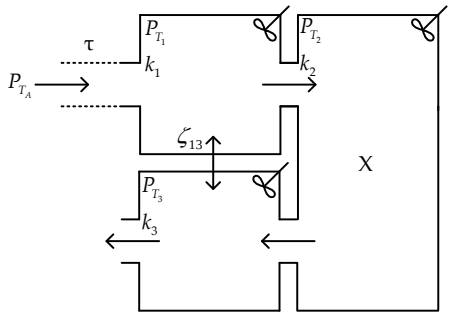 	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right)-P_{T_1}\right)-\zeta_{13}\left(P_{T_1}-P_{T_3}\right)$ $\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1}-P_{T_2}\right)$ $\frac{dP_{T_3}}{dt} = k_3\left(P_{T_2}-P_{T_3}\right)+\zeta_{13}\left(P_{T_1}-P_{T_3}\right)$
PKDDEPDCCD 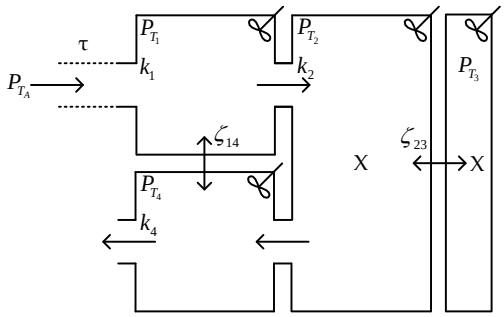 	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right)-P_{T_1}\right)-\zeta_{14}\left(P_{T_1}-P_{T_4}\right)$ $\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1}-P_{T_2}\right)-\zeta_{23}\left(P_{T_2}-P_{T_3}\right)$ $\frac{dP_{T_3}}{dt} = \zeta_{23}\left(P_{T_2}-P_{T_3}\right)$ $\frac{dP_{T_4}}{dt} = k_4\left(P_{T_2}-P_{T_4}\right)+\zeta_{14}\left(P_{T_1}-P_{T_4}\right)$
PKDDECS2T(3) 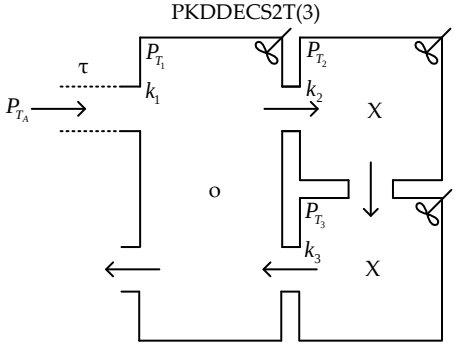 	$\frac{dP_{T_1}}{dt} = k_1\left(P_{T_A}\left(t-\tau\right)-P_{T_1}\right)+k_3\left(P_{T_3}-P_{T_1}\right)$ $\frac{dP_{T_2}}{dt} = k_2\left(P_{T_1}-P_{T_2}\right)$ $\frac{dP_{T_3}}{dt} = k_3\left(P_{T_2}-P_{T_3}\right)$

For each model, 32 solutions were generated using the NM algorithm, and the maximum LL value from each of the ten models was used for model comparison. Unlike in Chapter 3, simulated annealing was not employed on the NM optimized solutions. All model parameter distributions were assumed to be uniform. The bounds on the distributions from which each random initial guess for parameter optimization were defined using the technique described in Chapter 3. To reduce the number of fitted parameters, all model gain parameters were identified using the optimal gain method [9, 15]. Model selection through AICc weighting was used to allow comparison between the delay PK model group [36]. The Likelihood Ratio Test (LRT) was used to determine if the addition of delay parameters was statistically justified [2, 39].

Models were optimized on the NMRI98 dive data set which includes 223 DCS cases in 1304 profiles with 4335 individual exposures [48]. Marginal DCS incidents were ignored as their inclusion has been shown negatively affect model fit quality [31]. The NMRI98 data set contains enriched air, saturation, bounce, and oxygen decompression dive profiles.

Statistical model failure defined as a model calculated DCS probability of 0 on a profile which resulted in DCS was ignored (this concept is examined more closely in Chapter 5). Normally after optimization parameters are examined and if statistical model failure occurs the parameters are rejected. However, because we had no a priori knowledge of the delay parameter distributions and this was intended to be an

exploratory study to determine if the delay functionality warranted further examination, the statistical model failure restriction was relaxed.

4.2 Results and Discussion

AICc weighting results from optimized model LL values are shown in Table 14. The optimized model gas transfer rate terms and compartmental gain terms were similar in magnitude to our previous explorations of similar model formations [2, 39, 47]. Transfer delay values ranged from 1.13 to 7.26 minutes. We were surprised to find that the scale of the optimized delay term was quite short, as there are a large number of DCS incidents in the data set with T_1 times after surfacing (as was shown in Figure 15) with an average post surfacing T_1 time of 85 minutes. A possible explanation for the short optimized delay time is the fact that a significant portion of the total DCS cases had symptoms reported before surfacing (90 profiles) and another significant proportion of cases had symptoms reported between 0 and 40 minutes after surfacing (44 profiles). For these profiles, a large delay parameter would shift accumulated risk away from the ascent phase of the profile into the surface decompression phase (in which a diver remains in an ambient pressure of 1 ATM after a pressure exposure) which would probably not improve the overlap of model predicted risk with the $[T_1, T_2]$ interval.

Table 13: Optimized parameters for PKDDE Models

Model	g_1 (min^{-1})	g_2 (min^{-1})	g_3 (min^{-1})	k_{11} (min^{-1})	k_{12} (min^{-1})	k_{13} (min^{-1})	k_{14} (min^{-1})	k_{22} (min^{-1})	k_{23} (min^{-1})	k_{32} (min^{-1})	k_{44} (min^{-1})	τ (min)	P_{DCS}	LL
PKDDEPLB	7.61E-04	-	-	0.003	-	-	-	-	-	-	-	5.16	220.5	-1271.48
PKDDES2T	-	7.20E-04	-	0.004	-	-	-	0.098	-	-	-	7.07	220.31	-1277.08
PKDDES3T	-	1.02E-03	7.20E-04	3.4	-	-	-	1.88	-	-	0.003	2.35	223.77	-1170.13
PKDDEPDB	4.75E-04	7.48E-04	-	0.69	0.18	-	-	0.004	-	-	-	3.25	227.46	-1172.82
PKDDEPLCCD	-	7.60E-04	-	3.05	-	3.2	-	0.005	-	-	-	4.96	220.5	-1271.48
PKDDEPDCCD	-	3.40E-04	3.89E-04	0.02	-	-	0.12	0.49	0.014	0.57	-	3.47	222.04	-1247.58
PKDDES2LP	-	8.33E-04	7.21E-04	6.84	-	-	-	3.16	-	-	0.003	3.05	222.77	-1170.21
PKDDES2LPD	-	4.17E-04	3.43E-04	0.003	-	-	-	0.73	0.04	0.063	0.9	3.93	220.6	-1271.52
PKDDECS2T	-	3.88E-04	6.19E-04	0.38	-	0.2	-	0.21	-	-	0.005	7.26	222.52	-1203.3
PKDDECS2T3	5.07E-04	5.85E-04	3.82E-05	0.061	-	0.017	-	0.004	-	-	0.09	1.13	223.73	-1174.06

PKDDES3T and PKDDES2LP accounted for more than 95% of the total evidence ratio. Optimized parameters and the corresponding confidence intervals for the two AICc selected models are shown in Table 15. PDCS comparisons between PKS3T and PKDDES3T revealed that the only profiles for which a clear pattern of difference in predicted risk existed were from the submarine escape dive set SUBX87 (see Figure 20). A similar but less pronounced relationship can be seen in the comparison between PKS2LP and PKDDES2LP in Figure 21.

Table 14: AICc results from DDE model maximum LL optimization

Model	LL	Parameters	AICc Weight
PKDDES3T	-1170.13	6	0.501
PKDDES2LP	-1170.21	6	0.464
PKDDEPDB	-1172.82	6	0.034
PKDDECS2T3	-1174.06	8	0.001
PKDDECS2T	-1203.3	7	>0.000
PKDDEPDCCD	-1247.58	8	>0.000
PKDDEPLB	-1271.48	3	>0.000
PKDDEPLCCD	-1271.48	5	>0.000
PKDDES2T	-1277.08	4	>0.000
PKDDES2LPD	-1271.52	8	>0.000

Table 15: AICc weight selected PKDDE model LL optimized parameters with confidence intervals

Model	g_2 (min^{-1})	g_3 (min^{-1})	T_{11} (min)	T_{22} (min)	T_{33} (min)	τ (min)
S2LP	8.33E-4 ± 5.51E-4	7.21E-4 ± 1.01E-4	6.84 ± 30.64	3.16 ± 7.36	3.34E-3 ± 2.10E-4	3.05 ± 0.96
S3T	10.22E-4 ± 7.44E-4	7.20E-4 ± 1.01E-4	3.40 ± 11.85	1.88 ± 4.33	3.32E-3 ± 2.00E-4	3.05 ± 0.72

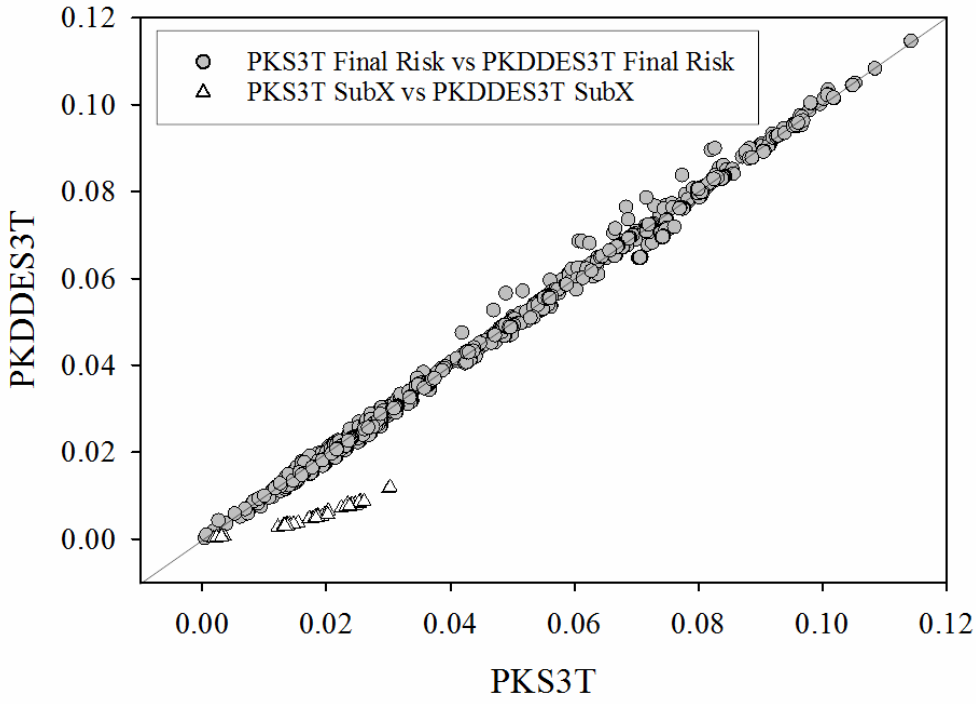

Figure 20: A graph showing the predicted DCS risk for PKS3T compared to that of PKDDES3T. The most significant difference in the model predictions occurred for the submarine escape profiles (shown as white triangles in the figure), for which PKS3T predicted a noticeably larger risk per profile.

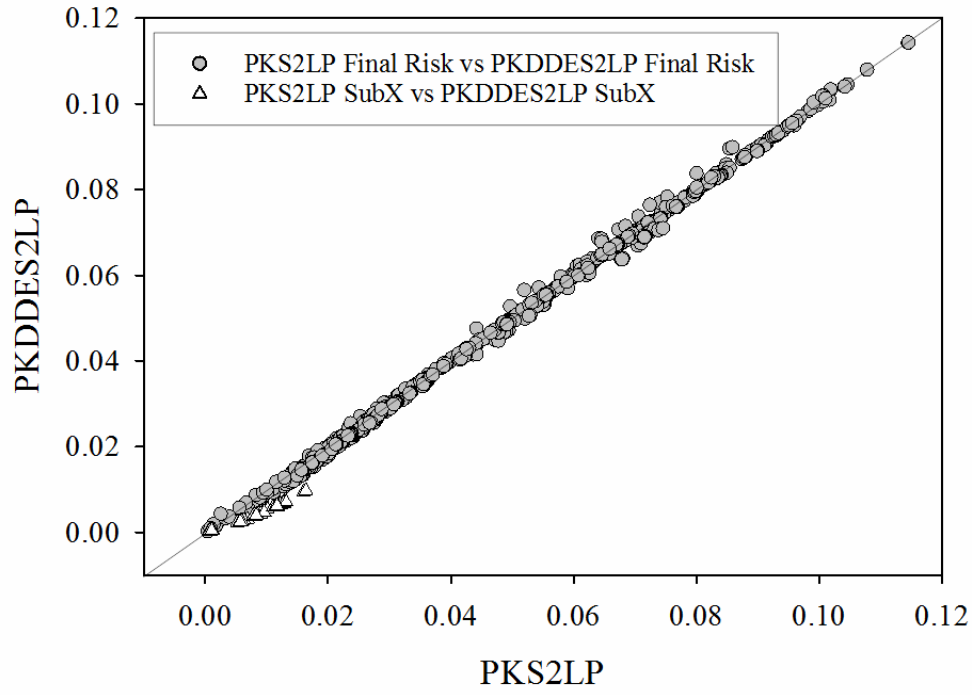

Figure 21: A graph showing the predicted DCS risk for PKS2LP compared to the delay model PKDDES2LP. The overall predictions are very similar. We included the submarine escape profiles in this figure (depicted as white triangles) to aid in comparison to the results shown in Figure 20.

AICc weighted comparison was performed between the optimized PKDDE results and the PK model PKCS2T3 (LL = -1166.87, degrees of freedom 7). In this test, the PKCS2T3 model was selected with an AICc weight of 83%, while PKDDES3T and PKDDES2LP accumulated 8% and 9% respectively. Generally, AICc model selection techniques require a single model to account for greater than 90% of the evidence ratio. The lack of a clearly selected model in this collection suggests that the additional delay mechanics do not improve the PK model fit sufficiently to justify their inclusion on these models fitted to the full NMRI98 set.

Table 16: Results from AICc model weighting using the best performing models from Chapter 4 with the best performing model from an exploration of similar pharmacokinetic compartmental models without delay. The model PKCS2T3, which did not include delay mechanics, significantly outperformed the models with delay.

Model	LL	Parameters	AICc Weight
PKDDES3T	-1170.13	6	0.08
PKDDES2LP	-1170.21	6	0.09
PKCS2T3	-1166.87	7	0.83

We used the LRT and the Wald test to determine if the delay mechanics from either of the two best performing models were statistically justified compared to identical models structures without delay. Test results using the LRT indicated that the addition of delay components for both PKDDES3T and PKDDES2LP were not statistically justified when compared to the previously explored PKS3T and PKS2LP models (PKDDES3T α = 0.30, PKDDES2LP α = 0.34, where α is the significance level of the LRT). LRT results indicate the addition of explicit delay mechanics for the two best performing models from this exploration. Results from the Wald test (with the null hypothesis that $\tau = 0$) indicated that the delay parameter was statistically justified for both models (Wald statistics for PKDDES3T and PKDDES2LP were 6.23 and 8.30 respectively). This result indicates that the when delay is included in the model, a nonzero delay is preferred, but does not significantly improve the model fit.

4.3 Conclusion

In this chapter, we presented the results from optimization of ten pharmacokinetic models which had a fitted explicit delay applied to the arterial gas inflow. Model selection using AICc weighting indicated that PKDDES3T and PKDDES2LP accounted for 50% and 46% of the evidence ratio respectively of the ten models tested. AICc weighting of the two selected delay models and the best performing pharmacokinetic model fitted without delay to the same data set (PKCS2T3) indicated that the model without delay fit the data significantly better. Additionally, results from the likelihood ratio test for the two best performing models compared to identical model structures without delay indicated that the addition of explicit delay mechanics were statistically unjustified. Overall, we found that the addition of explicit delay mechanics to the inflow of the compartmental systems did not improve model predictive quality.

5. On Identification of Allowable Model Parameter Regions

Models which predict the incidence of DCS are often parameterized using likelihood maximization to best match model predictions to training data. Parametric optimization for such models is a computationally expensive problem. Optimization methods including gradient ascent, simplex, and genetic algorithms have been employed for parametric LL optimization. In our experience, parameterization commonly takes months, even for simple models with only a few parameters. One source of difficulty for parametric optimization is the selection of regions in which parameters are examined. In this chapter, we present two methods that can be used to reduce the size of the parameter search region, which should reduce the difficult task of model optimization.

Optimized model parameters must be rejected if they result in a statistical model failure (SMF), which is designated as a predicted risk of zero on a dive profile in the training data set which resulted in a DCS incident. The reader may recall that in Chapter 4, we relaxed this restriction in order to reduce difficulty for optimization of the delay pharmacokinetic models. SMF occurs as a result of a parameterization of the model which causes the integration of the risk function over the risk accumulation interval to evaluate to zero. From an optimization standpoint, as a model parameterization approaches SMF, the value of the LL drops off sharply, as Howle *et al*. demonstrated in [11]. This can result

in locally nonanalytic regions of the LL function around SMF points, increasing the difficulty of optimization.

Because model failure is only defined for profiles with known DCS outcomes, we only need to examine those profiles which had divers experience symptoms, which generally (and thankfully) account for a small fraction of training data sets. Reducing the number of profiles to only those with DCS incidents will greatly reduce the computation requirements needed to identify parameters where model failure does not occur. Identification of the regions of parameters which do not result in SMF will significantly improve model optimization.

5.1.1 The Instantaneous Hazard Function

The total DCS risk of a profile is determined by integration of the instantaneous hazard function over the exposure time for the profile, described for incidence only risk definitions as $P(E)$ in Eq.(3), while the time of onset risk definition is defined in Eq.(9). For incidence only predictive DCS models, as long as the output of the instantaneous hazard function is greater than zero at some point over the length of the calculated profile time (generally from the beginning of an exposure to 24 hours after surfacing), SMF will not occur. For time of onset predictive DCS models the window of risk accumulation is greatly reduced to be just the time interval between T_1 the last known time without DCS symptoms and T_2 the first known time when DCS symptoms existed.

Recall from Chapter 2, we indicated that the hazard function is normally clamped to have nonnegative values, described by Eq.(5). In this chapter, we will use the unsigned hazard function and consider negative instantaneous hazard function values. Allowing negative hazard function values was necessary in order to identify the 'minmax,' a variable defined by model parameters which is discussed below.

If the output of the instantaneous hazard function is greater than zero anywhere over the risk accumulation window ($T_1 \le t \le T_2$), the model will calculate some non-zero risk for that profile. As long as the maximum value of the signed instantaneous hazard function for each profile, a

$$\alpha = \max_{h_i(t)} \left(h_i(t) \Big|_{t_1}^{t_2} \right) \tag{22}$$

is greater than zero in the risk accumulation interval, model predicted risk will be greater than zero and SMF will not occur. In the set of profiles with DCS symptoms, we need only consider the signed maximum value of the instantaneous hazard function for each profile to determine if that set of parameters results in SMF. We refer to parameter sets which do not lead to SMF as allowable parameters. We can simplify the determination of SMF further by recognizing that the set of profiles with DCS symptoms must have at least one profile with the smallest maximum instantaneous hazard function value. If the smallest maximum value is positive, the parameter set is allowable. We refer to the smallest maximum instantaneous hazard value as the minmax for that parameter set, defined as

$$\text{minmax} = \min_{\alpha_n}(\vec{\alpha}) \tag{23}$$

where α_n is the n^{th} member of $\vec{\alpha}$, the set of maximum instantaneous hazard values in the set of profiles associated with DCS outcomes. Thus, the minmax is a single value which is a proxy for the smallest integrated risk over the entire set of tested profiles for a given model parameterization. The reader may now question why we do not simply consider the integrated profile risk directly. The event probability of the survival model is by definition a positive value, as negative probability has no definition. If we examined the integrated risk, identification of the parameter value at which SMF occurs would be exceedingly difficult because parameters in the disallowed range would give no indication of where the boundary is (whereas a negative minmax values may inform how close to the allowable range a parameter is).

We introduce a test dive profile, called '60forf0,' to demonstrate the minmax value. The test profile is a hypothetical pressure exposure described by the series of depth and time nodes shown in Table 17. A graphical representation depicting the pressure (in FSW) ranging from 0 to 100 minutes is shown for the '60for60' profile in Figure 22. 60for60 was encoded to include a simulated DCS incident with T_1 = 220 minutes and T_2 = 1180 minutes.

Table 17: Table of time and depth nodes describing the hypothetical test profile '60for60.' A graphical representation for depth and time nodes is shown in Figure 22.

Time (min)	Depth (FSW)
0.0	0.0
1.0	60.0
61.0	60.0
62.0	0.0
2880.0	0.0

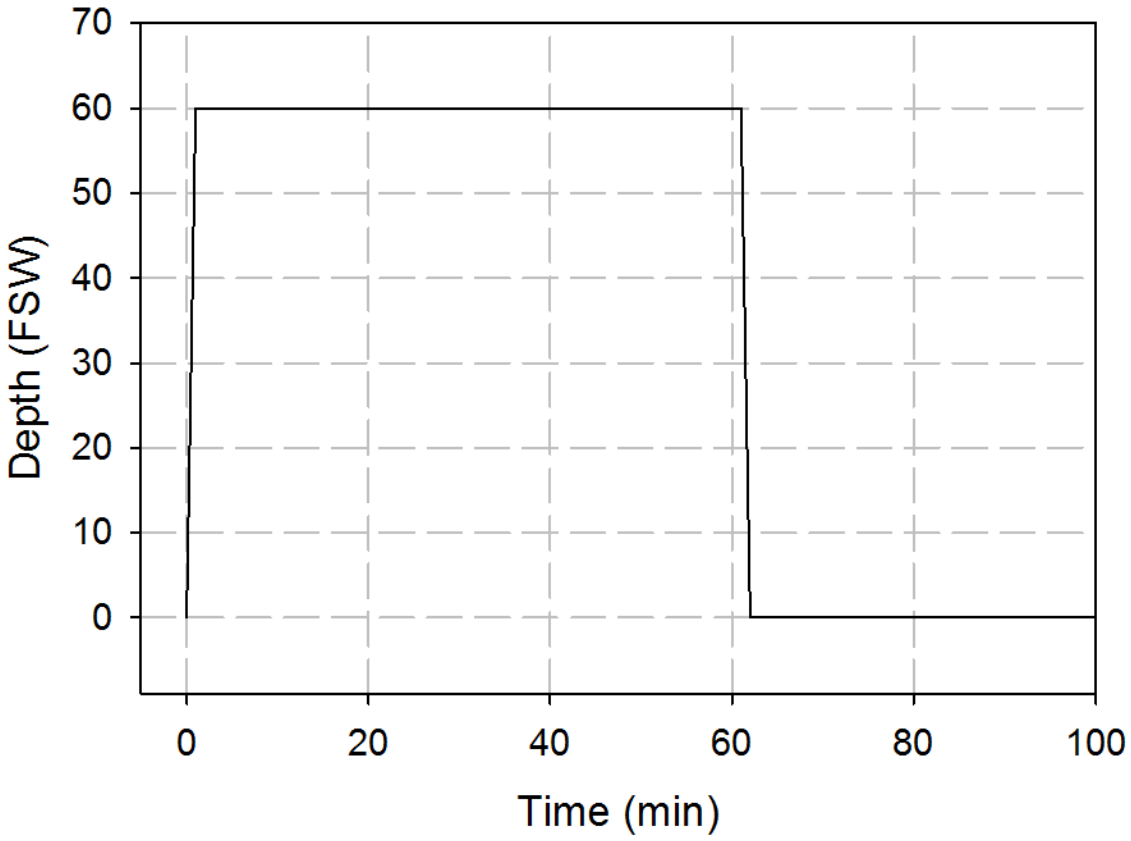

Figure 22: Time vs Depth plot for the hypothetical 60for60 dive profile. Note that after 100 minutes, the depth remains at 0 FSW until 2880 minutes.

We use the PKPLBX3 model as an example model for exploring the minmax values on sets of profiles using the 60for60 dive profile. PKPLBX3 is a three parallel compartment, perfusion-limited gas content model that is functionally identical to the DSR EE model used in Chapter 3 and is defined by the differential equations

$$\begin{aligned} \frac{dp_1}{dt} &= T_{11}^{-1}\left(p_{art} - p_1\right) \\ \frac{dp_2}{dt} &= T_{22}^{-1}\left(p_{art} - p_2\right) \\ \frac{dp_3}{dt} &= T_{33}^{-1}\left(p_{art} - p_3\right) \end{aligned} \tag{24}$$

in which T_{ij} is the tissue time term for inert gas transfer from the i^{th} compartment to the j^{th}. Note that in Chapter 4, we used the tissue rate term (simply the inverse of the tissue time). In this section, we use the tissue time, as the units of tissue time lends itself better to comprehension (minutes as opposed to minutes^{-1}). As stated earlier in Chapter 3, the tissue time characterizes gas transfer mechanics for each compartment in PKPLBX3. To demonstrate the differences in the minmax value resulting from changes in the compartment tissue time, we fix the tissue times for compartments C_1 and C_2 to T_{11} = 5 minutes and T_{22} = 50 minutes, while examining two tissue times for C_3 (500 minutes and approximately 65 minutes). Figure 23 shows the signed instantaneous hazard function values over the first 500 minutes for the 60for60 profile with the C_3 tissue time set to 500 minutes, while Figure 24 shows the function values with C_3 set to approximately 65 minutes.

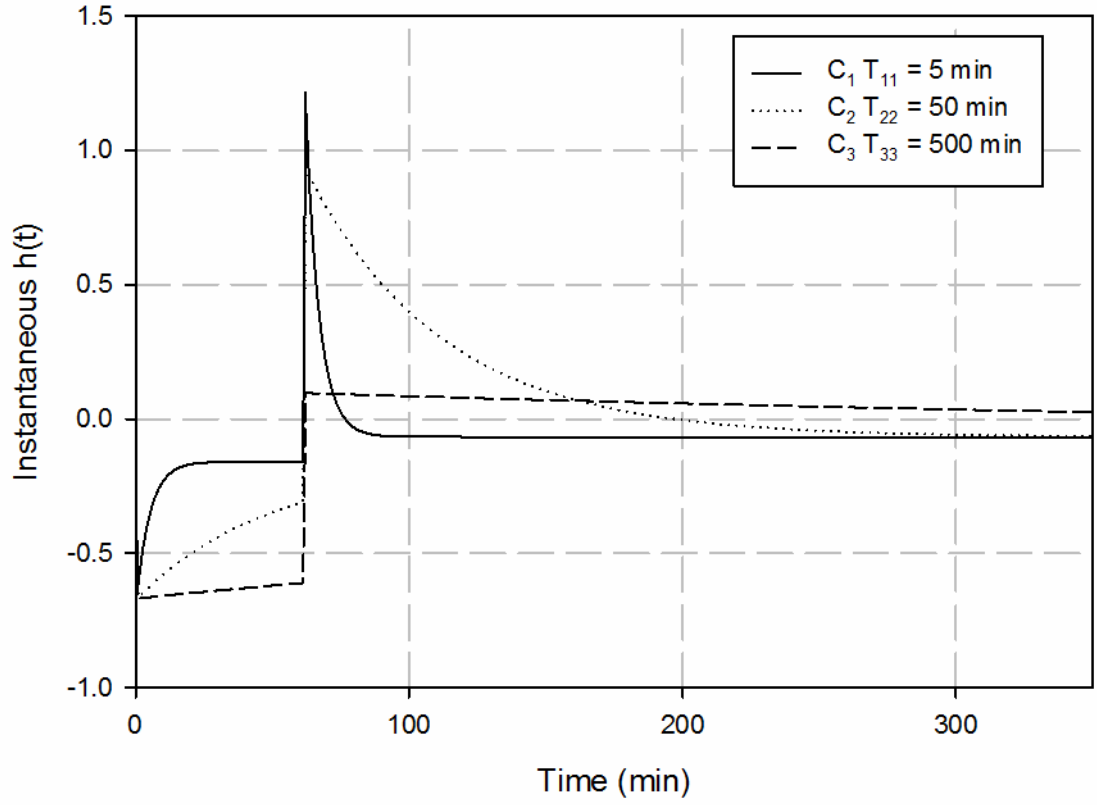

Figure 23: A comparison of PKPLBX3 between parameters resulting in statistical model failure and allowable parameters on the 60for60 test profile.

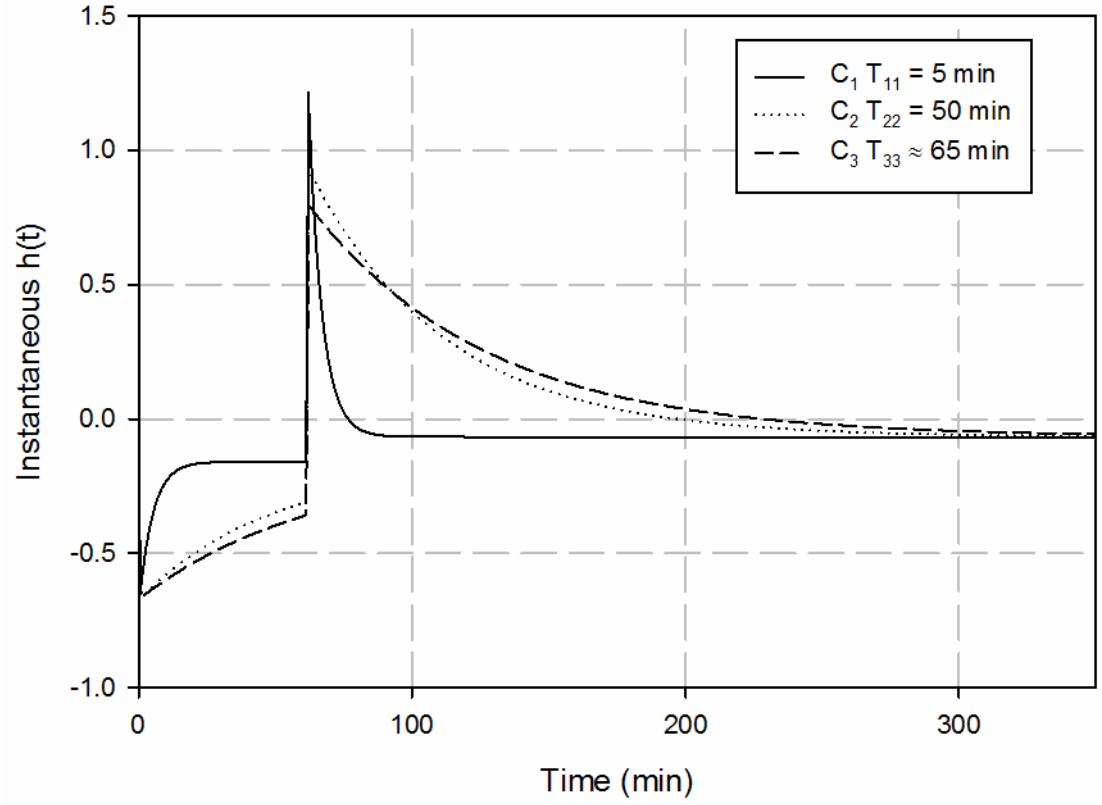

Figure 24: A graphical representation of the instantaneous hazard function for PKPLBX3 on the hypothetical '60for60' exposure described in Table 17. The C_3 tissue time is approximately 65 minutes, resulting in the instantaneous hazard function value dropping below zero after 220 minutes, the T_1 time of '60for60'.

Note that the instantaneous hazard function values for C_1 and C_2 are both less than zero at 200 minutes and do not increase in risk within the $[T_1, T_2]$ interval, resulting in neither compartment accumulating risk for the 60for60 profile in this parameterization. For T_{33} = 500 min, instantaneous hazard function values remain small over the entire dive profile and has a nearly linear slope after ascent at 60 minutes, while $C_3 \approx 65$ min shows a

high instantaneous hazard value at 60 minutes and an exponential decrease thereafter. Figure 25 shows the values of the instantaneous hazard functions for both C_3 parameterizations over the $[T_1, T_2]$ interval. The instantaneous hazard function for $C_3 = 500$ min is greater than zero between T_1 and about 400 minutes, while $C_3 \approx 65$ is zero at T_1 and negative between T_1 and T_2, as shown in Figure 26. This means that for $C_3 = 500$ minutes, risk is accumulated in the $[T_1, T_2]$ interval, but for $C_3 \approx 65$ it is not. Due to this, the 65 minute parameterization results in SMF. For tissue times greater than 65 minutes, risk will be accumulated within the $[T_1, T_2]$ interval, and SMF will not occur. For this parameterization, we refer to the 65 minute boundary as the SMF parameter boundary, as it is an identified parameter value at which SMF occurs.

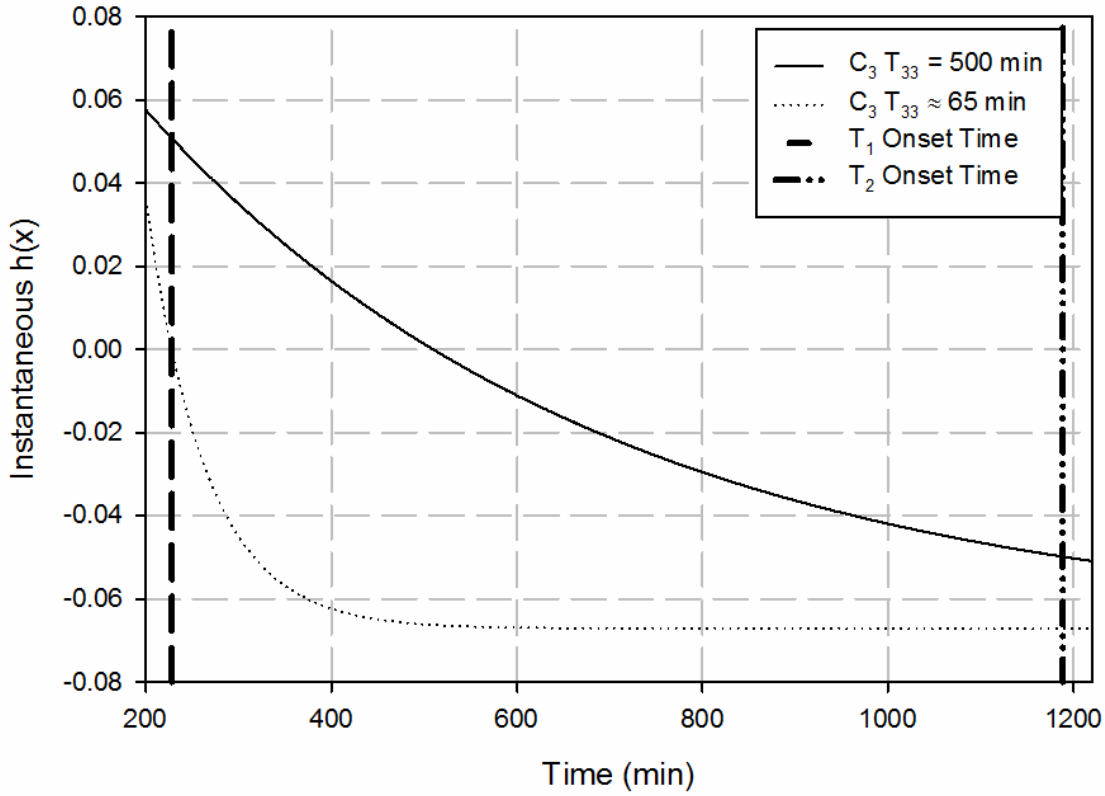

Figure 25: Comparison of C_3 = 500 min (allowable) with C_3 ≈ 65 min (statistical failure) within the [T_1, T_2] time of onset interval. For this example profile (60for60) C_3 ≈ 65 min never accumulates risk within the time of onset interval because the instantaneous hazard function value is less than zero during the entire window.

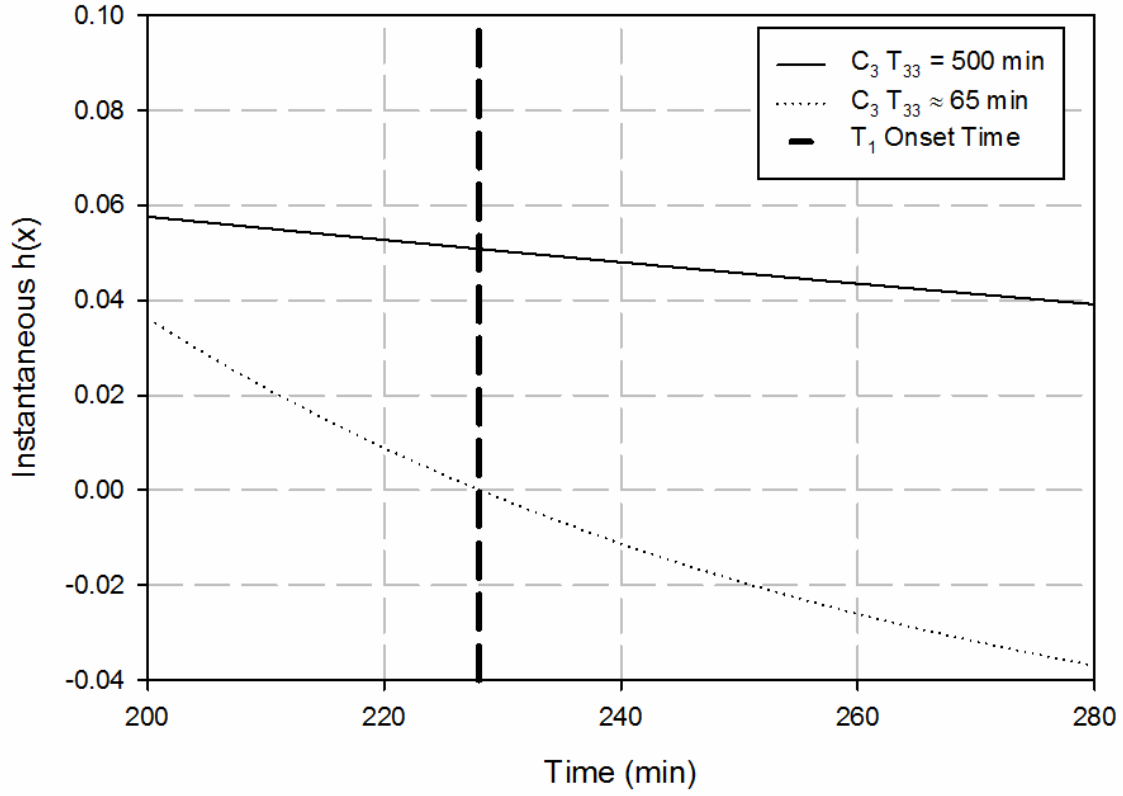

Figure 26: Instantaneous hazard function for compartment 3 of PKPLBX3 model. A tissue time of 500 minutes has an instantaneous hazard function greater than zero for part of the onset time region, resulting in an allowable parameter. A tissue time of approximately 65 minutes results in the instantaneous hazard function crossing zero exactly at T1, resulting in no risk accumulation within the onset time region and thus statistical model failure.

5.1.2 SMF Parameter Boundary Identification

Locating model SMF boundary parameters will expose the allowable regions of the model's parameter space in which optimal LL parameters must exist. After identifying the allowable parameter space, constrained optimization strategies can be employed to identify optimal local LL maxima. As an example, consider the identification of the

approximately 65 minute C_3 SMF parameter boundary in Figure 26. The maximum instantaneous hazard function value within the [T_1, T_2] interval will be negative for any tissue times smaller than the 65 minute SMF parameter boundary. Knowledge of that limit would improve model optimization strategies because parameters outside the allowable region could be excluded from examination.

5.2 Methods

5.2.1 Models

Our previously developed pharmacokinetic models served as a base for SMF boundary exploration [2, 39, 47]. Delay mechanics (explored in Chapter 4) were not included as the delay term was not found to significantly improve model fit quality. Models used for analysis were PKPLB, PKPLBX2, PKPDB, PKS2T, PKPLBX3, and PKCS2T3 (model differential equations are presented as they are discussed below). These models were selected to demonstrate SMF boundary identification in one to four dimensions. The instantaneous DSR hazard function was used for all models

$$h_i(t) = \frac{P_{T_i} - P_{amb} + P_{FVG}}{P_{amb}} \tag{25}$$

5.2.2 Data

We used the data set same discussed in Chapter 4, which includes 223 DCS cases in 1304 total dive profiles (more specific information regarding the data set can be found

in Chapter 4) [48]. As previously stated, exposures that resulted in either marginal DCS cases or no DCS were not included for model SMF parameter boundary exploration. Ignoring the 1081 non-hit profiles greatly reduced the computational load and increased parameter boundary identification throughput by reducing the computational requirements of analyzing the instantaneous risks for those profiles.

5.2.3 SMF Boundary Identification Strategies

Approaches for numerical continuation can be classified categorically into either predictor corrector (PC) methods or simplex methods. Simplex methods for manifold identification (also referred to as piecewise linear methods) represent a class of general curve tracing techniques which creates a map between the function surface of dimension Ω and a $\Omega+1$ dimensional triangulation simplex. This method allows identification of an arbitrary zero-curve with point accuracy on the order of the size of the simplex mesh [49]**.** Simplex methods require selection of a simplex grid which can make implementation very challenging. We instead focus on PC continuation methods, which are based off of numerical algorithms that can be used to integrate ordinary differential equations. PC methods function by taking an initial point, making a prediction for what the next function value will be, and the correcting that prediction. PC methods can be explicit or implicit, and many algorithms exist for both approaches. A simple example of an explicit PC method is the Euler-Trapezoidal method in which a predicted point $\hat{y}_{i+1}$ is generated

using the Euler Method on an initial point y_i with a step size of h and the derivative $\dot{y}_i$, shown functionally as

$$\hat{y}_{i+1} = y_i + h\dot{y}_i \tag{26}$$

The initial prediction is then corrected using the trapezoidal rule, including information from the initial point and the predicted point such that

$$y_{i+1} = y_i + \frac{1}{2}h\left(\dot{y}_i + \dot{y}_{i+1}\right) \tag{27}$$

where $\dot{y}_{i+1}$ is the derivative at the initially predicted point $\hat{y}_{i+1}$. Repeated corrections using the predicted point y_{i+1} in Eq.(27) until numerical stability is achieved improves the prediction quality. The corrected prediction point can be used as an initial point to solve for y_{i+2} using the same approach, and on ad-infimum. More complex PC methods exist which may use local manifold information, often including already identified solution points, to make a prediction as to where the next solution point will be and then correct the estimate to the true boundary value [49].

PC methods can be difficult to implement for higher order problems, especially in situations involving bifurcation or turning points where PC methods select "the next branch more or less at random" [50]. In this context, a solution branch refers to a bifurcation point upon which two or more solution directions exist. A graphical representation of this is shown in . No generalizable, computationally efficient homotopic methods currently exist for solving for solution manifolds of greater than degree five [49,

50]. The limitation for solution manifold identification comes from the dimensionality of the PK models, which range from degree 1 (PKPLB) to degree 5 (PKPDCCD). Note that the dimensionality does not include the gain parameter as it is uninformative from the perspective of the allowable solution manifold (gains scale output probability estimates and do not affect the value of the instantaneous hazard function). In order to avoid dimensionality issues for more complex models, we elected to use a discretization approach for identification of SMF parameter boundaries. A root bracketing method (described below) was used on one chosen model parameter while all other model parameters were fixed. Upon identification of the SMF parameter boundary or after a set number of solution identification attempts, fixed parameters were iterated by a preselected step size. Smaller step sizes resulted in higher resolution examinations of the parameter space but increased the computation necessary to complete the examination. In this work, we used low resolution examinations as our principal goal was to demonstrate the validity of identification of SMF boundaries using the minmax value.

SMF parameters can be identified through the use of root bracketing methods that employ the intermediate value theorem [51], which states

> If f is continuous on the closed interval $\left[a,b\right]$ and k is and number between $f\left(a\right)$ and $f\left(b\right)$, then there is at least one number c in $\left[a,b\right]$ such that
>
> $$f\left(c\right)=k \tag{28}$$

Provided that $f(a)$ and $f(b)$ have opposite signs, it follows that a root of the function exists such that $f(c)=0$. Many root bracketing methods exist, each with different convergence rates and properties. Out of these methods, we selected the bisection method primarily because of its convergence properties, which guarantee it will identify the roots of a differential system regardless of the surface continuity [27]. Guaranteed convergence is due to the fact that convergence using the method depends on the difference in sign between the ends of the closed interval $[a,b]$. Because we lack *a priori* information about the minmax surface continuity for an arbitrary PK model, the convergence properties of the bisection method are advantageous. Additionally, the bisection method is relatively simple to implement compared to other root bracketing methods, which reduced development time for the SMF boundary identifier.

For each sweep of the fixed model parameters, two initial input points were used to generate search intervals, a smaller valued term α, and a larger value term β. In our minmax root identifying routine, α and β were user defined values. If a positive minmax value was detected at α, the method would decrease the tested numerical value by a user set size percentage adjustment value until a minmax less than zero was identified. The last identified positive minmax and the first identified negative minmax would be used as an input interval for the bisection solver. The same process would be employed to find the β SMF parameter value by increasing the numerical value. If a positive value was not identified for either α or β, test parameters were adjusted in size by a user defined change

percentage. Large percentage size changes allow a larger area of parameter space to be tested with lower resolution, but it was found that a small percentage change of 3% was necessary to identify allowable regions reliably. A flowchart representation of the method can be seen in Figure 27.

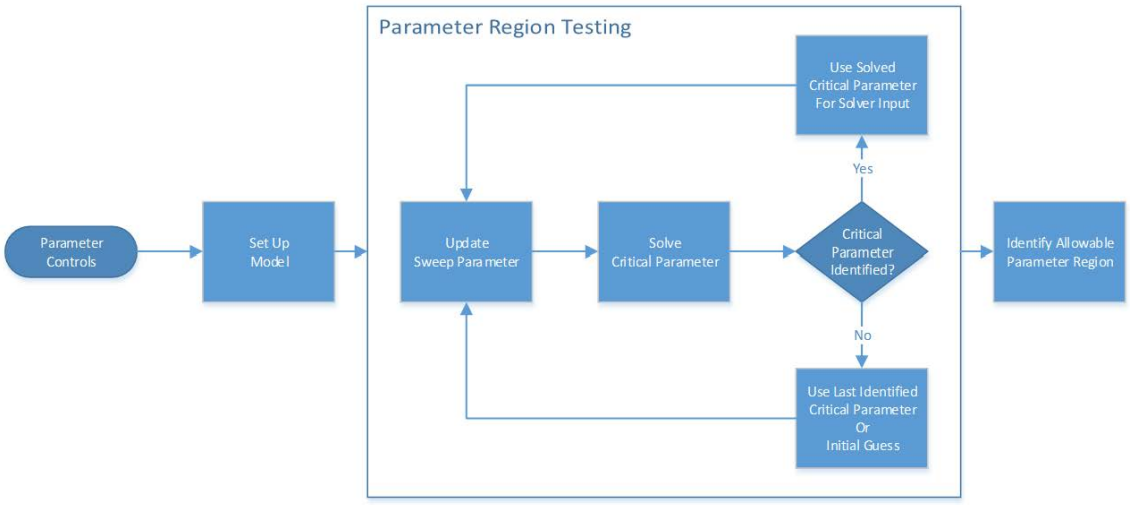

Figure 27: Flowchart of the bisection SMF parameter identification method.

5.3 Results and Discussion

5.3.1 Parameter Minmax Sweep Results

Use of our minmax sweep routine allowed exploration of models to see if any allowable parameters existed by providing a low resolution representation of the allowable parameter space, if any existed. It is important to note that the sweep method cannot guarantee that a candidate model does not have a region of allowable parameters, as there may exist individually allowable parameter sets outside in regions of SMF

parameters, or very small regions of allowable space that could be easily stepped over (an example is shown in Figure 28). Additionally, the parametric sweep does not provide any insight into the maximum LL surface, although the maximum LL parameter set must exist at a point of allowable model parameters.

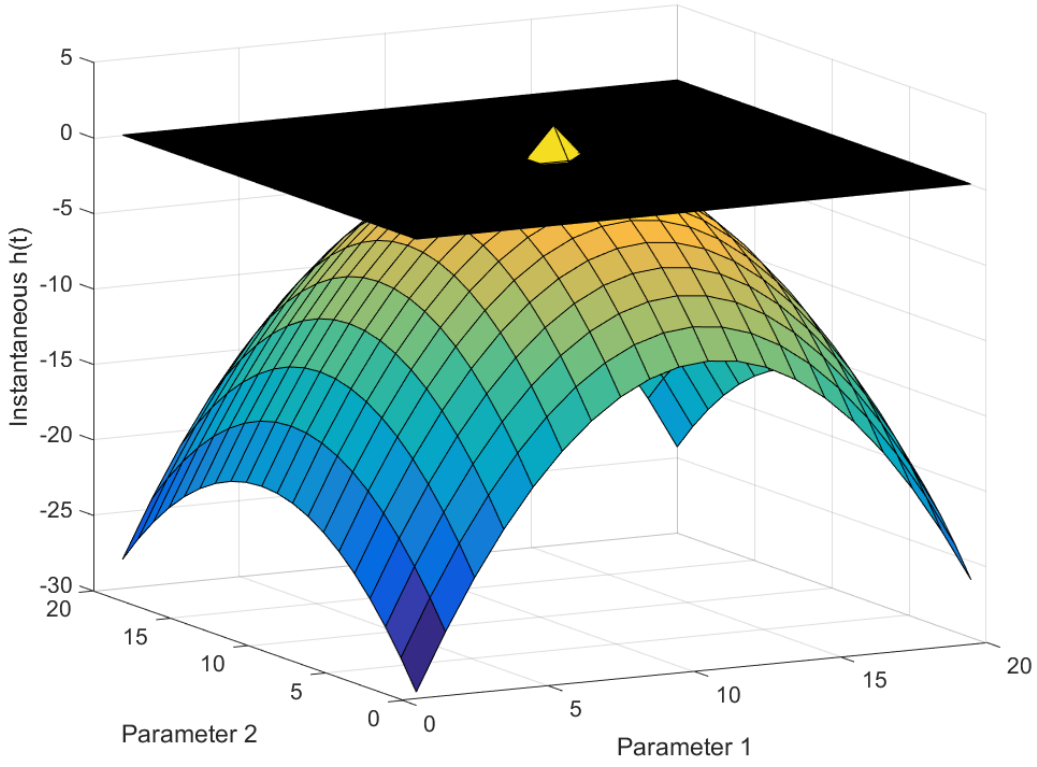

Figure 28: A hypothetical model instantaneous risk surface displaying one set of allowable model parameters. The minmax sweep method may have difficulty identifying that this model has any allowable parameters.

5.3.1.1 Investigation of the PKPLB Model Parameter Space

PKPLB is a simple one compartmental pharmacokinetic model defined by the differential equation

$$\frac{dP_{T_1}}{dt} = T_{11}^{-1}\left(P_{T_A} - P_{T_1}\right) \tag{29}$$

A one dimensional parametric sweep on the PKPLB model from 1 minute to 500 minutes with 1 minute increments resulted in SMF for all tested tissue times. A plot of the instantaneous risk value as a function of the compartment tissue time can be seen in Figure 29. Tissue times beyond 200 minutes were omitted to retain graphic resolution. The largest instantaneous risk value over the entire examined parameter range was -0.071, indicating that no parameterizations of the PKPLB model will be allowable for the explored data set. The nonlinear behavior of the instantaneous risk minmax is a product of changeovers between controlling profiles and also the nonlinearity of individual profiles minmax surfaces.

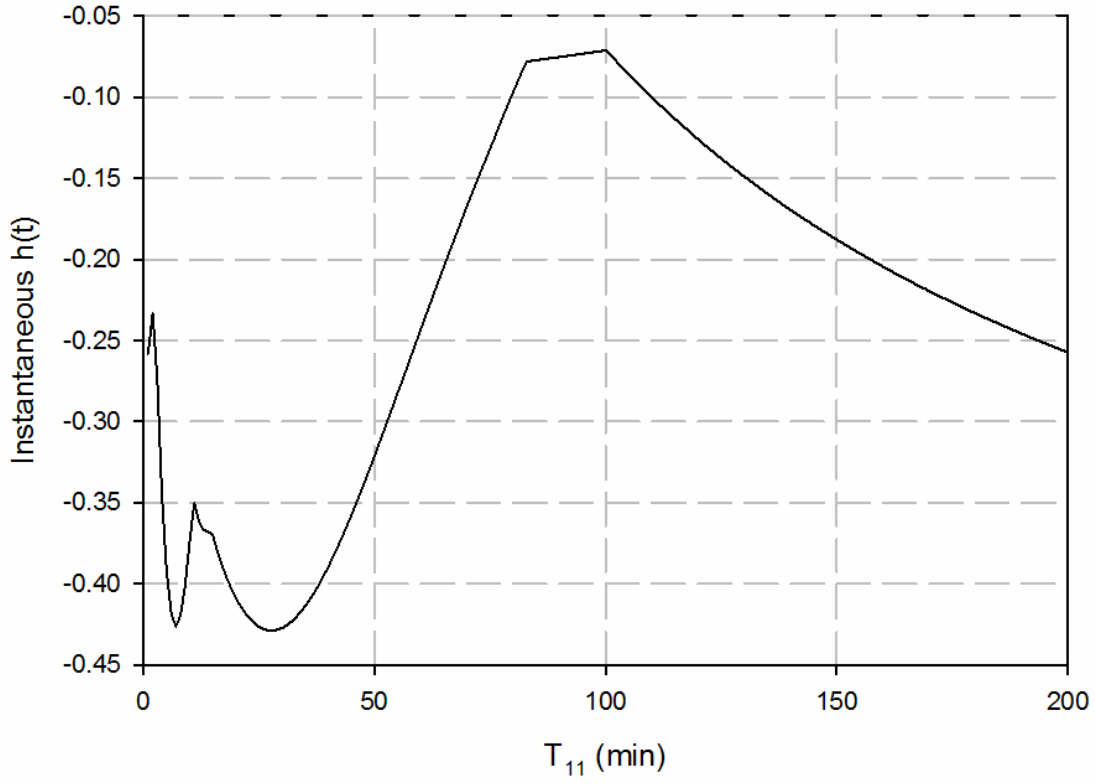

Figure 29: Instantaneous hazard function minmax for PKPLB for tissue times between 1 and 200 minutes. The largest minmax value was detected around 100 minutes. Note that each integer minute tissue time was calculated on this plot, producing a high resolution sweep of the parameter space.

The lack of allowable parameters for this model was not surprising as our group has previously tried to optimize the model and were unable to find allowable parameters. A similar single compartment model to PKPLB explored by Weathersby *et al.* was outperformed by LL ratio test by a constant risk null model, although the data set examined in the study did not include saturation diving [4]. We hypothesized that removing saturation profiles from the data set may create a more homogeneous collection

of profiles, allowing a single tissue time to account for the accumulated risk for all profiles. Figure 30 shows the instantaneous risk for PKPLB without saturation profiles for tissue times ranging from 1 minute to 250 minutes. The tissue time associated with the largest instantaneous risk was 94 minutes without saturation profiles (it was 100 minutes with the full data set). The largest instantaneous risk for the system (the minmax) increased from -0.071 for the full data set to -0.053 for the set without saturation profiles.

Tissue gas kinetic transfer rates that work for long decompression phases would fail to accumulate risk over a very short decompression phase, and rates appropriate for short decompression phases will under predict risk accumulation in the [T_1, T_2] interval for incidents associated with slow decompressions.

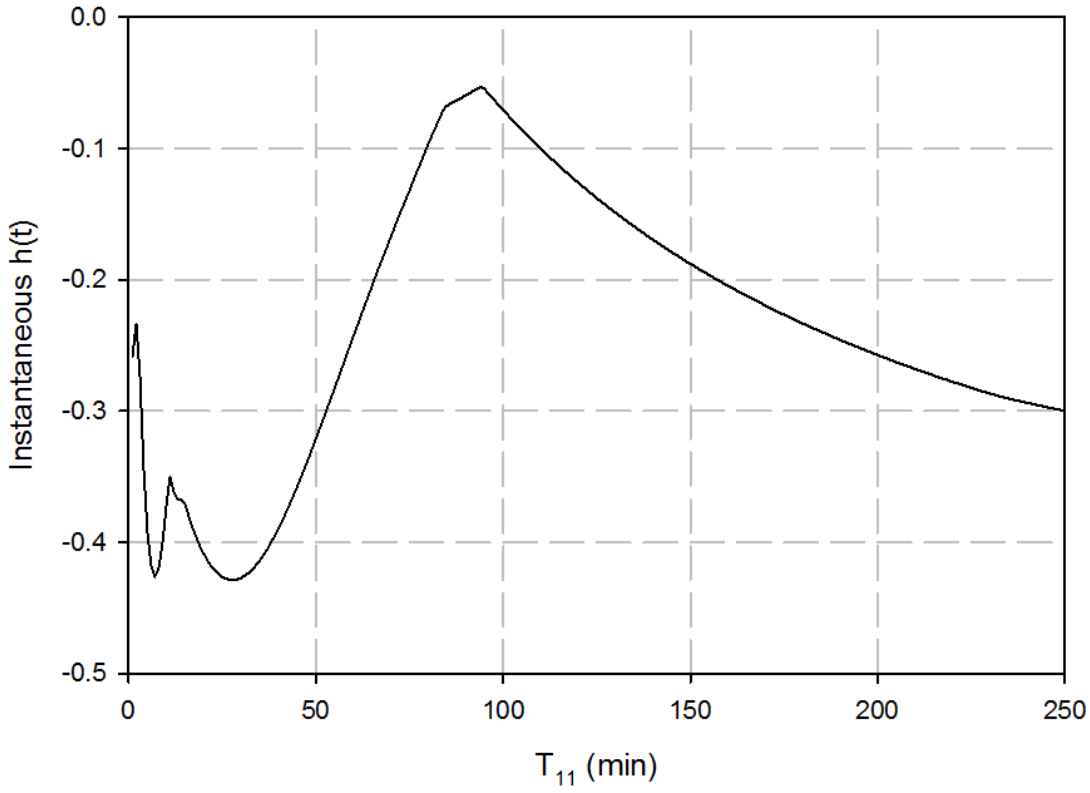

Figure 30: Instantaneous risk by tissue time for PKPLB without saturation profiles. The maximum instantaneous risk occurs slightly earlier than it did for the full data set (94 minutes vs 100 minutes) and is slightly larger in magnitude (-0.053 vs -0.071).

5.3.1.2 Investigation of the PKS2T Model Parameter Space

The two parameter serially linked two compartment model PKS2T is defined as

$$\begin{aligned} \frac{dP_{T_1}}{dt} &= T_{11}^{-1}\left(P_{T_A} - P_{T_1}\right) \\ \frac{dP_{T_2}}{dt} &= T_{22}^{-1}\left(P_{T_1} - P_{T_2}\right) \end{aligned} \tag{30}$$

was explored using the parametric sweep method for tissue time values between 1 and 500 minutes. A plot of the model minmax surface can be found in Figure 31. No allowable parameter set was identified, and the general trend of the model minmax surface suggests that all parameters outside the explored region will result in SMF. The largest minmax value was identified at T_{11} = 101 minutes and T_{22} = 1 minute of -0.074. The minmax surface for the PKS2T model is nearly symmetric, which was unexpected from the serial two compartmental structure (as opposed to the parallel two compartment structure of the PKPLBX2 model, which sensibly has a symmetric minmax surface).

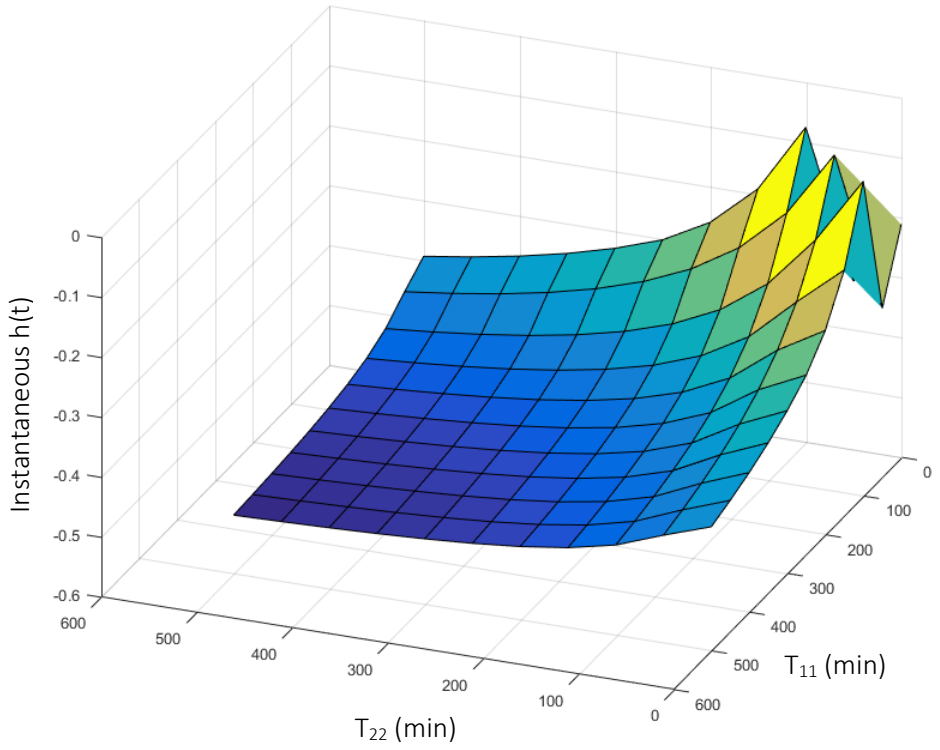

Figure 31: Minmax surface for PKS2T over the reduced Parker data set. The surface is symmetric across the axis of symmetry. No allowable parameters were identified for the PKS2T model.

We examined the minmax surface with a higher resolution near the edge T_{22} = 1 minute with T_{11} ranging from 1 minute to 200 minutes to determine if the edge of allowable space for PKS2T was similar to PKPLB. In Figure 32, the minmax surface for T_{22} = 1 minute is shown. The overall shape of the PKS2T is nearly identical to that of the monoexponential PKPLB for T_{22} = 1 minute. PKPLB's minmax value of -0.071 was found at the same tissue time value as PKS2Ts maximum minmax (around T_{11} = 100 minutes).

The similarity between the PKPLB and PKS2T models indicates that for small compartment tissue times for either compartment in PKS2T causes the model to behave nearly identically with respect to the minmax surface.

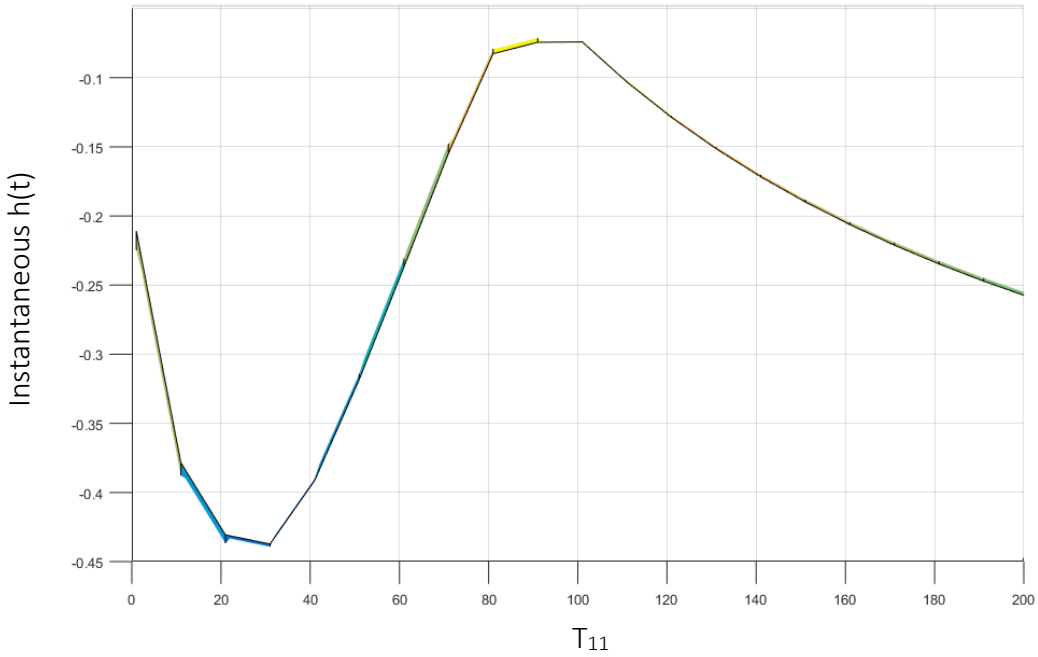

Figure 32: PKS2T Minmax surface on a reduced parameter interval with a higher resolution plot of the surface from T_{11} = 1:200 minutes with T_{22} = 1 minute. The minmax surface is very similar to that of PKPLB (see Figure 29). Note that the parameter spacing was 11 minutes while in PKPLB 1 minute spacing was used.

5.3.1.3 Investigation of the PKPDB Model Parameter Space

PKPDB is a two compartmental differential system which simulates diffusion mechanics using the set of differential equations

$$\frac{dP_{T_1}}{dt} = T_{11}^{-1}\left(P_{T_A} - P_{T_1}\right) - T_{12}^{-1}\left(P_{T_1} - P_{T_2}\right)$$
$$\frac{dP_{T_2}}{dt} = T_{12}^{-1}\left(P_{T_1} - P_{T_2}\right) \tag{31}$$

and was explored using the minmax sweep technique for tissue times and gas transfer times of 1 to 250 minutes with an increment step resolution of approximately 12 minutes. We arbitrarily chose our examined parameter boundaries as a proof of concept. In Figure 33, four surface plots for increasing values of T_{11} while identical regions of parameter space in T_{12} and T_{22} are shown. As the value of T_{11} increases, the highest minmax value peak (the orange ridge in Figure 33 A) shifts toward lower values of T_{22}. For T_{11} values greater than 85 minutes, the largest minmax value calculated in the region begins to drop off smoothly.

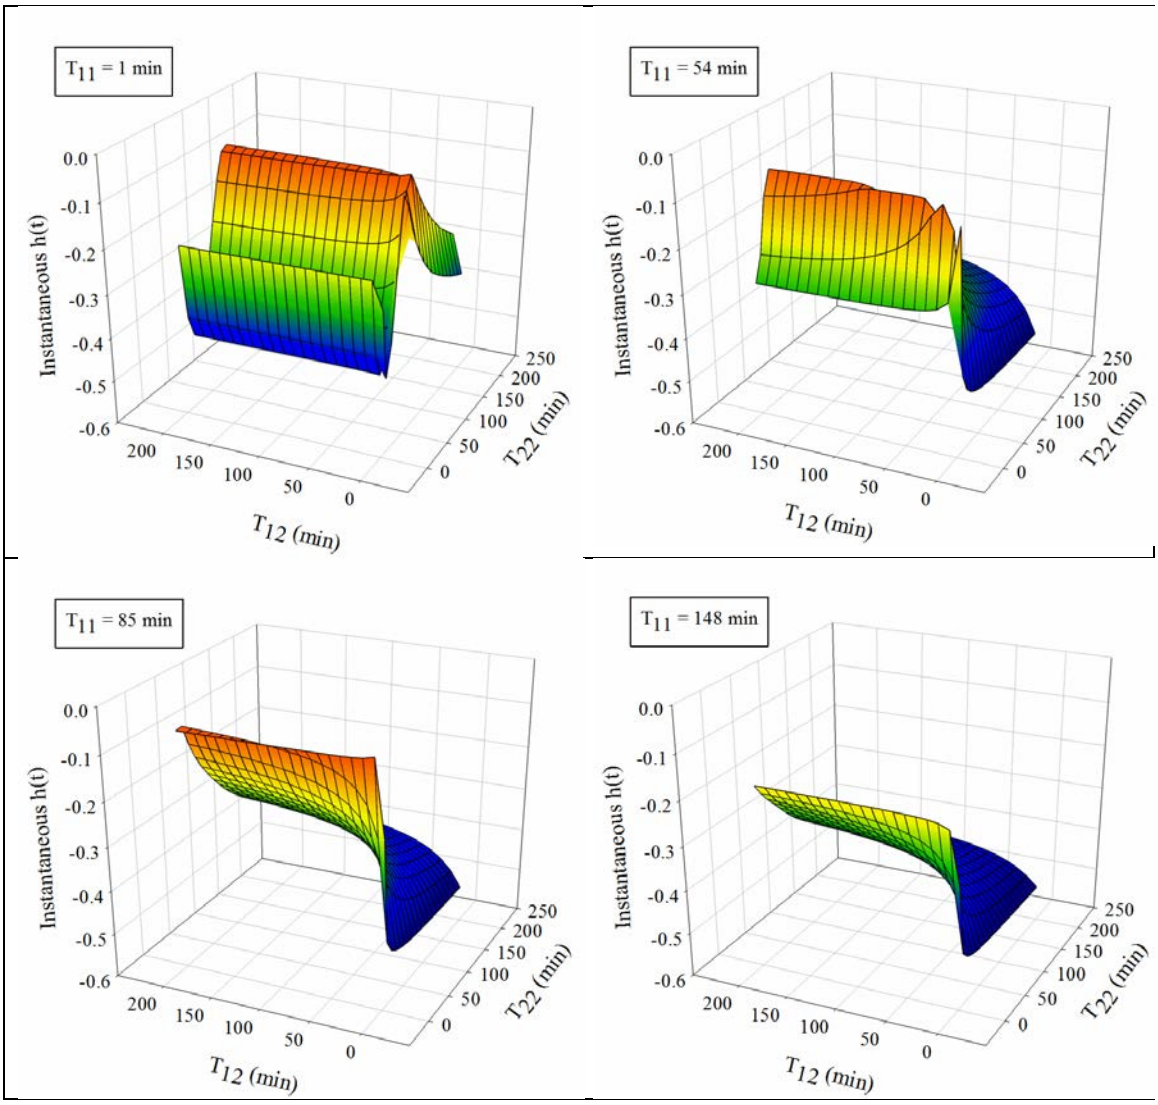

Figure 33: Plot for PDB with T_{11} adjusted between panels. The minmax value never crosses the 0.0 threshold, indicating model failure for the entirety of the explored parameter space.

In Figure 33 no allowable parameters are identified, and for the initial parameter sweep model none were identified for any model parameterization. Our group successfully identified solutions for the PKPDB model using model LL maximization (the model exploration is not included in this dissertation). To show the allowable region for

this model, we performed a minmax sweep in the vicinity of the maximum identified LL parameter set, shown in Table 18. T_{11} was held fixed at the optimal value (3.25 minutes) while T_{12} and T_{22} were adjusted as depicted in Figure 34. We used a 21 point sweep for T_{12} between 1 and 50 minutes and for T_{22} between 200 and 500 minutes. The black surface in the figure depicts the minmax value of 0. Areas of the colored minmax depicted above the black surface are allowed parameters, while parameters below the surface will result in SMF. The figure depicts a curved area of allowable parameters for small values of T_{22}.

Table 18: PKPDB maximum LL parameter set optimized on the Parker *et al.* data set () with LL = -1169.86 and PDCS = 223.12

Model	g_1 (min^{-1})	g_2 (min^{-1})	T_{11} (min)	T_{12} (min)	T_{22} (min)
PKPDB	2.19 E-3	6.70 E-4	3.25	15.5	249.3

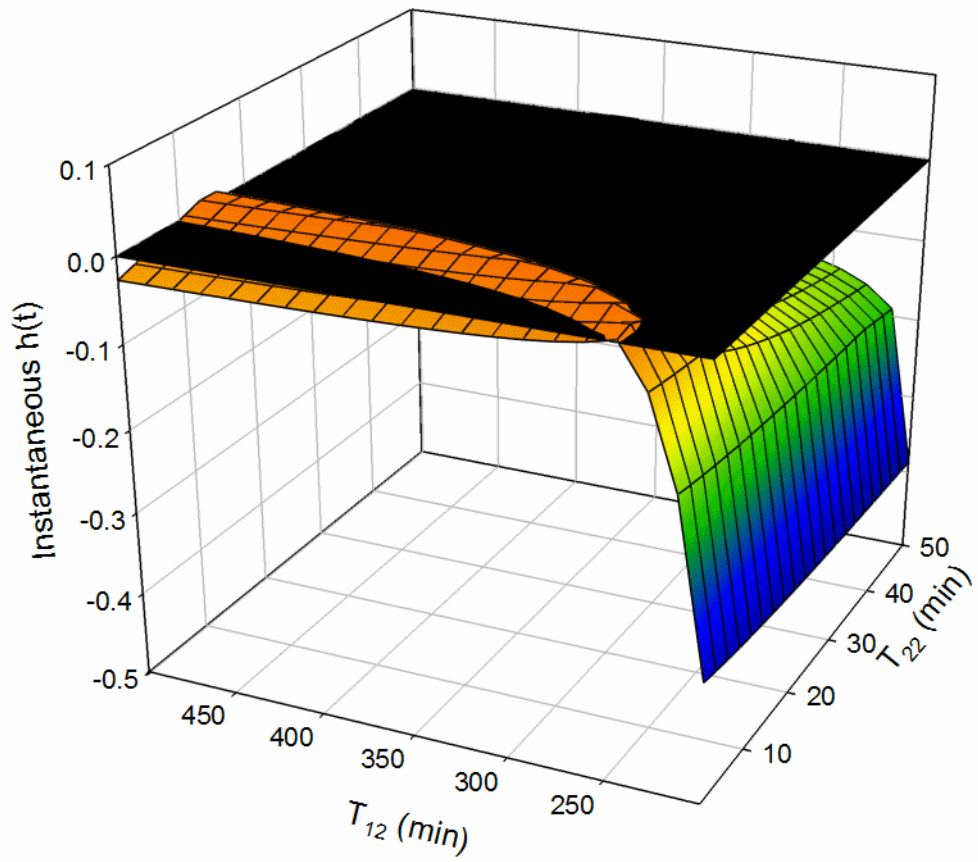

Figure 34: PKPDB allowable parameter region with T_{11} = 3.25 minutes. The colored area depicted above the black surface represents the allowable parameters for the PKPDB model.

Identification of the allowable model parameter space is important because it allows us to identify that the model can be parametrically optimized. If no allowable parameters exist, LL maximization will fail, and time spent optimizing the model will have been wasted. If constrained model optimization is employed in the region of the

allowed space, a local LL maxima could be discovered much more quickly due to bounding the search space to a small area around the maximal LL parameter set.

In our exploration of the PKPDB model, we identified a region of allowable parameters, but only after we refined our parametric sweep to the space immediately around a previously identified solution. In our initial parametric sweep depicted Figure 33, no allowable parameters were identified. This point illustrates an important limitation of the parametric sweep of the minmax surface – the resolution of the sweep can significantly impact identification of allowable surfaces. Without sufficiently high resolution parameter sweeps, small regions of allowable parameters can be overlooked.

5.3.1.4 Investigation of the PKCS2T3 Model Parameter Space

Investigation of the PKCS2T3 model using the four dimensional parameter sweep method revealed an interesting allowable parameter surface. PKCS2T3 is a three compartment pharmacokinetic model which has a feedback loop from the third compartment back into the first, defined by the set of differential equations

$$\begin{aligned}
\frac{dP_{T_1}}{dt} &= T_{11}^{-1}\left(P_{T_A} - P_{T_1}\right) + T_{33}^{-1}\left(P_{T_3} - P_{T_1}\right) \\
\frac{dP_{T_2}}{dt} &= T_{22}^{-1}\left(P_{T_1} - P_{T_2}\right) \\
\frac{dP_{T_3}}{dt} &= T_{33}^{-1}\left(P_{T_2} - P_{T_3}\right)
\end{aligned} \tag{32}$$

Parameter analysis ranged between 1 minute to 400 minutes for each of the four model parameters (T_{11}, T_{22}, T_{33}, and T_{13}) using a 50 minute increment. Four example figures for the PKCS2T3 model are shown in Figure 35. On the abscissa and the ordinate axes the

model parameters T_{13} and T_{33} are shown, while the model parameter T_{11} changes between subplots. For simplicity, each plot used the same T_{22} time of 401 minutes. Other PKCS2T3 minmax surfaces for T_{22} ranging from 1 minute to 401 minutes were overall graphically similar. Allowable parameter regions are shown in the T_{11} = 1 and T_{11} = 41 but were not identified for values of T_{11} greater than 41 minutes. Values of the gas transfer parameter T_{13} of less than 41 minutes never yielded allowable parameters. There was a pronounced model surface shoulder for all parameter sweeps that occurred around T_{13} of approximately 81 minutes. Beyond this shoulder, the effect of increasing T_{13} on the output minmax seemed to diminish.

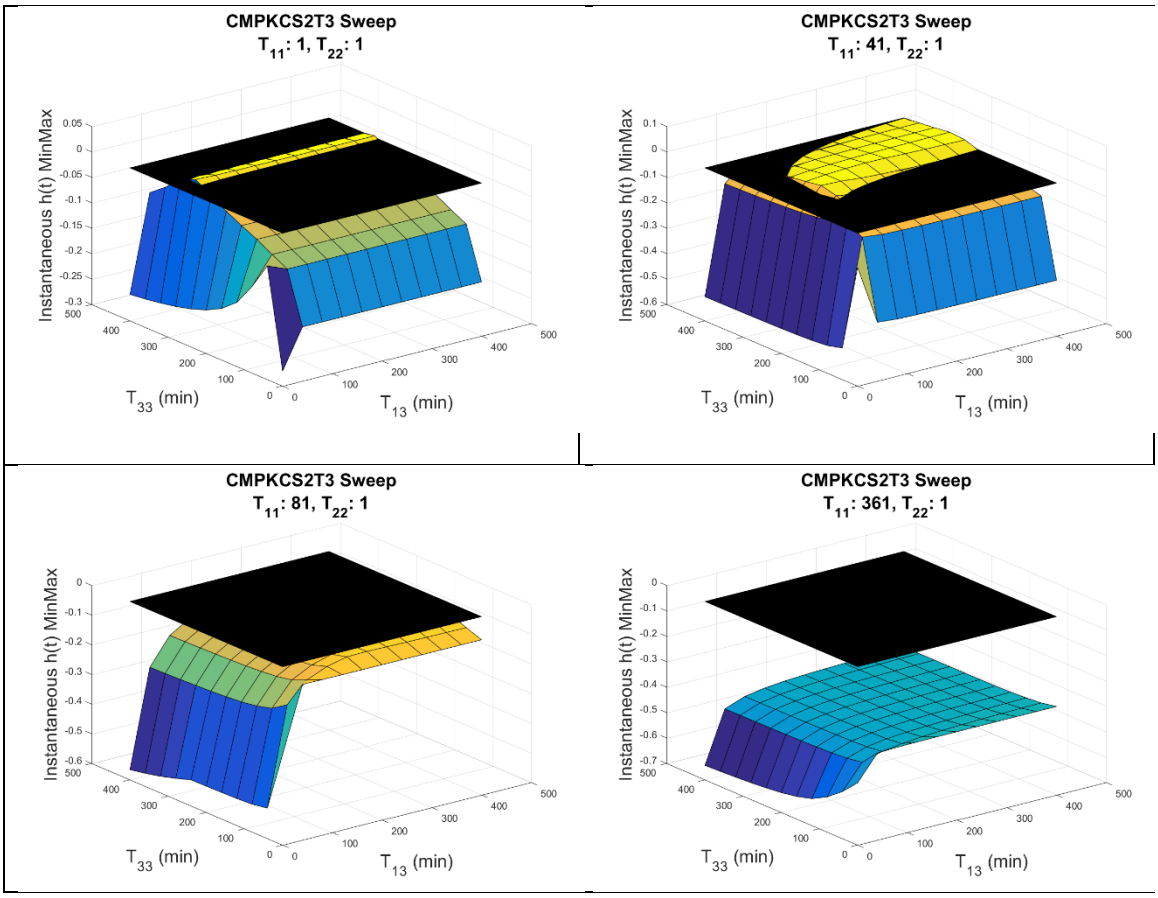

Figure 35: PKCS2T3 SMF plots for T_{11} ranging from 1 minute to 281 minutes. The allowable region for this model changes shape from a small band of allowable parameters to a large region for various values of T_{11}. The value of T_{22} was found to have a less significant effect on the shape of the allowable surface than the other parameters so the same value of 1 minute was used for each plot.

In addition to simply investigating the general parameter space, we investigated the region close to a maximized parameter set for the PKCS2T3 model. The optimization was performed using the NM algorithm described in Chapter 4 on the data set used by Parker *et al* described earlier. The maximal LL parameter set identified can be seen below in Table 19.

Table 19: Maximum LL parameter set for the PKCS2T3 model

Model	g_1 (min^{-1})	g_2 (min^{-1})	g_3 (min^{-1})	T_{11} (min)	T_{13} (min)	T_{22} (min)	T_{33} (min)
PKCS2T3	7.22 E-3	8.83 E-5	6.17 E-4	1.02	38.61	25.17	298.69

In Figure 36, the minmax surface for T_{13} and T_{33} is shown with fixed values for T_{11} and T_{22}. The optimal parameter set is indicated by the white dot. The minmax value of the optimal parameter set is approximately 0.022, which was close to the maximum minmax value for the region of 0.0233. The parameter space around T_{11} and T_{22} was explored and is shown in Figure 37. The maximum identified LL parameter set appears near the highest edge of the minmax surface. Interestingly, the optimized parameter set exists on the boundary of the allowable space with respect to the T_{13} and T_{22} parameters.

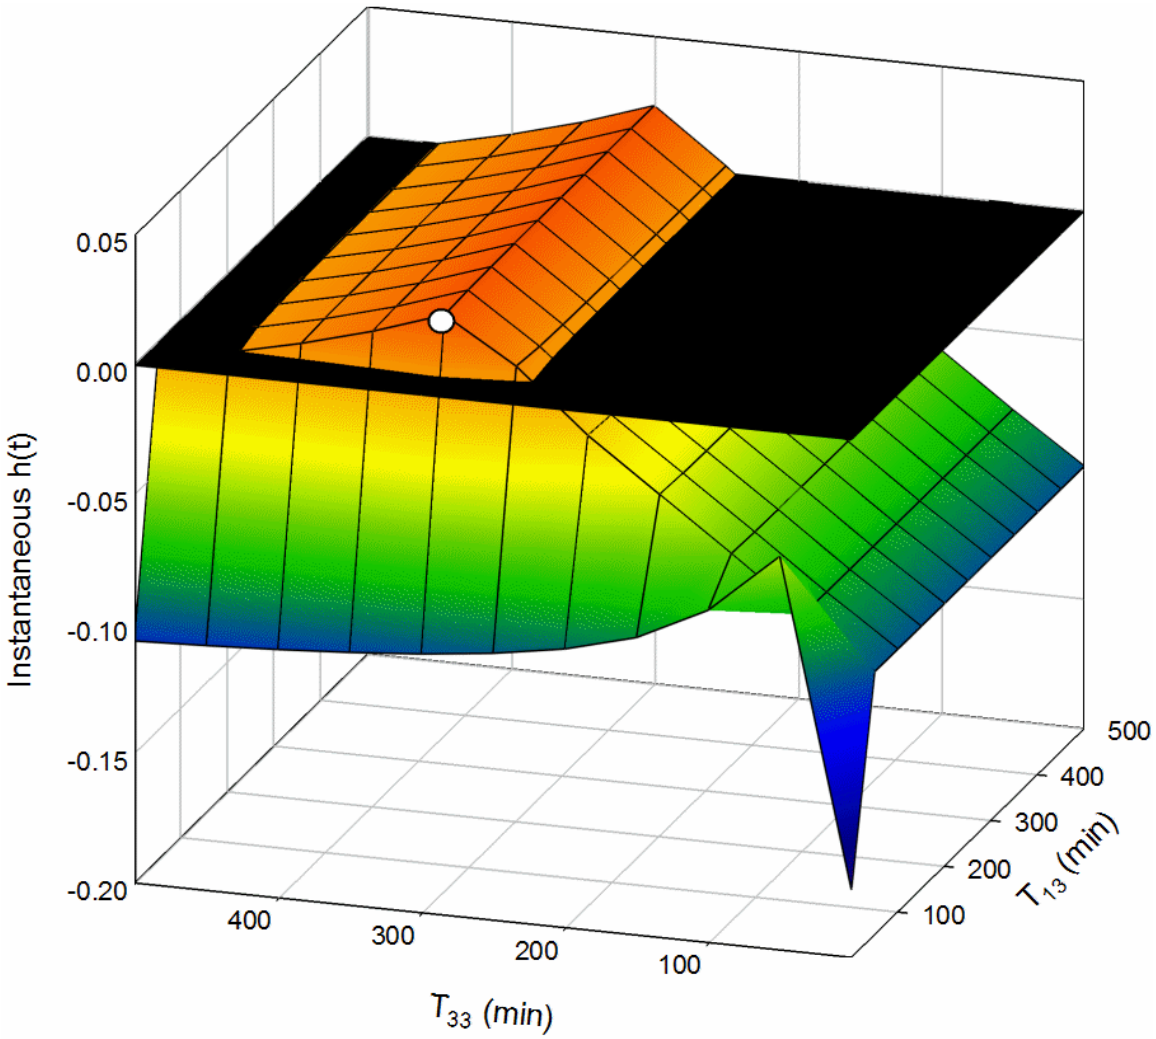

Figure 36: The optimal LL parameter set visualized on the T11 and T22 surface. The minmax value is 0.022 of the optimal point (depicted as a white dot). The optimal point is on the edge of the allowable T_{13} parameter set.

The minmax surface plotted in the T_{11} T_{22} parameter space around the maximum LL identified parameter set can be seen in Figure 37. Similarly to the T_{11} T_{13} parameter space, the location of the maximum LL parameter set is near the peak of the manifold. All parameters in the region were allowable, although the minmax reduction with increasing

T_{11} value indicates that for larger values the region will contain parameter values resulting in SMF. The minmax value for the parameter space falls much more rapidly for changes in the T_{11} parameter than the T_{22}. This change is significantly more rapid than the minmax reduction with decreased T_{33} that can be seen in Figure 36, which drops from 0.022 to 0 in around 100 minutes of change for T_{33}.

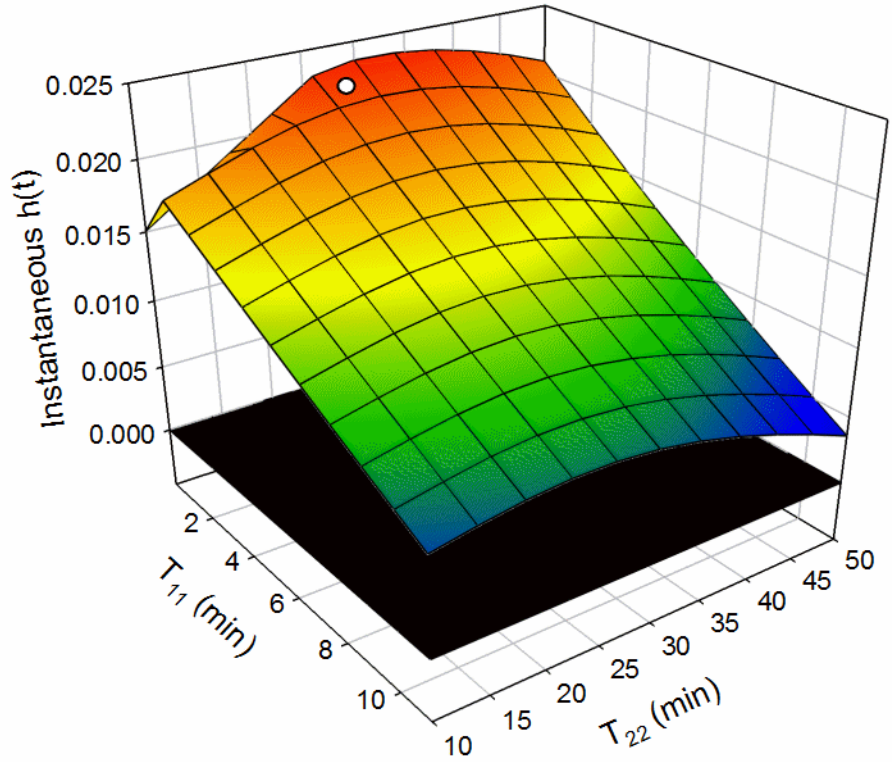

Figure 37: PKCS2T3 T_{11} T_{22} optimal point minmax surface. The white point represents the maximum LL parameter set identified for the model using the NM optimization method described in Chapter 3. Similarly to Figure 36, we find the optimal parameter set near the edge of the allowable parameter space.

Because of the observation that the optimized LL parameter set was located near the boundaries of both the T_{11} parameter and the T_{13} parameter in the 2 dimensional minmax sweeps, it seemed advantageous to explore the T_{11} T_{13} parameter region (Figure 38). The minmax surface beyond $T_{13} \approx 38$ minutes is relatively flat, and the optimal LL parameter set again appears on the edge of the allowable region.

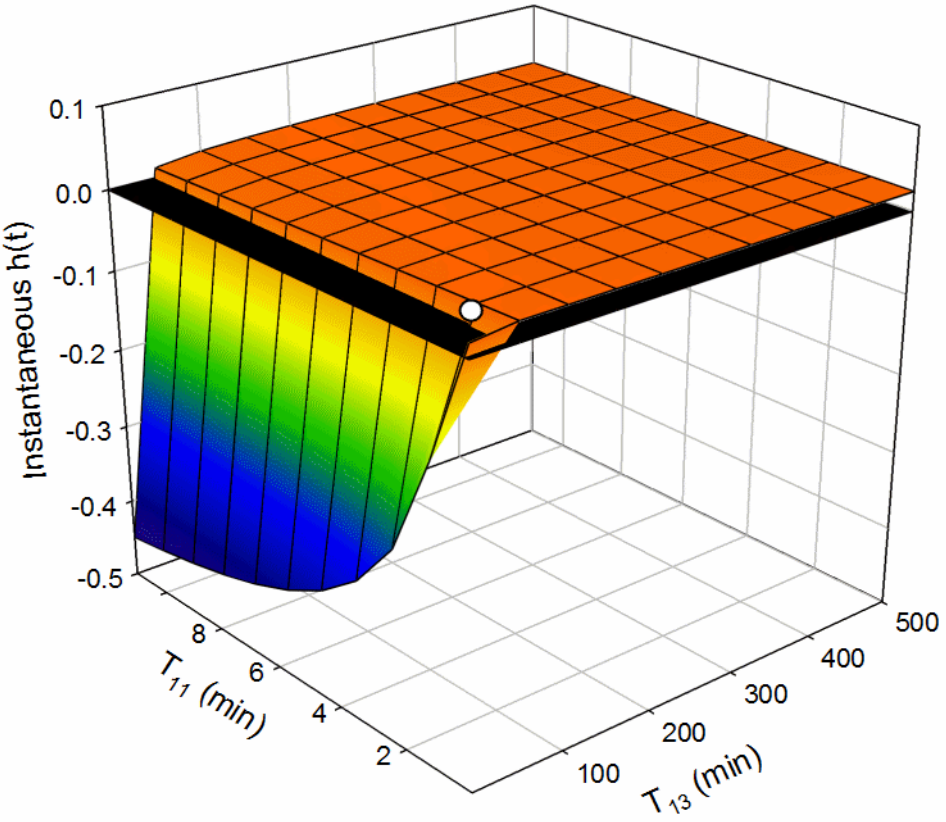

Figure 38: PKCS2T3 T_{11} T_{13} minmax surface. Again, we discover that optimal parameter set (depicted as a white dot) exists near the edge of the allowable space.

Although the minmax defines the region of allowed parameters, it does not necessarily inform about the shape of the likelihood manifold over the region. Investigation of the likelihood manifold around the optimal point revealed that for the PKCS2T3 model, the shape of the manifold seems to be dominated by the T_{11} parameter (see Figure 39). Unexpectedly, the LL value on the manifold decreases nearly linearly in the region explored. The T_{22} parameter seems to have substantially less pronounced effect on the LL hyper-surface, even though the range of the parameter space explored for T_{22} was four times larger than that of T_{11}. The general shape of the T_{11} T_{22} minmax surface and LL surface are similar, with a general downward trend with increasing T_{11} value and approximately flat changes with respect to T_{22}. Because of the need for fast compartment times to handle profiles with rapid decompressions, decreasing the transfer rate of the fast compartment should reduce the model predicted risk on those dive profiles and thus penalize the model LL value.

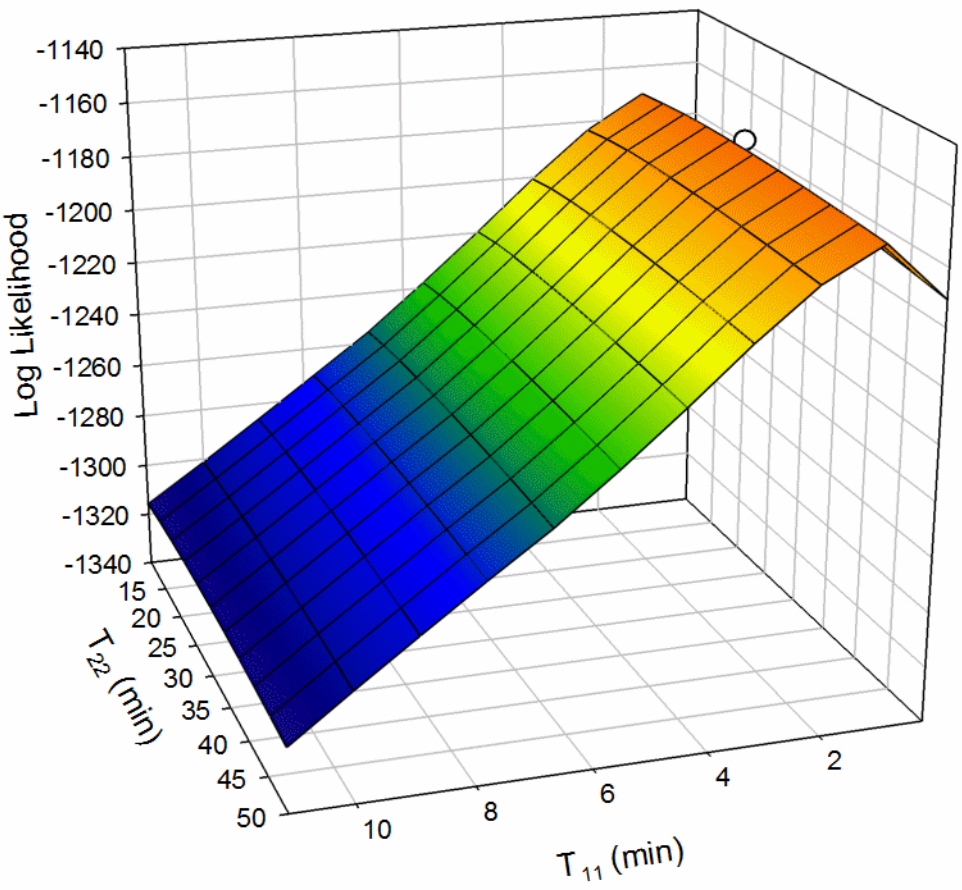

Figure 39: PKCS2T3 T_{11} T_{22} LL surface. The surface is largely consistent for the values of T_{22} explored, while T_{11} appears to have a more pronounced effect on the value of the LL.

5.3.2 SMF Parameter Bisection Solver

5.3.2.1 PKPLBX2

PKPLBX2 is a two parallel compartment system defined as

$$\begin{aligned} \frac{dP_{T_1}}{dt} &= T_{11}^{-1}\left(P_{T_A} - P_{T_1}\right) \\ \frac{dP_{T_2}}{dt} &= T_{22}^{-1}\left(P_{T_A} - P_{T_2}\right) \end{aligned} \tag{33}$$

A graphical representation of the parametric space for the two parameter PKPLBX2 model can be seen in Figure 40. The two compartmental rate parameters T_{11} and T_{22} are represented on the abscissa and ordinate axes, and the instantaneous hazard function minmax is represented in the applicate. On the left the minmax surface for the time of onset hazard definition is depicted and on the right the incidence only minmax surface is shown. In our earlier sweep of the PKPLB model, we discovered that no parameter would avoid SMF for the model using tissue times between 1 and 500 minutes on the NMRI98 data set. This finding explains the lack of allowable parameters for identical T_{11} and T_{22} values for the PKPLBX2 model (shown in Figure 40). For nearly identical parameters compartmental uniqueness is lost and predicted risk for both compartments becomes effectively equal. This suggests that PK models with disconnected compartments should have tissue time distributions that do not overlap so as to reduce the size of the parameter search space and increase model optimization throughput.

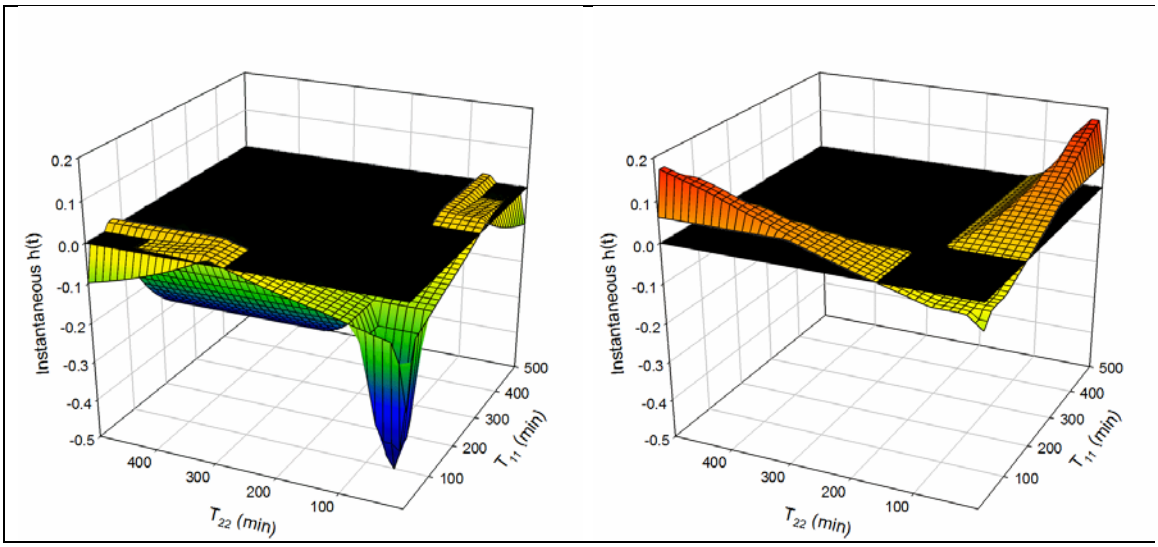

Figure 40: Comparison of incidence only (right) vs time of onset (left) allowable parameter regions on the PKPLBX2 model. The allowable region (the surface above the black plane) is much larger for the incidence only risk definition than for the time of onset version, resulting in reduced difficulty in identifying allowable parameters.

The SMF parameter bisection solver method was implemented on the PKPLBX2 model with a tissue time ranging from 1 minute to 500 minutes and a step size of approximately 40 minutes. Using the results from our parameter minmax sweep, we selected initial SMF boundary parameter estimates of 250 and 400 minutes. Although the solver would have identified the SMF boundaries without initial parameter estimates as close to the true values as what we selected, computation time is greatly reduced if relatively close estimates are selected. In Figure 41, the identified SMF points have been overlaid on the minmax surface for the total sweep time. SMF parameters identified using the bisection solver are depicted as white dots on the plane. The solved SMF boundary points fall on the black surface where they were expected to be found. In the SMF region

between the two allowable surfaces, no SMF parameters were identified. In Figure 41, no SMF points were identified along the SMF boundaries for fixed values of $T_{11.}$ This is due to the lack of uniqueness for SMF parameter values along those edges of parameter space. Overall, the solver worked as expected and located SMF boundaries effectively across the 2-D sweep of the PKPLBX2 model.

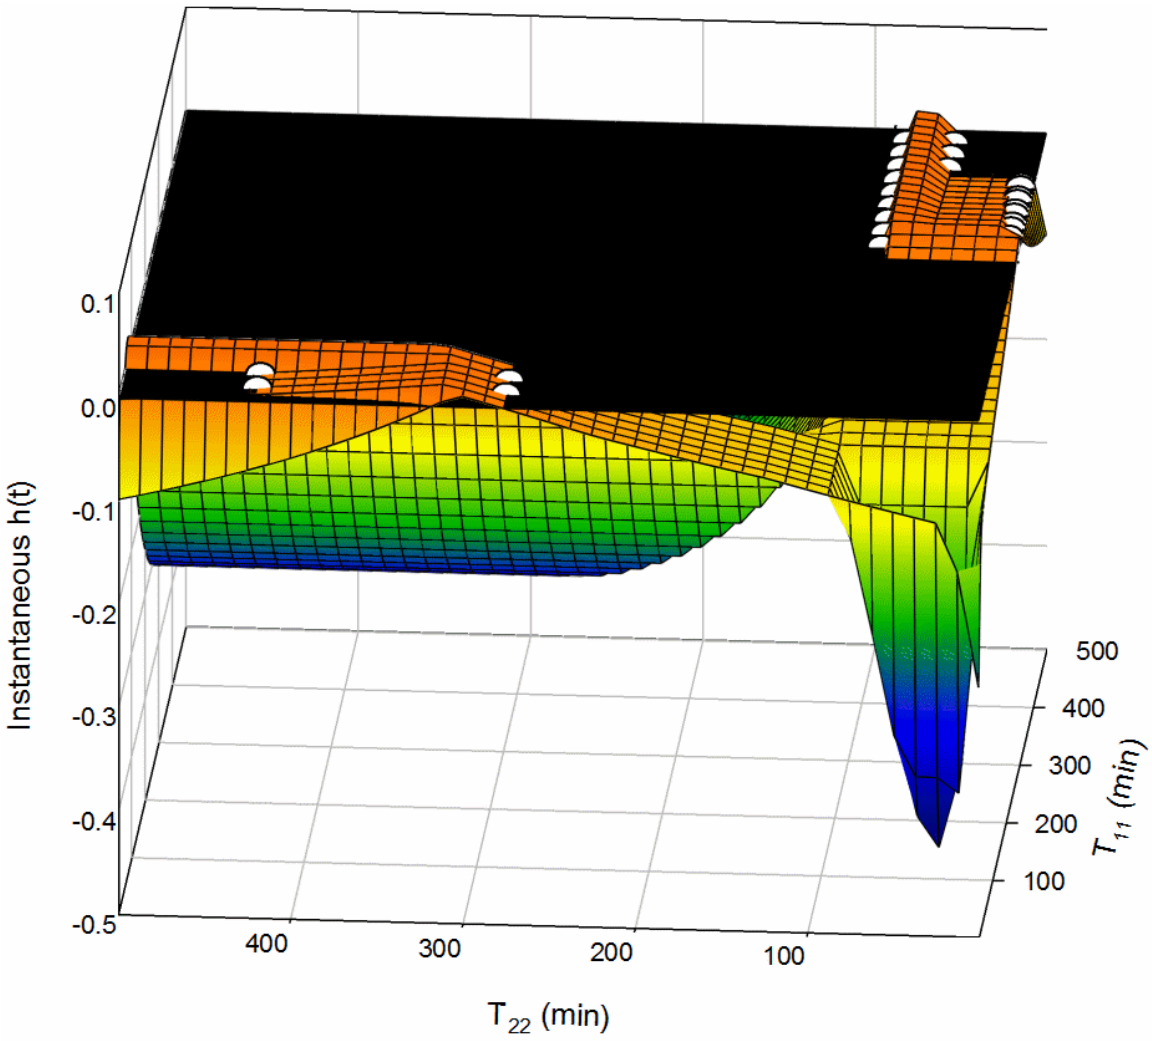

Figure 41: PKPLBX2 SMF Parameter boundary. White dots indicate SMF parameters solved for T_{22} while T_{11} was swept from 1 to 500 minutes. The bisection solver was implemented by solving for SMF parameters on T_{22} while T_{11} was swept and successfully identified parameters on both sides of the allowable region.

5.3.2.2 PKPLBX3

PKPLBX3 is a three compartment pharmacokinetic model defined by Eq.(24). SMF boundaries for the PKPLBX3 model were identified using the bisection solver. To aid in visualization, parameter sweeps for the first and second compartments were performed using tissue times ranging from 1 minute to 1000 minutes. Figure 42 shows a cross section of the solution space for all three compartments with the minmax value indicated by color where, for simplicity, only one permutation of the parameter space is displayed visually. The compartmental independence of PKPLBX3 causes parameter permutations which contain identical parameter values to have identical minmax values. No allowed tissue times for T_{11} were identified beyond approximately 75 minutes or for T_{33} times less than 300 minutes in the 1 to 1000 minute range explored. The red ball in the bottom right of Figure 42 indicates the maximum LL parameter set identified using the optimization scheme outlined in Chapter 4. As with the PKCS2T3 model, we observe that the maximal LL parameter set exists near an edge of allowable space in multiple dimensions, providing further evidence that the optimal points for models often exist near the boundary of the allowable parameter region.

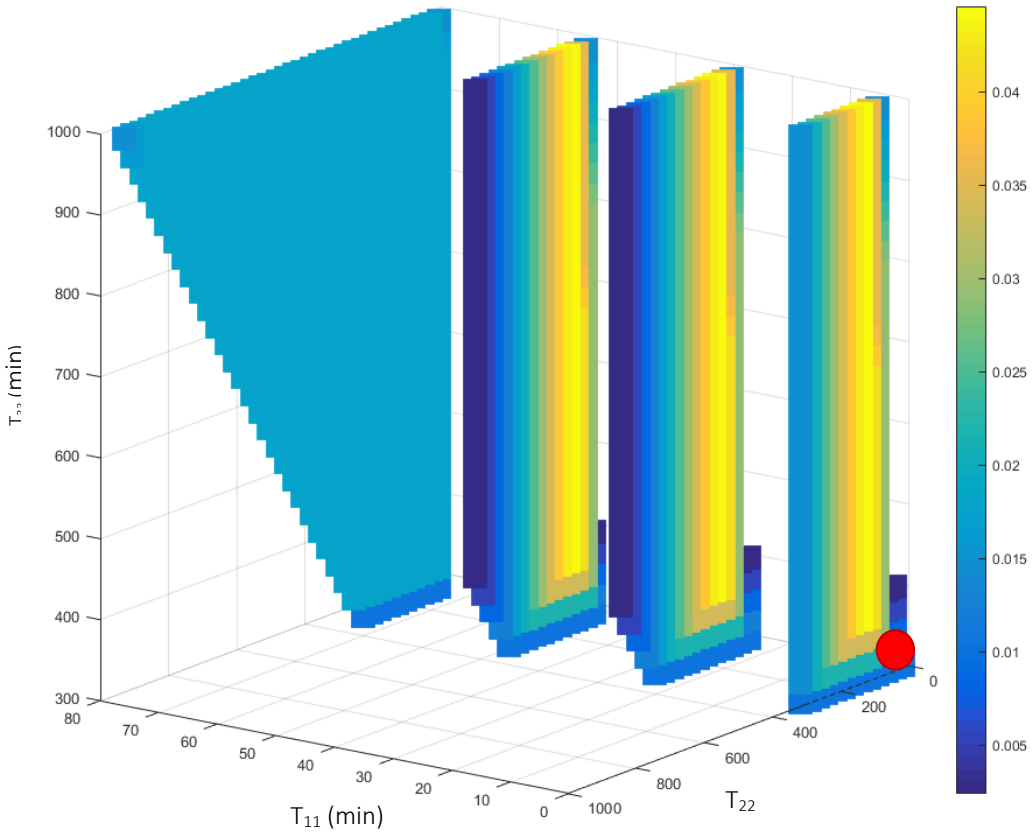

Figure 42: Allowable region of PKPLBX3 with minmax value indicated by color. The red dot in the lower left region indicates the maximum LL parameter set identified for the model (minmax = 0.0125).

The bisection SMF parameter solving routine was applied to the PKPLBX3 model with a fixed value for T_{33} of 600 minutes (arbitrarily chosen for illustrative purposes). Figure 43 displays the identified SMF boundary as a black line, while the colored regions depict the minmax value of the allowed parameter region. The SMF parameter method identified both edges of the allowable parameter space (one around T_{22} = 400 minutes and

one with T_{22} = 50 minutes). The SMF boundary solver correctly identified both edges for the parameter range explored.

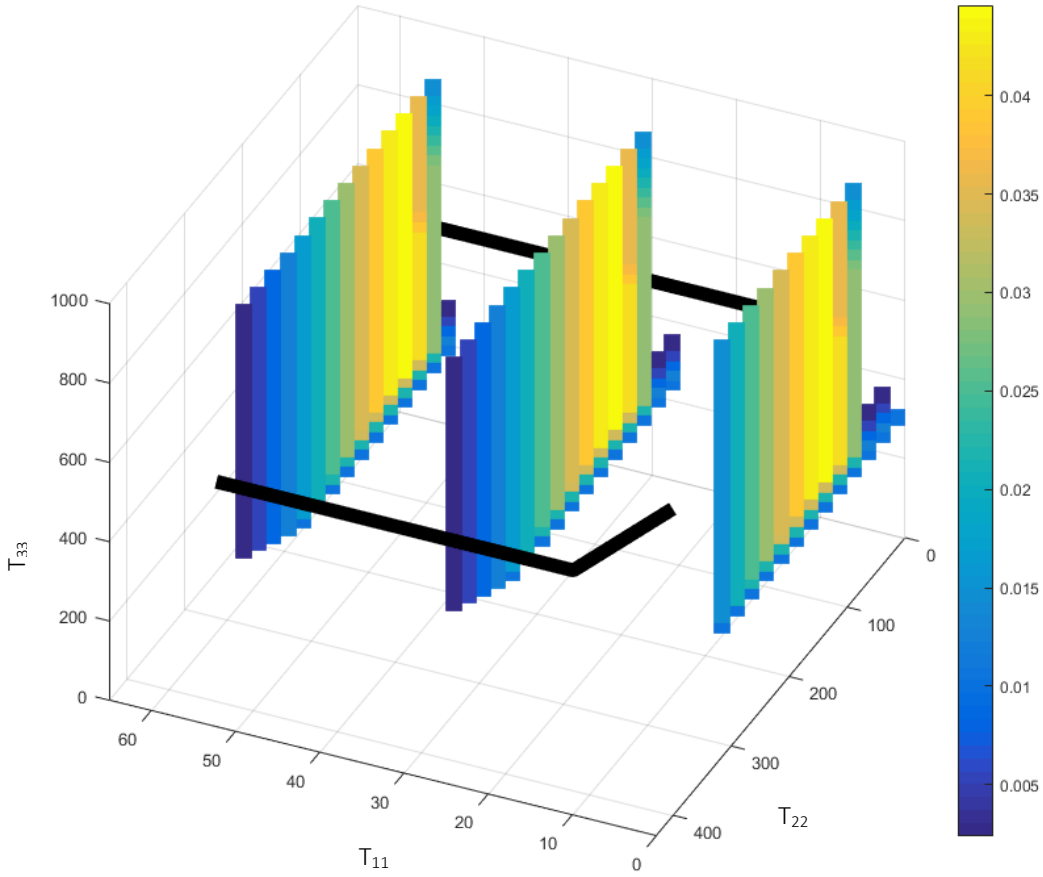

Figure 43: Depiction of the solved SMF parameter using the bisection method on PKPLBX3 with T_{33} = 600 minutes. The black line tracing the edge of the allowable space depicts the identified SMF boundary.

5.4 Conclusion

PK model SMF boundaries investigated using the methods developed within this work revealed that the overall shape of the allowable parameter spaces is often simple. SMF parameter boundaries were frequently identified as straight edges for one fixed parameter value (see Figure 40, Figure 35, Figure 36, Figure 38). No allowable parameter regions for the single compartment PK model was not unexpected, as mono-exponential models have been previously shown to have model failures for two profiles with T_1 times greater than 2 hours [19]. Results from this investigation suggest that the more parameters a model has, the greater the chance it will have some parameterization that will be allowable, however this hypothesis was not exhaustively tested.

We also identified that for at least two models (PLBX3 and PKCS2T3), the maximum identified LL parameter set existed on the boundaries of the allowable parameter region, a result has been suggested in the past [22]. This finding could be highly significant from a model optimization standpoint, as identification of the allowable parameter regions could simultaneously inform the location of maximum LL parameters. It is unclear if this finding will extend to additional PK model types and should be investigated further.

Although the SMF boundary identification techniques presented in this work cannot definitively rule out the possibility of either very small allowable parameter ranges or allowable parameters outside the tested parameter intervals, we believe that the SMF

boundary identification techniques offer investigators a useful method for quickly identifying the allowable parameter space (if one exists), reducing the difficulty for model optimization. It is important to state that while the maximum log likelihood parameter set must exist within the allowable region, identification of the boundaries of the region informs investigators little about what the shape of the likelihood hyper-surface is within it. Parametric optimization for probabilistic pharmacokinetic DCS models currently represents a significant hurdle in the examination of new candidate models. Substantial computational resources are still required to find optimal LL model parameters, regardless of the optimization scheme used. Through examination of parameter minmax values, the effort to determine viability of candidate models with respect to SMF may be greatly reduced. If allowable parameters are identified, the parameter search space can be significantly restricted by performing parametric optimization on only allowable regions.

6. Concluding Remarks

The objective of this exploration was to aid in the advancement of probabilistic decompression models. Chapters 1 and 2 introduced the motivation behind the work detailed in this dissertation and provided relevant background material necessary for understanding probabilistic models applied to predicting the incidence of decompression sickness. Our motivations for the reminder of the dissertation focused on improving the predictive quality of probabilistic decompression models using functional modifications and reducing the difficulty of model parameter optimization through the identification of allowable parameter regions.

In Chapter 3, we systematically compared the performance of 16 hazard functions variants using the EE model optimized using a combination of Nealder-Mead and simulated annealing optimization algorithms with the BIG292 training data set. Predictive quality was analyzed using Xi square testing by segmenting profiles by both maximum depth and by dive type. We were unable to identify an individual hazard function that outperformed alternatives across all dive groups examined. Using AICc model selection, a weighted multimodel was developed using the predictions of two selected models. Prediction improvements using the hazard function according to AICc model selection indicated that DSR, a common hazard function for simple gas content probabilistic decompression models, is a suboptimal hazard functional form for the EE model type, but overall improvements to the predictive quality of the EE model were modest. We

identified a hazard function with ambient pressure scaling and individually fitted exponential parameters (DSRE111N) was the best overall functional form. Performance between the explored models and previously published US Navy probabilistic models was not identified to be statistically significant, suggesting predictive quality between existing models and our explored model forms may be similar.

Explicit delay mechanics were explored using delay differential equations in a pharmacokinetic modeling framework. Ten models were developed which included a fitted delay component between the arterial gas input and the compartmental system. Delay models were compared to compartmental structures that did not have explicit delay mechanics using AICc model selection. Results selected two of the model forms PKDDES3T and PKDDES2LP with 50% and 46% of the total weight respectively. Comparison of delay models using likelihood ratio testing indicated that explicit delays were not statistically justified when used for modeling the entire NMRI98 data set. Additionally, comparison to the AICc selected model PKCS2T3 (that did not include delay mechanics) indicated that the addition of delay was not advantageous for model parsimony with respect to the data. Overall, the addition of delay mechanics on the arterial inputs for our pharmacokinetic models was found to be statistically unjustified.

The parameter bounding techniques demonstrated in Chapter 5 can be used to reduce the time required to find optimized solutions, allowing more rapid model investigations. Using the parameter minmax, a value related to the instantaneous hazard

function, we developed two methods for exploring the allowable parameter space and demonstrated both on a series of previously published pharmacokinetic models. We identified that the allowable region for several of our pharmacokinetic models appears to be simple, with parametric boundaries of a single value being discovered for most of the models explored. Additionally, we identified maximum LL parameter sets near the edge of the allowable parameter space for two models, lending support to the hypothesis that optimal LL parameters may often exist near a model failure boundary. Identification of these boundaries can be used to reduce the computational time required for model fitting, allowing future models to be more quickly explored.

Our underlying objective in the course of this work was the development of techniques to improve both the predictive quality and the ease of parametric optimization of probabilistic decompression models. Our investigation of hazard function scaling indicated that significant improvements for predictive quality can be made to the EE model through the employment of a function that used individually fitted nonlinear exponential scaling terms. Our implementation of delay mechanics in a group of pharmacokinetic models revealed that explicit delay applied to the arterial input did not improve the quality of model predictions. Two techniques were developed that can be used to identify the allowable parameter region and determine its boundaries were presented through demonstrations on several pharmacokinetic models. The authors

believe that the works documented within this dissertation demonstrate an advancement of probabilistic models of decompression sickness.

References

1. James, T., R. Francis, and S.J. Mitchell, *Manifestations of decompression disorders.* Bennett and Elliott's physiology and medicine of diving. 5th ed. London: Saunders, 2003: p. 578-599.

2. Murphy, F.G., E.A. Hada, and L.E. Howle, *The Case for Mixed Pharmacokinetic Models as a Descriptor of Decompression Sickness*, in *UHMS Scientific Meeting*. 2015: Montreal.

3. Vann, R.D., et al., *Resolution and Severity in Decompression Illness.* Aviation Space and Environmental Medicine, 2009. **80**(5): p. 466-471.

4. Weathersby, P.K., et al., *Statistically Based Decompression Tables. I. Analysis of Standard Air Dives: 1950-1970.* 1985.

5. West, V.R. and J. Parker, J. F., *A Review of the Influence of Physical Condition Parameters on a Typical Aerospace Stress Effect: Decompression Sickness.* 1973.

6. Workman, R.D., *Calculation of Air Saturation Decompression Tables*. 1957, DTIC Document.

7. Workman, R.D., *Calculation of decompression schedules for nitrogen--oxygen and helium- oxygen dives. Res Rep 6-65.* Rep US Navy Exp Diving Unit, 1965: p. 1-33.

8. Boycott, A.E., G.C. Damant, and J.S. Haldane, *The Prevention of Compressed-air Illness.* J Hyg (Lond), 1908. **8**(3): p. 342-443.

9. Tikuisis, P. and W.A. Gerth, *Decompression theory.* Bennett and Elliott's physiology and medicine of diving, 2003.

10. Weathersby, P.K., L.D. Homer, and E.T. Flynn, *On the likelihood of decompression sickness.* J Appl Physiol, 1984. **57**(3): p. 815-25.

11. Howle, L.E., P.W. Weber, and R.D. Vann, *A computationally advantageous system for fitting probabilistic decompression models to empirical data.* Computers in Biology and Medicine, 2009. **39**(12): p. 1117-1129.

12. Thalmann, E.D., et al., *Improved probabilistic decompression model risk predictions using linear-exponential kinetics.* Undersea Hyperb Med, 1997. **24**(4): p. 255-74.

13. Gerth, W. and R. Vann, *Probabilistic gas and bubble dynamics models of decompression sickness occurrence in air and nitrogen-oxygen diving.* Undersea & hyperbaric medicine: journal of the Undersea and Hyperbaric Medical Society, Inc, 1996. **24**(4): p. 275-292.

14. Rodríguez, G., *Lecture notes on generalized linear models*. 2010: Princeton University.

15. Howle, L.E., *Analytic gain in probabilistic decompression sickness models.* Comput Biol Med, 2013. **43**(11): p. 1739-47.

16. Weathersby, P.K. and W.A. Gerth. *Survival Analysis & Maximum Likelihood Techniques as Applied to Physiological Modeling*. in *Fifty-first Workshop of the Undersea and Hyperbaric Medical Society*. 2002. Seattle, Washington: Undersea and Hyperbaric Medical Society.

17. Goldman, S., *A new class of biophysical models for predicting the probability of decompression sickness in scuba diving.* J Appl Physiol, 2007. **103**(2): p. 484-93.

18. Tikuisis, P., R.Y. Nishi, and P.K. Weathersby, *Use of the maximum likelihood method in the analysis of chamber air dives.* Undersea Biomed Res, 1988. **15**(4): p. 301-13.

19. Weathersby, P.K., et al., *Predicting the time of occurrence of decompression sickness.* J Appl Physiol, 1992. **72**(4): p. 1541-8.

20. Hugon, J., *Decompression models: review, relevance and validation capabilities.* Undersea & hyperbaric medicine: journal of the Undersea and Hyperbaric Medical Society, Inc, 2013. **41**(6): p. 531-556.

21. Homer, L. and P. Weathersby, *Statistical aspects of the design and testing of decompression tables.* Undersea biomedical research, 1985. **12**(3): p. 239-249.

22. Weber, P.W., *Probabilistic Modeling of Decompression Sickness, Comparative Hydrodynamics of Cetacean Flippers, Optimization of CT/MRI Protocols and Evaluation of Modified Angiocatheters: Engineering Methods Applied to a Diverse Assemblage of Projects*. 2010, Duke University.

23. Weathersby, P.K. and K.A. Gault, *Decompression sickness predictive models for unsafe human exposure.* Undersea Hyperb Med, 2013. **40**(3): p. 247-66.

24. Weathersby, P.K., et al., *Role of oxygen in the production of human decompression sickness.* J Appl Physiol, 1987. **63**(6): p. 2380-7.

25. Loveman, G. and e. al., *Submarine Tower Escape Decompression Sickness Risk Estimation.* Journal of the Undersea and Hyperbaric Medical Society, 2014. **41**(4): p. 315-329.

26. Press, W.H., et al., *Numerical recipes in C*. Vol. 2. 1996: Citeseer.

27. Press, W.H. and e. al., *Numerical recipes : the art of scientific computing*. 3rd ed. 2007, Cambridge, UK ; New York: Cambridge University Press. xxi, 1235 p.

28. Bertsimas, D. and J. Tsitsiklis, *Simulated Annealing.* Statistical Science, 1993. **8**(1): p. 10-15.

29. Ingber, L., *Simulated Annealing - Practice Versus Theory.* Mathematical and Computer Modelling, 1993. **18**(11): p. 29-57.

30. Kirkpatrick, S., C.D. Gelatt, and M.P. Vecchi, *Optimization by Simulated Annealing.* Science, 1983. **220**(4598): p. 671-680.

31. Howle, L.E., et al., *Marginal DCS events: their relation to decompression and use in DCS models.* Journal of Applied Physiology, 2009. **107**(5): p. 1539-1547.

32. Burnham, K.P. and D.R. Anderson, *Multimodel inference - understanding AIC and BIC in model selection.* Sociological Methods & Research, 2004. **33**(2): p. 261-304.

33. Burnham, K.P., D.R. Anderson, and K.P. Huyvaert, *AIC model selection and multimodel inference in behavioral ecology: some background, observations, and comparisons.* Behavioral Ecology and Sociobiology, 2011. **65**(1): p. 23-35.

34. Wagenmakers, E.J. and S. Farrell, *AIC model selection using Akaike weights.* Psychonomic bulletin & review, 2004. **11**(1): p. 192-196.

35. Singer, S. and J. Nelder, *Nelder-Mead algorithm.* Scholarpedia, 2009. **4**: p. 2928.

36. Burnham, K.P. and D.R. Anderson, *Model selection and multimodel inference : a practical information-theoretic approach*. 2nd ed. 2002, New York: Springer. xxvi, 488 p.

37. Parker, E., et al., *Statistically based decompression tables. VIII. Linear-exponential kinetics.* NMRI Report, 1992: p. 92-73.

38. Doolette, D.J., R.N. Upton, and C. Grant, *Perfusion–diffusion compartmental models describe cerebral helium kinetics at high and low cerebral blood flows in sheep.* The Journal of physiology, 2005. **563**(2): p. 529-539.

39. Murphy, F.G., *Pharmacokinetic Models Incorporating Diffusion as a Predictor of Decompression Sickness in Humans.* 2015.

40. Haddad, W.M. and V. Chellaboina. *Stability theory for nonnegative and compartmental dynamical systems with time delay*. in *American Control Conference, 2004. Proceedings of the 2004*. 2004. IEEE.

41. Kidd, D.J. and R.A. Stubbs, *The use of the pneumatic analog computer for divers.* The Physiology and Medicine of Diving and Compressed Air Work, 1969: p. 386-413.

42. Nishi, R.Y. and G.R. Lauckner, *Development of the DCIEM 1983 Decompression Model for Compressed Air Diving*. 1984, Defense and Civil Inst. of Environmental Medicine Downsview (Ontario).

43. Bove, A.A., *Bove and Davis' diving medicine*. 2004: WB Saunders Philadelphia.

44. Jacquez, J.A. and C.P. Simon, *Qualitative theory of compartmental systems with lags.* Mathematical Biosciences, 2002. **180**(1): p. 329-362.

45. Jacquez, J.A. and C.P. Simon, *Qualitative theory of compartmental systems.* Siam Review, 1993. **35**(1): p. 43-79.

46. Jacquez, J.A., *Compartmental analysis in biology and medicine*. 2nd ed. 1985, Ann Arbor: University of Michigan Press. x, 560 p.

47. Murphy, F.G., et al., *Perfusion-Diffusion Gas Content Compartmental Models As A Precitor Of Decompression Sickness.* 2014 UHMS Annual Scientific Meeting, 2014.

48. Parker, E., et al., *Probabilistic models of the role of oxygen in human decompression sickness.* Journal of Applied Physiology, 1998. **84**(3): p. 1096-1102.

49. Allgower, E.L. and K. Georg, *Numerical continuation methods: an introduction*. Vol. 13. 2012: Springer Science & Business Media.

50. Kubicek, M. and M. Marek, *Computational methods in bifurcation theory and dissipative structures.* Springer Series in Computational Physics. Springer Verlag, New York-Berlin-Heidelberg-Tokyo, 1983. **1**.

51. Larson, R., R.P. Hostetler, and B.H. Edwards, *Calculus of a Single Variable*. 2005: Houghton Mifflin College Division.

Biography

Ethan Alexander Hada was born in Putnam, CT on the 11th of July, 1989. His family moved to Toms River, NJ shortly after his birth, where he lived until he turned eighteen. Ethan obtained a bachelor's degree in Applied Physics at Stockton University (formerly The Richard Stockton College of New Jersey) in May of 2012. In July 2012, Ethan moved on to pursue a doctorate in Mechanical Engineering under the supervision of Dr. Laurens E. Howle at Duke University.

In the course of his works at Duke Univeristy, Ethan presented graduate student seminars including "Geometric Intuitions of 'EigenThingys and an Introduction to Principal Component Analysis" and "Probabilistic Modeling of Decompression Sickness: From Haldane to Howle and Hada." Ethan presented the abstract "Depth Scaling of Several Gas Super-Saturation Decompression Models" and gave a talk of the same title, and published the abstract "Delay Differential Equations as an Explicit Method of Aligning DCS Model Prediction with Dive Trial Outcome" at UHMS 2015. At the time of the writing of this document, Ethan was an author on three publications: "The Probability and Severity of Decompression Sickness" (in press), "Hazard Function Scaling for Probabilistic Decompression Models – Part 1 Basic Gas Content Scaling"

(submitted), and "Hazard Function Scaling for Probabilistic Decompression Models – Part 2 Exponential Gas Content Scaling" (submitted).