Generalized Sensitivities and Optimal Experimental Design

H. T. Banks, Sava Dediu, Stacey L. Ernstberger and Franz Kappel
Center for Research in Scientific Computation
North Carolina State University
Raleigh, NC 27695-8205
and
Institute for Mathematics and Scientific Computation
University of Graz
Graz, Austria A8010

November 11, 2009

Abstract

We consider the problem of estimating a modeling parameter θ using a weighted least squares criterion $J_d(y,\theta) = \sum_{i=1}^{n} \frac{1}{\sigma(t_i)^2}\big(y(t_i) - f(t_i,\theta)\big)^2$ for given data y by introducing an abstract framework involving generalized measurement procedures characterized by probability measures. We take an *optimal design* perspective, the general premise (illustrated via examples) being that in any data collected, the information content with respect to estimating θ may vary considerably from one time measurement to another, and in this regard some measurements may be much more informative than others. We propose mathematical tools which can be used to collect data in an *almost optimal* way, by specifying the duration and distribution of time sampling in the measurements to be taken, consequently improving the accuracy (i.e., reducing the uncertainty in estimates) of the parameters to be estimated.

We recall the concepts of *traditional* and *generalized sensitivity functions* and use these to develop a strategy to determine the "optimal" final time T for an experiment; this is based on the time evolution of the sensitivity functions and of the condition number of the Fisher information matrix. We illustrate the role of the sensitivity functions as tools in optimal design of experiments, in particular in finding "best" sampling distributions. Numerical examples are presented throughout to motivate and illustrate the ideas.

AMS subject classifications: 62G08, 62H99, 90C31, 65K10, 93B51, 62B10.
Key Words: Least squares inverse problems, sensitivity and generalized sensitivity functions, Fisher Information matrix, design of experiments.

1 Introduction

We assume that we have a dynamical system which models some physical, sociological or biological phenomenon, e.g.,

$$\begin{aligned} \dot{x}(t) &= \mathcal{F}(t, x(t), \theta), \quad x(0) = x_0(\theta), \\ \eta(t) &= h(t, x(t), \theta), \quad t \in [0, T], \end{aligned} \tag{1.1}$$

where $x(t) \in \mathbb{R}^N$ is the vector of state variables of the system, $\eta(t) \in \mathbb{R}^M$ is the vector of observable or measurable outputs, $\theta \in \mathbb{R}^p$ is the vector of system parameters, while $\mathcal{F}$ and h are mappings $\mathbb{R}^{1+N+p} \to \mathbb{R}^N$ and $\mathbb{R}^{1+N+p} \to \mathbb{R}^M$, respectively. Of course, in general $\mathcal{F}$ and h will be defined on appropriate subsets of $\mathbb{R}^{1+N+p}$. Furthermore, we assume throughout that $\mathcal{F}$ and h are sufficiently smooth to carry out the arguments we make. Note that the formulation of (1.1) allows for components of the initial state vector to be considered as part of the system parameters and that additional components of the parameters may also arise in the output function h. The output variables could, of course, also be just (as encountered frequently in practice) some components of the state vector x of the system.

Solving the initial value problem (1.1), we obtain $x = x(t, \theta)$, $0 \le t \le T$, and

$$\eta(t) = f(t, \theta), \quad 0 \le t \le T, \tag{1.2}$$

where f is defined as $f(t, \theta) = h\big(t, x(t, \theta), \theta\big)$. The output model (1.2) will be the basis for the investigations in this paper. It can be also obtained from dynamical systems other than ordinary differential equations such as (1.1); for instance delay systems or partial differential equations may in principle also be treated with our formulation. That η is the vector of observable/measurable outputs of the physical, sociological or biological system implicitly assumes that we have measurements $y(t)$ for $\eta(t)$ at time instances in some subset of $[0, T]$. For the purpose of this paper we assume that we have a single output system, i.e., we have $M = 1$.

There are typically two problems of interest in the context of a model of type (1.1) – (1.2):

- The *forward* or the *direct problem*, which directly manifests the dependence of the output variable η on the parameter vector θ. An important task is to identify those parameters to which the model output is the most/least sensitive.
- The *inverse problem* or *parameter estimation problem*, which consists of estimating the parameter vector θ from a set of time observations y of the output variable η.

The sensitivity of the model outcome to the various parameters provides important information on where to interrogate or observe the system in order to optimize the outputs or to avoid undesirable system outcomes. These are very important questions in simulation studies and *traditional sensitivity functions* (TSF) as defined below are frequently used to investigate them (see [2, 15, 18, 21, 26, 27, 32] and Section 3.1).

In order to solve the parameter estimation problem for a mathematical model of type (1.1) – (1.2) one requires a set of measurements of a quantity in the real system which corresponds to the model output and an error functional for the comparison of model output and data. Finally, one needs a numerical algorithm in order to obtain estimates for the parameter vector θ as minimizers of an error functional, such as for instance least squares or maximum likelihood, etc., (see [14, 16, 22, 24, 33]). In order to ascertain the quality of the parameter estimates one computes standard errors and establishes confidence intervals (see [7, 10, 16]). It is obvious that the quality of the results depends on all the factors involved, namely the quality of the mathematical model, the uncertainty in the data, and the performance of the estimation algorithm.

In this paper we consider the problem of estimating the parameter vector θ from an *optimal design* perspective. Our premise is that the information content on the parameters θ may vary considerably from one time measurement to another. Intuitively, data with a higher information

content with respect to the parameters to be estimated should yield more accurate estimates. Our goal in the following will be to develop and illustrate mathematical tools which can be used to specify the duration, number and distribution of times for sampling data in order to obtain estimates with improved accuracy. The basis for our approach is the model (1.1) – (1.2), which is assumed to model the dynamics of the real system sufficiently well in the sense that time behavior of the output η for some nominal parameter vector θ_0 coincides with the time behavior of the corresponding quantity of the real system (the existence of a "true" parameter assumption typically encountered in statistical analysis). This frequently made assumption has the implication that in the ideal case of measurements without measurement errors, a measurement taken at any time $\tilde{t}$ is given by $y(\tilde{t}) = \eta(\tilde{t}, \theta_0)$. The problem of estimating the "true" or nominal parameter vector θ_0 is formulated as a least squares approximation problem. That is, in order to compare measurements and corresponding model output we use the (weighted) sum of squares of the differences between measurements and outputs at the chosen measurement times. As a measure of the quality (which we want to optimize), of the estimates, we frequently use the standard errors of the estimates and/or other design criteria.

Specifically we want to address the following important questions in parameter estimation problems when considering models of type (1.1) – (1.2):

1. How do we choose the duration T of the experiment and the number n of measurements in the interval $[0, T]$, in order to obtain estimates of a high quality for θ_0?
2. Once the duration T and the number n of measurements are established, what is the optimal sampling distribution in the interval $[0, T]$ in order to obtain the most accurate estimates?

The first question pertains to the minimum time interval $[0, T]$ and to the minimal number of measurements taken in this time interval such that these measurements provide sufficient information on θ_0 in order to achieve an acceptable accuracy in the estimates.

The second question constitutes the main topic of the optimal experimental design literature, and the traditional way to answer it is to find sampling distributions for measurements which optimize some specific design criterion. This criterion is typically a given function of the Fisher information matrix. However, this matrix depends on the nominal parameter θ_0, which in practice is unknown. In order to overcome this difficulty one usually substitutes an initial guess θ^* for θ_0. The resulting designs are called locally optimal designs. Of the numerous design strategies found in the literature (for example, see [12, 19, 20, 31]) we mention the following three among the more popular ones:

a) The *D-optimal design*, which requires one to maximize the determinant $\det F(\hat{\theta})$ of the Fisher information matrix. It is largely accepted in the literature because of its appealing geometrical interpretation involving the asymptotic confidence regions for a maximum likelihood estimate $\hat{\theta}$ of θ which are ellipsoids. One can show that an optimal design with respect to this criterion yields approximately a minimal value for the volume of the error ellipsoid of the estimates;

b) The *c-optimal design*, for c a given vector in $\mathbb{R}^p$, which requires one to minimize the variance of the linear combination $g(\theta) = c^{\mathsf{T}}\theta$ of the parameters to be estimated. Up to a multiplicative constant, the asymptotic covariance of the maximum likelihood of g is given by $c^{\mathsf{T}}F^{-1}(\hat{\theta})c$, and an optimal design with respect to this criterion minimizes this variance. In the particular case when c is the i-th unit vector, i.e., $c = e_i$, the c-optimal design minimizes the variance of the least squares estimate for the i-th parameter θ_i;

c) The *E-optimal design*, which requires one to maximize the smallest eigenvalue $\lambda_{\min}\big(F(\hat{\theta})\big)$ of the Fisher information matrix. This is equivalent to minimizing the maximum eigenvalue $\lambda_{\max}\big(F^{-1}(\hat{\theta})\big)$ of the inverse F^{-1} (the covariance matrix), or to minimizing the worst variance among all estimates $c^{\mathsf{T}}\hat{\theta}$ for the linear combinations $c^{\mathsf{T}}\theta$ with a given norm $c^{\mathsf{T}}c = 1$.

Geometrically, the E-optimal design minimizes the maximum diameter of the asymptotic confidence ellipsoids for θ.

The practical importance of the issues addressed by the previous two questions is obvious for experiments where the cost of measurements (invasive procedures, expensive assays and technology, necessity of highly trained personnel, etc.) is high and where one wishes to avoid running the experiments longer than necessary and/or taking more measurements than needed. The latter point also can be important in cases where too many measurements may change the dynamics of the system (for instance, when measurements require one to take repeated blood or tissue samples).

In order to answer the questions raised above we introduce a general setting which involves idealized measurement procedures (see Subsection 2.1). This setting will allow one to prove the existence of idealized measurement procedures as solutions of the problems stated above. Although in general these idealized measurement procedures cannot be realized by concrete measurements, this approach, in addition to providing a sound basis for our investigations, reveals what may be achieved as the limit of practical measurement procedures. As we shall see, any idealized measurement procedure can be approximated arbitrarily closely by discrete (and therefore realizable) measurement procedures. Of course, the degree of approximation one can achieve in specific cases is in general limited by factors such as costs per measurement or influence of measurements on the system dynamics.

In Section 3 we introduce *traditional* and *generalized sensitivity functions* (GSF) and demonstrate that they can indicate regions of high information content on the parameters. Thus the corresponding parameter estimates may be improved if additional data points are sampled in these regions. (To our knowledge, GSFs were introduced first in this connection by Thomaseth and Cobelli, see [8, 11, 34]).

In Section 4 we present the basic results on existence of optimal sampling strategies and on approximation issues for the theoretical framework introduced in Subsection 2.1. Our approach involves spaces of probability distributions endowed with the Prohorov metric. In Section 5 we show that the Fisher information matrix, which is essential when we define the GSFs, plays a central role for our strategy to determine an appropriate final time T for duration of an experiment. This strategy is based on the time evolution of the sensitivity functions and that of the condition number of the Fisher information matrix.

In Section 6 we also illustrate the role of the sensitivity functions as tools in optimal design; more specifically, we exemplify their use in finding enhanced sampling distributions. Finally, in Section 6, we also present an optimal design criterion which is based on our abstract approach of obtaining almost optimal time sampling strategies by approximating an intuitively optimal idealized (but in general not realizable) strategy.

Throughout the paper we use the Verhulst-Pearl logistic population model as an example to illustrate our results. This example is widely used in the research literature, and has an analytical solution, few parameters, and well-known dynamics (see [8, 10, 27]). This logistic model, which approximates the size of a saturation limited population over time, is described by the differential equation

$$\dot{x}(t) = rx(t)\left(1 - \frac{x(t)}{K}\right), \quad x(0) = x_0, \tag{1.3}$$

where the constants K, r and x_0 represent the carrying capacity of the environment, the intrinsic growth rate in the population and the initial population size, respectively. Since one typically treats the initial condition also as a parameter, the parameter vector we wish to estimate is $\theta = (K, r, x_0)^\mathsf{T} \in \mathbb{R}^3$. The analytical solution of (1.3) is given by

$$x(t) = \frac{K}{1 + (K/x_0 - 1)\, e^{-rt}}, \tag{1.4}$$

and this approaches the steady state $x \equiv K$ as $t \to \infty$ (see Figure 1(a)) and will be taken as the output of the system, i.e., the right side of equation (1.4) defines the function $f(t, \theta)$ in this case.

2 A General Formulation of the Parameter Estimation Problem

2.1 Measurement Procedures

Given $T > 0$ we assume that we can take measurements for any $t \in [0, T]$. Following standard statistical theory, we consider

$$y(t) = f(t, \theta_0) + \varepsilon(t), \quad t \in [0, T], \tag{2.1}$$

as the realization model for the observation process, where θ_0 is the "true" or nominal parameter vector, which is unknown to us. The nonlinear regression function f is assumed to be continuous and twice differentiable with respect to the parameters θ. Later in our development of the theory (see Section 3.2 on generalized sensitivities) we shall need the assumption that the time behavior of the output η coincides with the time behavior of the corresponding quantity of the real system not only for the nominal parameter vector θ_0 but for all θ in a neighborhood of θ_0.

The function y which gives the measurements at any time $t \in [0, T]$ is a realization (sample path) of a stochastic process (the *statistical model*)

$$Y(t) = f(t, \theta_0) + \mathcal{E}(t),$$

which governs the measurement process. As usual $f(t, \theta_0)$ is the (deterministic) model output corresponding to the true parameter θ_0 and $\mathcal{E}$ is a noisy random process for measurement errors. The measurement error $\varepsilon(t)$ in (2.1) is simply a realization of $\mathcal{E}(t)$. Throughout the paper we assume that, for each $t \in [0, T]$, the random variable $\mathcal{E}(t)$ has zero mean, known, possibly time-dependent variance $\sigma^2(t)$ and that, for any $t \neq s$ in $[0, T]$, $\mathcal{E}(t)$ and $\mathcal{E}(s)$ are independent. That is, we have

$$\begin{aligned} E\big(\mathcal{E}(t)\big) &= 0, \quad t \in [0, T], \\ \operatorname{Var} \mathcal{E}(t) &= \sigma^2(t), \quad t \in [0, T], \\ \operatorname{Cov}\big(\mathcal{E}(t)\mathcal{E}(s)\big) &= \sigma(t)\sigma(s)\delta(t - s), \quad t, s \in [0, T], \end{aligned} \tag{2.2}$$

where δ is the Dirac distribution with support $\{0\}$. Let $m(t)$ denote the measurement density at time $t \in [0, T]$, i.e., $\int_s^t m(\tau)\, d\tau$ is the number of measurements in the interval $[s, t]$, $s < t$. In order to define the least squares error $J(y, \theta)$ for a given measurement procedure $y(\cdot)$, and for θ in a neighborhood of θ_0, we start with the weighted error for one measurement at time t, which is given by $\sigma(t)^{-2}\big(y(t) - f(t, \theta)\big)^2$, i.e., we put more weight on measurements where the variance of the measurement error is smaller.

Next we choose mesh points $t_i = i\Delta t$, $i = 0, \ldots, n$, with $n = T/\Delta t$. We assume that $\Delta t > 0$ is chosen sufficiently small so that the number of measurements in the time intervals $[t_i, t_i + \Delta t]$, $i = 0, \ldots, n - 1$, is given by $m(t_i^*)\Delta t$ and the average weighted error on $[t_i, t_i + \Delta t]$ by $\sigma(t_i^{**})^{-2}\big(y(t_i^{**}) - f(t_i^{**}, \theta)\big)^2$ for some $t_i^*, t_i^{**} \in [t_i, t_i + \Delta t]$. Here we assume for the moment that m and σ are continuous. Then the weighted errors in the intervals $[t_i, t_i + \Delta t]$ are given by

$$\frac{m(t_i^*)}{\sigma(t_i^{**})^2}\big(y(t_i^{**}) - f(t_i^{**}, \theta)\big)^2 \Delta t, \quad i = 0, \ldots, n - 1.$$

Taking $\Delta t \to 0$ we obtain

$$J(y,\theta) = \int_0^T \frac{m(t)}{\sigma(t)^2}\big(y(t) - f(t,\theta)\big)^2\, dt.$$

The integral above can be viewed as an integral with respect to an absolutely continuous measure P on $[0,T]$ with density $m(t) = dP(t)/dt$. This motivates us to consider, more generally, error functionals of the form

$$J(y,\theta) = \int_0^T \frac{1}{\sigma(t)^2}\big(y(t) - f(t,\theta)\big)^2\, dP(t), \tag{2.3}$$

where P is a general measure on $[0,T]$. Since we shall obtain the parameter estimate $\hat{\theta}$ by minimizing $J(y,\theta)$ for θ in a neighborhood of θ_0, we can, without restriction of generality, assume that P is a probability measure on $[0,T]$.

If, for points $t_1 < \cdots < t_n$ in $[0,T]$, we take

$$P_d = \sum_{i=1}^n \delta_{t_i}, \tag{2.4}$$

where δ_a denotes the Dirac delta distribution with support $\{a\}$, we obtain

$$J_d(y,\theta) = \sum_{i=1}^n \frac{1}{\sigma(t_i)^2}\big(y(t_i) - f(t_i,\theta)\big)^2, \tag{2.5}$$

which is the weighted least squares cost functional for the case where we take a finite number of measurements in $[0,T]$. Of course, the introduction of the measure P allows us to change the weights in (2.5) or the weighting function in (2.3). For instance, if P is absolutely continuous with density $m(\cdot)$ the error functional (2.3) is just the weighted L^2-norm of $y(\cdot) - f(\cdot,\theta)$ with weight $m(\cdot)/\sigma(\cdot)^2$.

2.2 Least Squares Estimates

In order to estimate θ_0, we use a weighted least squares procedure, i.e., the *parameter estimate* $\hat{\theta}$ is an admissible parameter vector which minimizes the cost functional (2.3) with given measurements $y(\cdot)$,

$$\hat{\theta} = \operatorname*{argmin}_{\theta} J(y,\theta), \tag{2.6}$$

Since the specific data y is a realization of the stochastic process Y, we consider $\hat{\theta}$ as a realization of a random variable $\hat{\Theta}$, which usually is called the *weighted least squares estimator* for the parameter estimation problem. Symbolically we rewrite (2.6) as

$$\hat{\Theta} = \operatorname*{argmin}_{\theta} J(Y,\theta).$$

An important question to address is how the uncertainty in the measured data propagates into the parameter estimates. This is equivalent to determining the statistical properties of the estimator $\hat{\Theta}$ based on those of Y. Using the linearization argument presented in detail in Appendix A.3, we obtain that to a first order approximation we have the distribution results

$$\hat{\Theta} \sim \mathcal{N}_p(\theta_0, \Sigma_0),$$

i.e., $\hat{\Theta}$ is approximately a normally distributed p-dimensional random vector with expected value θ_0 and covariance matrix Σ_0, where Σ_0 is the inverse of the Fisher information matrix

(A.23). In order to quantify the accuracy of specific estimates $\hat{\theta}$, and to compare between different sampling strategies, we use the classical approach in asymptotic statistical analysis, and compute the standard errors SE_k given by

$$SE_k = \sqrt{(\Sigma_0)_{kk}}, \quad k = 1, 2, ..., p. \tag{2.7}$$

The standard errors can be used to compute the confidence intervals, which give further insight about the quality of the parameter estimates obtained.

It is important to note that the matrix Σ_0 depends explicitly on the measure P (see (A.23) in Section A.3). Changing P affects Σ_0, which in turn controls the accuracy of the least square estimates (2.6). At this point the potential use of the measure P in (2.3) as an optimal design tool becomes apparent. For a given class of measures P on the interval $[0, T]$ we want to find the one which yields the most accurate parameter estimates.

3 Traditional and Generalized Sensitivity Functions

3.1 Traditional Sensitivity Functions

The traditional sensitivity functions are frequently used in simulation studies (i.e., when we investigate forward problems) where one wants to assess the degree of sensitivity of a model output with respect to various parameters on which it depends, and to identify the parameters to which the model is most/least sensitive. We consider the output model (1.2). To quantify the variation in the output variable $\eta(t)$ with respect to changes in the k-th component θ_k of the parameter vector θ we are led to consider the *first order sensitivity functions*, also called the *traditional sensitivity functions* (TSF), defined naturally in terms of the partial derivatives

$$s_k(t, \theta) = \frac{\partial \eta}{\partial \theta_k}(t, \theta) \in \mathbb{R}^M, \quad k = 1, \ldots, p, \tag{3.1}$$

which assumes smoothness of f with respect θ (compare [2, 15, 18, 21, 26], [27, pp. 7 – 9], [32]). If the output model (1.2) originates from the dynamical system (1.1) and we assume sufficient regularity of the functions $\mathcal{F}$, $x_0(\cdot)$, h, then the $M \times p$ matrix (throughout ∇_θ is a row vector)

$$s(t, \theta) = \big(s_1(t, \theta), \ldots, s_p(t, \theta)\big) = \begin{pmatrix} \nabla_\theta \eta_1(t, \theta) \\ \vdots \\ \nabla_\theta \eta_M(t, \theta) \end{pmatrix}$$

is given by

$$s(t, \theta) = \frac{\partial h}{\partial x}(t, x(t, \theta), \theta) \frac{\partial x}{\partial \theta}(t, \theta) + \frac{\partial h}{\partial \theta}(t, x(t, \theta), \theta), \quad 0 \le t \le T,$$

where $\partial h/\partial x = \big(\partial h_i/\partial x_j\big)_{i=1,\ldots,M,\ j=1,\ldots,N}$, $\partial x/\partial \theta = \big(\partial x_j/\partial \theta_k\big)_{j=1,\ldots,N,\ k=1,\ldots,p}$ and $\partial h/\partial \theta = \big(\partial h_i/\partial \theta_k\big)_{i=1,\ldots,M,\ k=1,\ldots,p}$.

It is well known that the matrix $X(t, \theta) := (\partial x/\partial \theta)(t, \theta) \in \mathbb{R}^{N \times p}$ as a function of t satisfies the linear ODE-system (called the *sensitivity equations*)

$$\begin{aligned} \dot{X}(t, \theta) &= \mathcal{F}_x(t, x(t, \theta), \theta) X(t, \theta) + \mathcal{F}_\theta(t, x(t, \theta), \theta), \\ X(0, \theta) &= \frac{\partial x_0}{\partial \theta}(\theta), \end{aligned} \tag{3.2}$$

where $x(t, \theta)$ is the solution of the initial value problem in (1.1). The matrices $\mathcal{F}_x \in \mathbb{R}^{N \times N}$ and $\mathcal{F}_\theta \in \mathbb{R}^{N \times p}$ are defined analogous to the matrices $\partial h/\partial x$, etc., above. Equation (3.2) is very

useful in practice, since it provides in conjunction with system (1.1) a fast and efficient way to compute the sensitivity matrix s numerically for general systems of the form (1.1).

The sensitivity functions (3.1) are in fact simply the derivatives of the components of the output η with respect to the parameters. As long as one is only interested in the sensitivity of the output solely with respect to individual parameters (as in this paper), it is sufficient to use the functions given by (3.1). However, if we want to compare the sensitivity of outputs with respect to different parameters, then the derivatives (3.1) can be misleading. Instead of the derivatives (3.1) one should use the *relative sensitivity functions* σ_k defined by

$$\sigma_k(t,\theta) = \lim_{\Delta\to 0} \frac{\big(\eta(t,\theta+\Delta e_k) - \eta(t,\theta)\big)/\eta(t,\theta)}{\Delta/\theta_k} = \frac{\theta_k}{\eta(t,\theta)}\,\frac{\partial\eta}{\partial\theta_k}(t,\theta) = \frac{\theta_k}{\eta(t,\theta)} s_k(t,\theta),$$

where e_k is the usual unit vector in the k^{th} direction. A further advantage of the functions σ_k is that they are dimensionless.

Although the first order sensitivity functions (3.1) are among the most commonly used tools in simulation studies, sometimes it is useful to consider also the *second order sensitivity functions* defined by

$$s_{k,m}(t,\theta) = \frac{\partial^2\eta}{\partial\theta_k\partial\theta_m}(t,\theta), \quad k,m = 1,\dots,p.$$

The function $s_{k,m}$ is a measure for the variation of the sensitivity s_m when the parameter θ_k changes or, vice versa, for the variation of s_k when θ_m changes.

It is important to note that the (first and second order) sensitivity functions depend on both t and θ, and it is this double dependence that makes them suitable as optimal design tools. The TSF characterize the rate of change of the output of a model (as a function of time) with respect to a parameter, but being derivatives with respect to the parameters this characterization is only *local with respect to the parameters.* For example, if the sensitivity s_k is close to zero on a time subinterval $[\tilde{t}-\delta,\tilde{t}+\delta]$ and for θ in a neighborhood of the true value θ_0, then the output given by (1.2) is *insensitive* to the parameter θ_k on that particular time subinterval and in a neighborhood of θ_0. The same function s_k can take large values on a different time subinterval or in a neighborhood of a different nominal parameter vector $\tilde{\theta}_0$, indicating that the output variable η is *very sensitive* to the parameter θ_k on the latter time subinterval in a neighborhood of $\tilde{\theta}_0$.

While not directly relevant to our analysis in this section, we derive briefly an expression involving a second order sensitivity that will be useful in subsequent discussions (see discussions of the sets $S^*_{[5,10]}$ in Subsection 6.2 below). To this end, we consider the Taylor expansion of the regression function $f(t,\theta)$ in (2.1) around the true value θ_0, i.e.,

$$f(t,\theta) = f(t,\theta_0) + \nabla_\theta f(t,\theta_0)(\theta-\theta_0) + \frac{1}{2}(\theta-\theta_0)^{\mathsf{T}}\nabla^2_{\theta\theta} f(t,\theta_0)(\theta-\theta_0) + \dots\ .$$

Recall that f is assumed to be sufficiently smooth in θ. In a small neighborhood of θ_0, the second order Taylor polynomial is a good approximation of f, so we can substitute it in (2.5) and minimize instead this approximate cost functional $\tilde{J}(y,\theta)$, given by

$$\tilde{J}(y,\theta) = \int_0^T \frac{1}{\sigma^2(t)}\Big(\varepsilon(t) - \nabla_\theta f(t,\theta_0)(\theta-\theta_0) - \frac{1}{2}(\theta-\theta_0)^{\mathsf{T}}\nabla^2_{\theta\theta} f(t,\theta_0)(\theta-\theta_0)\Big)^2\, dP(t).$$

In a neighborhood of θ_0, $J(y,\theta)\approx\tilde{J}(y,\theta)$ and the value $\tilde{\theta}$ which minimizes $\tilde{J}(y,\theta)$ is a reasonable approximation for the estimate $\hat{\theta}$ given by (2.6). The optimality condition at $\tilde{\theta}$ is simply $\nabla_\theta\tilde{J}(y,\tilde{\theta})=0$, or equivalently

$$2\int_0^T \frac{1}{\sigma^2(t)}\Big(\varepsilon(t) - \nabla_\theta f(t,\theta_0)(\tilde{\theta}-\theta_0) - \frac{1}{2}(\tilde{\theta}-\theta_0)^{\mathsf{T}}\nabla^2_{\theta\theta} f(t,\theta_0)(\tilde{\theta}-\theta_0)\Big) \times \Big(\nabla_\theta f(t,\theta_0) + (\tilde{\theta}-\theta_0)^{\mathsf{T}}\nabla^2_{\theta\theta} f(t,\theta_0)\Big)^{\mathsf{T}}\, dP(t) = 0_{p\times 1}.$$

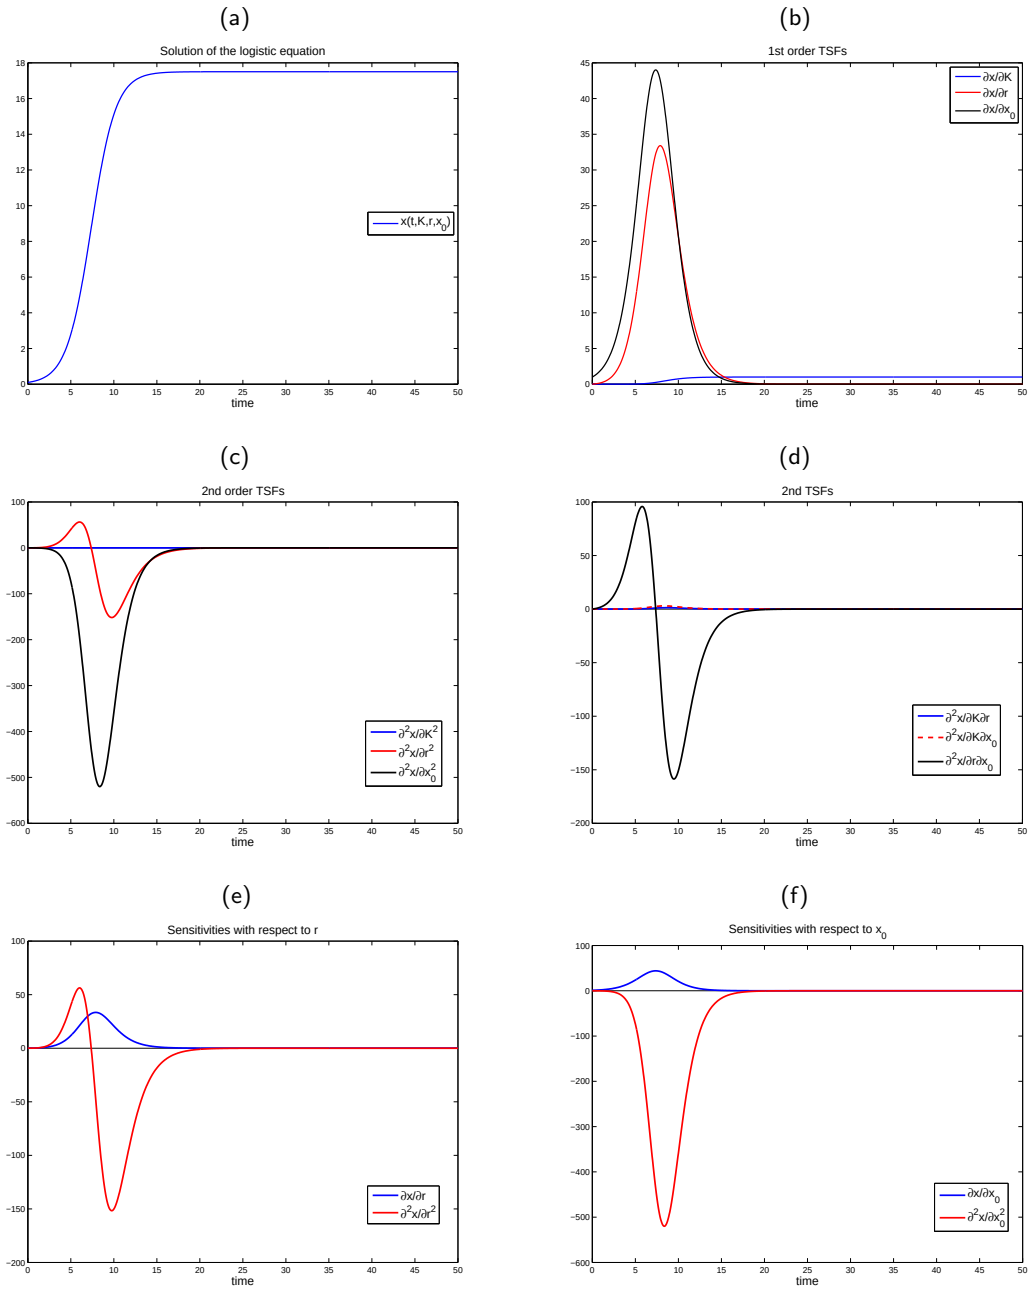

Figure 1: (a) Logistic curve; (b) TSF for the logistic model; (c) Second order TSF for the logistic model, diagonal entries; (d) Second order TSF for the logistic model, off-diagonal entries; (e) First and second order TSF with respect to r; (f) First and second order TSF with respect to x_0. All the figures above were plotted for the *true parameter vector* $\theta_0 = (17.5, 0.7, 0.1)$.

If we keep only the first order terms in $\tilde{\theta} - \theta_0$ in the equation above, we obtain

$$A(\tilde{\theta} - \theta_0) \approx \int_0^T \frac{\varepsilon(t)}{\sigma^2(t)} \nabla_\theta^\mathsf{T} f(t, \theta_0)\, dP(t), \tag{3.3}$$

where the matrix A is given by

$$A = \int_0^T \frac{1}{\sigma^2(t)} \Big(\nabla_\theta^{\mathsf{T}} f(t,\theta_0) \nabla_\theta f(t,\theta_0) - \varepsilon(t) \nabla^2_{\theta\theta} f(t,\theta_0) \Big)\, dP(t). \tag{3.4}$$

In Figure 1 we have plotted the logistic growth curve along with the first and second order traditional sensitivity functions for $\eta(t) = f(t,\theta) = x(t,\theta)$ from (1.4). The magnitudes of the second order sensitivity functions are much larger than those of the first order sensitivity functions. Moreover, their support is slightly wider, as can be seen in Figure 1, (e) and (f). These facts support our suggestions that the second order sensitivity functions may play an important role, and they should be used in conjunction with the first order sensitivity functions in parameter estimation problems.

We conclude this section with the intuitive comment that if we want to improve the accuracy of a parameter θ_k, we should take extra measurements at time points where the first order sensitivity functions (TSF) are large and the second order sensitivities are small (or where the relative magnitudes considerably favor the TSF). Based on our experience, we advocate the use of the first and second order sensitivity functions as optimal design tools.

3.2 Sensitivity of Estimates with Respect to Model Parameters and Generalized Sensitivity Functions

As we have already noted, the sensitivity functions introduced in the previous subsection are mainly used in simulation studies, where one wants to ascertain the change in model output with respect to changes in parameters. However, as we suggested they can be used also as optimal design tools. In this subsection we build mathematical tools which directly quantify the sensitivity of the parameter estimates $\hat\theta$ with respect to the model parameters θ. As with the traditional sensitivity functions, our approach will be motivated and carried out from an optimal design perspective. Our starting point will again be the optimality condition (2.3) for the least squares problem with general measurements procedures. Our discussions here are based directly on the ideas of Thomaseth and Cobelli as discussed in [11, 34].

Under reasonable assumptions, the cost functional J in (2.3) is differentiable with respect to θ. Therefore the etimate $\hat\theta$ corresponding to a given $y(\cdot,\theta_0)$ satisfies the optimality condition

$$\nabla_\theta J\big(y(\cdot,\theta_0), \hat\theta\big) = 0,$$

or equivalently

$$\int_0^T \frac{1}{\sigma^2(s)} \big(y(s,\theta_0) - f(s,\hat\theta)\big)\, \nabla_\theta f(s,\hat\theta)\, dP(s) = 0. \tag{3.5}$$

We consider nominal parameter vectors θ in a neighborhood $\mathcal{U}$ of θ_0 with corresponding estimates $\hat\theta = \hat\theta(\theta)$, $\theta \in \mathcal{U}$. Then condition (3.5) holds for all $\theta \in \mathcal{U}$, i.e.,

$$g(\theta) := \int_0^T \frac{1}{\sigma^2(s)} \big(y(s,\theta) - f(s,\hat\theta(\theta))\big) \nabla_\theta f(s,\hat\theta(\theta))\, dP(s) = 0, \quad \theta \in \mathcal{U}.$$

This implies that the derivative Dg of g with respect to θ is zero in $\mathcal{U}$ or

$$Dg(\theta) = 0, \quad \theta \in \mathcal{U}.$$

Since $\theta \mapsto g(\theta)$ defines a mapping $\mathcal{U} \to \mathbb{R}^{1,p} \cong \mathcal{L}(\mathbb{R}^p, \mathbb{R})$, we see that $Dg(\theta) \in \mathcal{L}\big(\mathbb{R}^p, \mathcal{L}(\mathbb{R}^p,\mathbb{R})\big) \cong \mathcal{L}(\mathbb{R}^p \times \mathbb{R}^p, \mathbb{R})$, i.e., $Dg(\theta)$ is a bilinear functional on $\mathbb{R}^p$ and thus can be represented by a $p \times p$ matrix (see, for instance, [17, Section VIII.12]). In order to obtain $Dg(\theta)$ we expand $g(\theta+\delta)$, $\delta \in \mathbb{R}^p$, as

$$g(\theta+\delta) = g(\theta) + Dg(\theta)\delta + \text{higher order terms},$$

and obtain after some straight forward computations

$$\begin{aligned} 0 = Dg(\theta) &= \nabla^2_{\theta\theta} J\big(y(\cdot,\theta),\hat\theta(\theta)\big)\frac{\partial\hat\theta}{\partial\theta}(\theta) + \nabla^2_{\theta_y} J\big(y(\cdot,\theta),\hat\theta(\theta)\big)\nabla_\theta y(\cdot,\theta) \\ &= \bigg(\int_0^T \frac{1}{\sigma^2(s)}\big(y(s,\theta) - f(s,\hat\theta(\theta))\big)\nabla^2_{\theta\theta} f(s,\hat\theta(\theta))\,dP(s) \\ &\qquad - \int_0^T \frac{1}{\sigma^2(s)}\nabla_\theta{}^{\mathsf T} f(s,\hat\theta(\theta))\nabla_\theta f(s,\hat\theta(\theta))\,dP(s)\bigg)\frac{\partial\hat\theta}{\partial\theta}(\theta) \\ &\quad + \int_0^T \frac{1}{\sigma^2(s)}\nabla_\theta{}^{\mathsf T} f(s,\hat\theta(\theta))\nabla_\theta y(s,\theta)\,dP(s), \quad \theta\in\mathcal{U}. \end{aligned} \tag{3.6}$$

From this equation we can obtain the sensitivity matrix $\partial\hat\theta/\partial\theta$ provided $\nabla^2_{\theta\theta} J\big(y(\cdot,\theta),\hat\theta(\theta)\big)$ is nonsingular. We are interested in how this sensitivity matrix changes if we progress with our measurement procedure in time. Thus we may intuitively view this methodology as one which attempts to measure how observations (data) contribute longitudinally to our ability to estimate $\hat\theta$ through the sensitivity of $\hat\theta$ with respect to a particular component θ_k. That is, we are describing a *sensitivity of the estimated parameters with respect to the observations or data* concept. In view of this goal we assume that, for some fixed $t\in[0,T]$, a variation of the nominal parameter vector θ only causes a change of the measurements taken during the time interval $[0,t]$, whereas measurements taken during $(t,T]$ remain fixed to their original values. Therefore we have (using also that $\varepsilon(t)$ is not dependent on θ)

$$\nabla_\theta y(s,\theta) = \begin{cases} \nabla_\theta f(s,\theta) & \text{for } 0\le s\le t, \\ 0 & \text{for } t< s\le T, \end{cases}$$

which implies

$$\nabla^2_{\theta_y} J\big(y(\cdot,\theta),\hat\theta(\theta)\big)\nabla_\theta y(\cdot,\theta) = \int_0^t \frac{1}{\sigma^2(s)}\nabla_\theta{}^{\mathsf T} f(s,\hat\theta(\theta))\nabla_\theta f(s,\theta)\,dP(s), \quad \theta\in\mathcal{U}.$$

Equation (3.6) relates the sensitivity of the estimates $\hat\theta$ with respect to θ, to the corresponding sensitivities of the output. The entries of the sensitivity matrix $\partial\hat\theta/\partial\theta$ are the quantities of interest. They depend on the model parameters θ, on $t\in[0,T]$ and on the final time T. Naturally, equation (3.6) is our starting point for investigating the behavior of sensitivities $\partial\hat\theta/\partial\theta = \big(\partial\hat\theta/\partial\theta\big)(t)$ as functions of t and θ. However, there is an important disadvantage associated with this equation, namely the fact that it is *realization-dependent*. The sensitivities $\partial\hat\theta/\partial\theta$ are given in terms of $\hat\theta(\theta)$ and $y(\cdot,\theta)$, which are realizations of the corresponding least squares estimator $\hat\Theta(\theta)$ and of the random variable $Y(\cdot,\theta)$ which governs the measurement process on the time interval $[0,T]$.

Our ultimate goal is to develop mathematical tools that indicate regions of high information content with respect to a parameter where the tools are *not realization-dependent*. Therefore, we next consider equation (3.6) formulated in terms of the random variables $\hat\Theta$ and Y, i.e.,

$$\Big(\nabla^2_{\theta\theta} J\big(Y(\cdot,\theta),\hat\Theta(\theta)\big)\Big)\,\frac{\partial\hat\Theta}{\partial\theta}(\theta) = -\nabla^2_{\theta_y} J\big(Y(\cdot,\theta),\hat\Theta(\theta)\big)\frac{\partial Y}{\partial\theta}(\cdot,\theta),$$

and instead of investigating the behavior of the sensitivities $\partial\hat\theta/\partial\theta$ for a particular realization, we will be interested in the behavior of the expected value

$$E\Big(\frac{\partial\hat\Theta}{\partial\theta}(\theta)(t)\Big), \quad 0\le t\le T,\ \theta\in\mathcal{U},$$

as a function of t and θ.

From asymptotic statistical theory, we know that the estimator $\hat{\Theta}$ is unbiased, i.e., $E\big(\hat{\Theta}(\theta)\big) = \theta$ in a neighborhood $\mathcal{U}$ of θ_0. Using this, we can argue that $E\big(\nabla_\theta f(t, \hat{\Theta}(\theta))\big) \approx \nabla_\theta f(t, \theta)$, and using additional technical assumptions in equation (3.6) (expected value commuting with the time integral, expected value of products approximately equal to the product of expected values), we obtain that the expected value of $\nabla^2_{\theta\theta} J$ is approximately given by

$$E\left(\nabla^2_{\theta\theta} J\big(Y(\cdot,\theta), \hat{\Theta}(\theta)\big)\right) \approx \int_0^T \frac{1}{\sigma^2(s)} \nabla_\theta^\intercal f(s,\theta) \nabla_\theta f(s,\theta)\, dP(s) := F(T,\theta), \tag{3.7}$$

which we assume to be nonsingular. We note that the $p \times p$ matrix $F(T,\theta)$ on the right side of (3.7), is simply the *Fisher information matrix* (FIM) for our problem with respect to general measurement procedure determined by P in the interval $[0,T]$. As one can see from its definition, the Fisher information matrix depends on the mathematical model f, the measure P and the final time T. An interesting question to address is if the evolution of the condition number of F as a function of T can be used to choose an appropriate stopping time for our experiments. We will carry out this analysis in Section 5.2.

Under similar conditions which led to equation (3.7) we obtain

$$E\left(\nabla^2_{\theta y} J\big(Y(\cdot,\theta), \hat{\Theta}(\theta)\big) \frac{\partial Y}{\partial \theta}(\cdot,\theta)\right) \approx \int_0^t \frac{1}{\sigma^2(s)} \nabla_\theta^\intercal f(s,\theta) \nabla_\theta f(s,\theta)\, dP(s) = F(t,\theta). \tag{3.8}$$

It is important to note that if $r = r(\theta)$ is a linear function then $E\big(r(\theta)\big) = r(E(\theta))$, which consequently holds only up to a first-order approximation when r is nonlinear; in general we only have $E(r(\theta)) \approx r(E(\theta))$.

Next we assume that $\big(\nabla^2_{\theta\theta} J(Y, \hat{\Theta})\big)^{-1}$ and $\nabla^2_{\theta y} J(Y, \hat{\Theta})\big(\partial Y/\partial\theta\big)$ are independent $p \times p$ random variables, which implies that

$$E\Big(\big(\nabla^2_{\theta\theta} J(Y,\hat{\Theta})\big)^{-1} \nabla^2_{\theta y} J(Y,\hat{\Theta}) \frac{\partial Y}{\partial \theta}\Big) = E\Big(\big(\nabla^2_{\theta\theta} J(Y,\hat{\Theta})\big)^{-1}\Big) E\Big(\nabla^2_{\theta y} J(Y,\hat{\Theta}) \frac{\partial Y}{\partial \theta}\Big).$$

If we now take expected values in equation (3.6), and use the equation above along with (3.7) and (3.8), we obtain

$$E\Big(\frac{\partial \hat{\Theta}}{\partial \theta}(\theta)\Big)(t) \approx F^{-1}(T,\theta) F(t,\theta) =: G(T,t,\theta), \quad 0 \le t \le T,\ \theta \in \mathcal{U}. \tag{3.9}$$

Equation (3.9) is *realization-independent* and gives the evolution of the expected value $E\big(\partial\hat{\Theta}/\partial\theta\big)$ as a function of t and θ when $t \in [0,T]$ and θ is in a neighborhood of θ_0. For $t = T$ we see that

$$E\Big(\frac{\partial \hat{\Theta}}{\partial \theta}(\theta)\Big)(T) = I_{p\times p},$$

which reflects the assumption that the estimator $\hat{\Theta}$ is unbiased (i.e., the variation in parameter estimates is equal to the variation in the true model parameters [34]).

We call the diagonal elements of the matrix $F^{-1}(T,\theta)F(t,\theta)$ the *generalized sensitivity functions* (GSF) with respect to the parameters θ,

$$\begin{aligned} gs(t,\theta) &= \operatorname{diag}\big(G(T,t,\theta)\big) \\ &= \operatorname{diag}\left(\int_0^t F(T,\theta)^{-1} \frac{1}{\sigma^2(s)} \nabla_\theta^\intercal f(s,\theta_0) \nabla_\theta f(s,\theta_0)\, dP(s)\right), \quad t \in [0,T]. \end{aligned} \tag{3.10}$$

If P is the discrete measure (2.4) then we obtain precisely the generalized sensitivity functions as introduced by Thomaseth and Cobelli in [34]. Of course, the vector $gs(t,\theta)$ can also be written as

$$gs(t,\theta) = \int_0^t \Big(F(T,\theta)^{-1} \frac{1}{\sigma^2(s)} \nabla_\theta^{\mathsf{T}} f(s,\theta_0) \Big) \bullet \nabla_\theta f(s,\theta_0)\, dP(s), \quad t \in [0,T],$$

where the symbol "$\bullet$" denotes the element-by-element multiplication of two vectors. From (3.9) we see that

$$gs(t,\theta) \approx \operatorname{diag}\Big(E\Big(\frac{\partial \hat{\Theta}}{\partial \theta}(T,\theta) \Big)(t) \Big).$$

Hence the GSF, approximate the diagonal elements of the expected value matrix of the sensitivities of the estimator $\hat{\Theta}$ with respect to θ in (3.9), regarded as time-dependent functions on $[0,T]$.

Like discrete counterparts introduced by Thomaseth and Cobelli, the generalized sensitivity functions (3.10) illustrate how the information content with respect to the parameters to be estimated is distributed throughout the experiment. By the nature of their definition, the GSF are cumulative functions, and due to this feature sometimes they may be misleading and exhibit false regions of high information content (the so called "forced-to-one" artifact [3, 8]; see also [25], where it is shown that if $G(T,\cdot,\theta)$ is close to a linear functions for t in some interval $[T_1,T]$, $T_1 < T$, then the parameter estimation problem with measurements taken only in $[T_1,T]$ is ill-posed). In order to avoid the potential misunderstanding caused by these artifacts, besides GSF, we also consider the first time derivative of the GSF, which in case that P is the Lebesgue measure on $[0,T]$ is given by

$$\frac{\partial}{\partial t} gs(t,\theta) = \Big(F(T,\theta)^{-1} \frac{1}{\sigma^2(t)} \nabla_\theta^{\mathsf{T}} f(t,\theta) \Big) \bullet \nabla_\theta f(t,\theta),$$

which is related to the *incremental generalized sensitivities* introduced in [34].

Remark 3.1. a) If $P = P_d$ as given by (2.4) then we obtain

$$F_d(t,\theta) = \begin{cases} 0 & \text{for } 0 \le t < t_1, \\ \displaystyle\sum_{i=1}^{k} \frac{1}{\sigma^2(t_i)} \nabla_\theta{}^{\mathsf{T}} f(t_i,\theta) \nabla_\theta f(t_i,\theta) & \text{for } t_k \le t < t_{k+1},\ k = 1,\dots,n-1, \\ \displaystyle\sum_{i=1}^{n} \frac{1}{\sigma^2(t_i)} \nabla_\theta{}^{\mathsf{T}} f(t_i,\theta) \nabla_\theta f(t_i,\theta) & \text{for } t_n \le t \le T. \end{cases} \tag{3.11}$$

In this case the generalized sensitivity functions $gs(t,\theta)$ are piecewise constant functions with $gs(0,\theta) = (0,\dots,0)$, $gs(T,\theta) = (1,\dots,1)$ and jumps of size

$$\operatorname{diag}\Big(\Big(\sum_{i=1}^{n} \frac{1}{\sigma^2(t_i)} \nabla_\theta{}^{\mathsf{T}} f(t_i,\theta) \nabla_\theta f(t_i,\theta) \Big)^{-1} \nabla_\theta{}^{\mathsf{T}} f(t_i,\theta) \nabla_\theta f(t_i,\theta) \Big)$$

at t_i, $i = 1,\dots,n$.

b) Let P be the Lebesgue measure and let F be the Fisher information matrix in this case (as given by (3.8)). Assume that $0 \le t_1^{(n)} < \cdots < t_{N(n)}^{(n)} = T$, $n = 1,2,\dots$, is a sequence of meshes in $[0,T]$ with $\max_{i=1,\dots,N(n)} \big(t_i^{(n)} - t_{i-1}^{(n)} \big) \to 0$ as $n \to \infty$ (here we set $t_0^{(n)} = 0$ for all n). Let F_n denote the Fisher information matrix corresponding to the discrete measure $P_n = \sum_{i=1}^{N(n)} (t_i^{(n)} - t_{i-1}^{(n)}) \delta_{t_i^{(n)}}$. Then it is easy to see that

$$\lim_{n\to\infty} F_n(t,\theta) = F(t,\theta), \quad 0 \le t \le T,\ \theta \in \mathcal{U}.$$

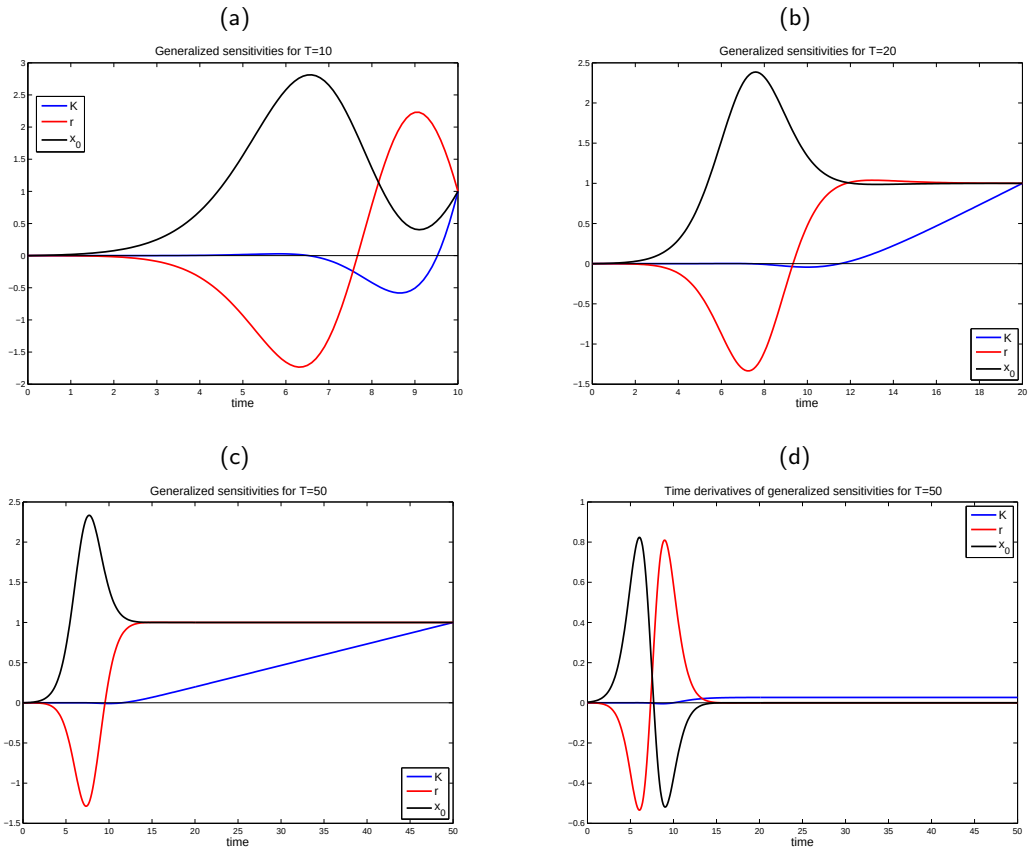

Figure 2: (a),(b) and (c): GSFs of the logistic model for $T = 10$, $T = 20$ and $T = 50$, respectively, with P being the Lebesgue measure; (d): time derivatives of the GSFs for $T = 50$. The *nominal parameter vector* is $\theta_0 = (17.5, 0.7, 0.1)$.

If the meshes are uniform with mesh size $t_i^{(n)} - t_{i-1}^{(n)} = T/N(n)$, $i = 1, \ldots, N(n)$, then

$$F_n(t, \theta) = \frac{T}{N(n)} F_d(t, \theta), \quad 0 \le t \le T, \ \theta \in \mathcal{U}, \ n = 1, 2, \ldots,$$

where F_d is given by (3.11) for the uniform mesh $\left(t_i^{(n)}\right)_{i=1,\ldots,N(n)}$.

For the logistic model (1.3) we plotted in Figure 2(a) – (c), the generalized sensitivity functions given by (3.10) with Lebesgue measure P for various values of the final time T. For the first value $T = 10$ (Figure 2(a)), we note that the GSF curves exhibit "unsettled" shapes, whereas for $T = 20$ and $T = 50$ (Figure 2(b) and (c)) the shapes stabilize, and common features like minimum, maximum, monotonicity no longer depend on T. We attribute the unsettled behavior of the GSF curves for $T = 10$ to the fact that the GSF are *cumulative functions*, and that the information contained in measurements taken in the interval $[0, 10]$ on all three parameters can be considerably improved by taking measurements beyond $t = 10$. On the other hand, starting with approximately $T = 20$, the information content relevant to the estimation of our parameters is well contained within $[0, T]$, a fact supported visually by the stabilization of the GSF curves as function of T, after we reach $T = 20$. As we will see in the next section, the time T where the GSF curves reach their steady shapes, provides a stopping time criterion for the experiments.

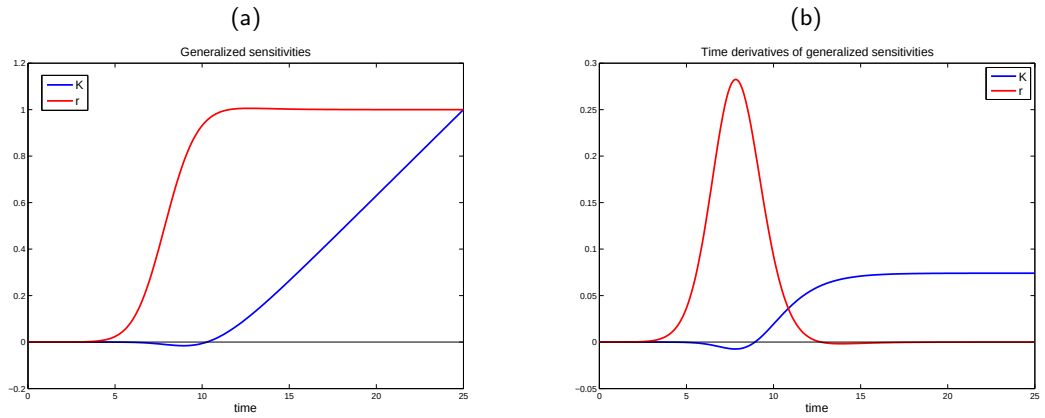

Figure 3: GSFs if only the parameters $\tilde{\theta} = (r, K)$ are considered (panel (a)) and the corresponding time derivatives (panel (b)). The nominal parameters are $\tilde{\theta}_0 = (17.5, 0.7)$

The high correlation between parameters r and x_0 can be clearly observed in Figure 2(c) and (d), and it is the main cause of the so called "forced-to-one" artifact, thoroughly discussed in [3] and [8]). In order to investigate how the shapes of the GSF curves change in the absence of high correlation, we hold x_0 constant at its true value 0.1 and consider the new parameter vector $\tilde{\theta} = (K, r)$. The corresponding set of generalized sensitivity functions and their time derivatives for the true values $(17.5, 0.7)$ are plotted in Figure 3(a) and (b).

We observe immediately that the GSF with respect to r is no longer decreasing; rather, gs_r is now increasing even on the time subinterval where it was decreasing when x_0 was included in the parametrization. The same dramatic change can also be seen when comparing the time derivatives of the GSF curve with respect to r which in the new case no longer has a negative part (see Figures 2(d), and 3(b)). On the other hand, we note that the GSF with respect to K (and its corresponding time derivative) remain very similar to the original shape. We explain this based on the fact K had a very low correlation with x_0 (and also wth r), and therefore is not affected very much by the removal of this parameter.

In conclusion, we note that the decreasing regions of the GSF are directly related to the high correlation between various parameters. These regions are still highly informative when attempting to recover parameters or improve standard errors, and should be included when taking additional data points.

As we shall see in the following, the regions of steepest increase and steepest decrease on the graphs of the generalized sensitivity curves correspond to data of high information content with respect to parameter estimation. Intuitively, if one samples additional data points in these regions, one typically obtains more accurate estimates as illustrated by the decrease in magnitude of the corresponding standard errors.

4 Theoretical Framework

The introduction of the measure P in Section 2 allows for a unified framework for optimal design criteria which incorporates all the popular design criteria mentioned in the introduction. The Fisher information matrix $F(T, \theta)$ introduced in Subsection 3.2 (see (3.7)) depends on the measure P. For this section we indicate this dependence by writing $F = F_P$. We also remind the reader that we can restrict ourselves to probability measures on $[0, T]$. Let $\mathcal{P}(0, T)$ denote the set of all probability measures on $[0, T]$ and assume that $\mathcal{J} : \mathbb{R}^{p \times p} \to \mathbb{R}^+$ is given. The *optimal design problem* associated with $\mathcal{J}$ is the problem of finding a probability measure $\hat{P} \in \mathcal{P}(0, T)$

such that

$$\mathcal{J}\big(F_{\hat{P}}(T,\theta_0)\big) = \min_{P\in\mathcal{P}(0,T)} \mathcal{J}\big(F_P(T,\theta_0)\big). \tag{4.1}$$

This formulation incorporates all strategies for optimal design which try to optimize a functional depending continuously on the elements of the Fisher information matrix. In case of the design criteria mentioned in the introduction, $\mathcal{J}$ is the determinant, the smallest eigenvalue, or a quadratic form, respectively, in terms of the inverse of the Fisher information matrix.

4.1 Existence of optimal sampling strategies

In order to solve problem (4.1), one needs a theoretical and computational framework that ensures the existence of a minimizer and provides tractable numerical algorithms to compute it. Unfortunately, the set $\mathcal{P}[0,T]$ does not have a linear space structure (a linear combination of two probability distributions is not a probability distribution) and we cannot take advantage of the plethora of optimization results available in Banach or Hilbert spaces. Hence, we are forced to solve the optimization problem (4.1) in a metric space setting. Fortunately, probability theory offers a topology which provides the theoretical framework we may utilize.

With the metric $d(x,y) = |x-y|$, $x, y \in [0,T]$, the interval $[0,T]$ is a complete, compact and separable metric space. On the set $\mathcal{P}(0,T)$ of Borel probability measures on $[0,T]$ (i.e., probability measures on the Borel subsets of $[0,T]$) we introduce the Prohorov metric ρ defined by (see [13, 23, 30])

$$\rho(P,\tilde{P}) = \inf_{A\subset[0,T] \text{ closed}} \big\{\epsilon \mid P(A) \le \tilde{P}(A^\epsilon) + \epsilon\big\}, \quad P,\tilde{P}\in\mathcal{P}(0,T),$$

where $A^\epsilon = \{y \in [0,T] \mid d(y,A) \le \epsilon\}$. In the following proposition we state properties of the metric space $\big(\mathcal{P}(0,T),\rho\big)$ which are inherited from the metric space $([0,T],d)$ and a result on ρ-convergence (see [5, 6]):

Proposition 4.1. a) *The metric space $\big(\mathcal{P}(0,T),\rho\big)$ is complete, compact and separable.*
b) *ρ-convergence is equivalent to weak*-convergence on $\mathcal{P}(0,T)$ (considering $\mathcal{P}(0,T)$ as a subset of $C_b([0,T])^*$).*

A consequence of statement b) of this proposition is (see, for instance, [13, 23, 29])

Lemma 4.2. *Let $P, P_n \in \mathcal{P}(0,T)$, $n = 1,2,\ldots$, be given. If $\lim_{n\to\infty}\rho(P_n,P) = 0$ then*

$$\lim_{n\to\infty}\int_0^T g(t)dP_n(t) = \int_0^T g(t)dP(t) \quad \textit{for all continuous functions } g \textit{ on } [0,T].$$

From the results stated above we obtain immediately the following theorem which is basic to our approach:

Theorem 4.3. *Assume that the functional $\mathcal{J}:\mathbb{R}^{p\times p}\to\mathbb{R}^+$ is continuous. Then problem (4.1) has a solution $\hat{P}\in\mathcal{P}(0,T)$.*

Proof. The assumption on $\mathcal{J}$ and the lemma imply that the mapping $P \to \mathcal{J}\big(F_P(T,\theta_0)\big)$ is continuous on $\mathcal{P}(0,T)$. Then the result follows from compactness of $\mathcal{P}(0,T)$. □

4.2 Approximation Issues

Having argued the existence of a solution problem (4.1), we next turn our attention to the computational aspects related to finding $\hat{P}$. We note that $(\mathcal{P}(0,T),\rho)$ is an infinite dimensional metric space, so we need to construct finite dimensional approximations $\hat{P}_n$ for $\hat{P}$ which to converge to $\hat{P}$ in the Prohorov metric when $n\to\infty$. The following density result (formulated

and proven for a general set Q in [4]) provides us the needed structure to build such finite dimensional approximations. We state here the result for $Q = [0,T]$.

Theorem 4.4. *Let $Q_0 = \{t_j\}_{j=1}^{\infty}$ be a countable, dense subset of $[0,T]$ and δ_{t_j} be the Dirac measure with atom at t_j. Then the set*

$$\mathcal{P}_0(0,T) := \Big\{ P \in \mathcal{P}(0,T) \Big| P = \sum_{j=1}^{k} p_j \delta_{t_j},\ k \in \mathbb{N}^+,\ t_j \in Q_0,\ p_j \geq 0,\ p_j \text{ rational},\ \sum_{j=1}^{k} p_j = 1 \Big\}$$

(i.e., the set of $P \in \mathcal{P}(0,T)$ with finite support in Q_0 and rational masses) is dense in $\big(\mathcal{P}(0,T), \rho\big)$ in the Prohorov metric ρ.

Let Q_d be a dense subset of $[0,T]$. The theorem above simply says that any probability distribution in $\mathcal{P}(0,T)$ can be approximated arbitrarily close in the Prohorov metric by finite convex combinations of Dirac measures supported in Q_d. Given $Q_d = \bigcup_{M=1}^{\infty} Q_M$ with $Q_M = \{t_j^M\}_{j=1,\dots,M}$ chosen such that Q_d is dense in $[0,T]$, we define, for $M = 1, 2, \dots,$

$$\mathcal{P}^M(0,T) := \Big\{ P \in \mathcal{P}(0,T) \Big| P = \sum_{j=1}^{M} p_j \delta_{t_j^M},\ t_j^M \in Q_M,\ p_j \geq 0,\ p_j \text{ rational},\ \sum_{j=1}^{M} p_j = 1 \Big\},$$

which is a compact subset of $\big(\mathcal{P}(0,T), \rho\big)$. Then we have that $\mathcal{P}^M(0,T) \to \mathcal{P}(0,T)$ in the ρ-topology; that is, elements in the infinite dimensional space $\mathcal{P}(0,T)$ can be approximated arbitrarily closely by elements from the finite dimensional space $\mathcal{P}^M(0,T)$ for M sufficiently large. In particular the minimizer $\hat{P}$ of $\mathcal{J}\big(F_P(T,\theta_0)\big)$ over $\mathcal{P}(0,T)$ can be approximated by a sequence of minimizers of $\mathcal{J}\big(F_P(T,\theta_0)\big)$ over the finite dimensional spaces $\mathcal{P}^M(0,T)$ as summarized by the following theorem:

Theorem 4.5. *Let $\mathcal{P}^M(0,T)$, $M = 1, 2, \dots,$ be defined as above. Suppose $\hat{P}_M$ is minimizer for $\mathcal{J}\big(F_P(T,\theta_0)\big)$ over $\mathcal{P}^M(0,T)$, i.e.,*

$$\hat{P}_M = \operatorname*{argmin}_{P \in \mathcal{P}^M(0,T)} \mathcal{J}\big(F_P(T,\theta_0)\big), \quad M = 1, 2, \dots\ .$$

Then we have $\rho(\hat{P}_M, \hat{P}) \to 0$ as $M \to \infty$.

The following corollary is an immediate consequence of this theorem:

Corollary 4.6. *Assume that $\mathcal{J} : \mathbb{R}^{p \times p} \to \mathbb{R}^+$ is continuous and let*

$$\hat{P} = \operatorname*{argmin}_{P \in \mathcal{P}(0,T)} \mathcal{J}\big(F_P(T,\theta_0)\big).$$

For $n = 1, 2, \dots$ we set $\Delta t = T/n$, $t_i^{(n)} = i\Delta t$, $i = 0, \dots, n$, and define $\mathcal{P}_n(0,T) \subset \mathcal{P}(0,T)$ by

$$\mathcal{P}_n(0,T) = \Big\{ P_n = \sum_{j=0}^{n} \mu_j^{(n)} \delta_{t_j^{(n)}},\ \mu_j^{(n)} \text{ rational},\ \mu_j^{(n)} \geq 0 \text{ and } \sum_{j=0}^{n} \mu_j^{(n)} = 1 \Big\}, \quad n = 1, 2, \dots,$$

where $\delta_{t_j^{(n)}}$ denotes the Dirac measure supported at $t_j^{(n)}$. Furthermore, let

$$\hat{P}_n = \operatorname*{argmin}_{P_n \in \mathcal{P}_n(0,T)} \mathcal{J}\big(F_{P_n}(T;\theta_0)\big).$$

Then we have

$$\lim_{n \to \infty} \rho(\hat{P}_n, \hat{P}) = 0, \tag{4.2}$$

$$\lim_{n \to \infty} \mathcal{J}\big(F_{\hat{P}_n}(T,\theta_0)\big) = \mathcal{J}\big(F_{\hat{P}}(T,\theta_0)\big). \tag{4.3}$$

Proof. Existence of $\hat{P}$ and $\hat{P}_n$, $n = 1, 2, \ldots$, follows from Theorem 4.3. It is easy to see that $Q = \bigcup_{n=1}^{\infty} Q_n$ with $Q_n = \{t_j^{(n)} | j = 0, \ldots, n\}$ is dense in $[0, T]$. Then (4.2) follows immediately from Theorem 4.5. Continuity of $P \to \mathcal{J}\big(F_P(T, \theta_0)\big)$ (compare Lemma 4.2) and (4.2) imply (4.3). □

5 Choosing a Final Time T

5.1 Eigenvalues of the Fisher Information Matrix

As we have seen before, the Fisher information matrix relates the variation in the parameter estimates to variations in the measurements. Also, its inverse F^{-1} is used to define the generalized sensitivity functions (3.10). Due to these facts, the condition number of the FIM is expected to play an important role in the statistical properties of $\hat{\theta}$, and numerical stability issues suggest that it is important to choose T such that the Fisher information matrix is well conditioned. In this subsection we present a theoretical result which shows that the minimum and maximum eigenvalues of the Fisher information matrix, as well as its spectral norm are monotonically increasing functions of the final time T. As we shall see in the next subsection, in the particular case where the condition number of the FIM is decreasing, which is true for the logistic model, we can use the condition number of the FIM to indicate an appropriate stopping time T. In order to simplify notation we do not indicate the dependence of the FIM and of quantities associated with the FIM on θ in this subsection.

Theorem 5.1. *Assume that* $\det F(T) \neq 0$. *Then* $\det F(T + \Delta) \neq 0$ *for all* $\Delta \geq 0$ *and*

i) *the minimum and the maximum eigenvalues of the Fisher information matrix are increasing functions of the final time T, i.e.,*

$$\lambda_{\min}(T) \leq \lambda_{\min}(T + \Delta) \quad \textit{and} \quad \lambda_{\max}(T) \leq \lambda_{\max}(T + \Delta), \quad \Delta \geq 0,$$

ii) *the spectral norm of the covariance matrix $F(T)^{-1}$ is a non-increasing function of the final time T, i.e.,*

$$\big\|F(T + \Delta)^{-1}\big\|_2 \leq \big\|F(T)^{-1}\big\|_2, \quad \Delta \geq 0$$

(*here* $\|\cdot\|_2$ *denotes the spectral norm of a matrix*).

Proof. i) The matrix $F(T)$ is positive definite. If $\lambda_{\min}(T)$ and $\lambda_{\max}(T)$ denote the smallest and the largest eigenvalue of $F(T)$, then we have

$$\lambda_{\min}(T)\|a\|^2 \leq a^{\mathsf{T}} F(T) a \leq \lambda_{\max}(T)\|a\|^2, \quad a \in \mathbb{R}^p.$$

For $a \in \mathbb{R}^p$ and $\Delta > 0$ we have

$$\begin{aligned}\lambda_{\min}(T)\|a\|^2 + \int_T^{T+\Delta} \frac{1}{\sigma^2(t)} \big\|\nabla_\theta f(t, \theta_0) a\big\|^2 dP(t) &\leq a^{\mathsf{T}} F(T + \Delta) a \\ &\leq \lambda_{\max}(T)\|a\|^2 + \int_T^{T+\Delta} \frac{1}{\sigma^2(t)} \big\|\nabla_\theta f(t, \theta_0) a\big\|^2 dP(t).\end{aligned}$$

The left side of this inequality implies

$$\lambda_{\min}(T) \leq \lambda_{\min}(T + \Delta). \tag{5.1}$$

If we choose a as an eigenvector of $F(T)$ corresponding to $\lambda_{\max}(T)$ then we obtain

$$a^{\mathsf{T}} F(T + \Delta) a = \lambda_{\max}(T)\|a\|^2 + \int_T^{T+\Delta} \frac{1}{\sigma^2(t)} \big\|\nabla_\theta f(t, \theta_0) a\big\|^2 dP(t) \leq \lambda_{\max}(T + \Delta)\|a\|^2, \tag{5.2}$$

and from the right side of this inequality we obtain

$$\lambda_{\max}(T) \leq \lambda_{\max}(T+\Delta).$$

ii) The largest and smallest eigenvalues of $F^{-1}(T+\Delta)$ are respectively $1/\lambda_{\min}(T+\Delta)$ and $1/\lambda_{\max}(T+\Delta)$. We note that $\|F(T+\Delta)^{-1}\|_2 = 1/\lambda_{\min}(T+\Delta)$, since $F(T+\Delta)^{-1}$ is symmetric and the spectral radius of $\left(F^{-1}(T+\Delta)\right)^{\mathsf{T}} F^{-1}(T+\Delta)$ is $\lambda^{-2}_{\min}(T+\Delta)$. Then (5.1) implies that $\det F(T+\Delta) \neq 0$, and ii) follows immediately from (5.1). □

Remark 5.2. The statements at the beginning of this subsection indicate that it would be desirable to show that the condition number $\rho_F(T) := \lambda_{\max}(T)/\lambda_{\min}(T)$ of the Fisher information matrix is a non increasing function of the final time T, i.e.,

$$\rho_F(T) \geq \rho_F(T+\Delta). \tag{5.3}$$

Fortunately, this is true for the logistic model, where it is the consequence of the fact that the minimum eigenvalue increases at a *significantly faster rate* than the maximum eigenvalue (see Table 1 and Figure 4, (a)). We also note that although both the minimum and the maximum eigenvalues increase, the rates at which they increase are directly related to the first order traditional sensitivity functions. This can be readily seen for example from (5.2) which provides some measure of the gap between $\lambda_{\max}(T)$ and $\lambda_{\max}(T+\Delta)$.

T	$\lambda_{\min}(T)$	$\lambda_{\max}(T)$	$\rho_F(T)$
15	2.6112	11403.3913	4367.0612
20	7.3132	11404.9422	1559.5081
25	12.2328	11404.9461	932.3246
30	17.1568	11404.9467	664.7493

Table 1: Values for $\lambda_{\min}(T)$, $\lambda_{\max}(T)$ and the condition number of the Fisher information matrix corresponding to the logistic model, as the final time T increases. The nominal parameter vector is $\theta_0 = (17.5, 0.7, 0.1)$.

However, although desirable, the relation (5.3) is not true in general. To see this we consider the initial value problem

$$\dot{x}(t) = -\mu_1 e^{-t} + \frac{\mu_2}{1+t}, \quad x(0) = \mu_1, \tag{5.4}$$

for positive t, whose solution is

$$x(t, \mu_1, \mu_2) = \mu_1 e^{-t} + \mu_2 \ln(1+t), \quad t \geq 0.$$

We take $f(t, \mu_1, \mu_2) = x(t, \mu_1, \mu_2)$ and observe that the Fisher information matrix does not depend on the parameters in this case, because f is a linear function of the parameters. Of course, the eigenvalues of the Fisher information matrix are increasing as required by Theorem 5.1 (see also Table 2), however the largest eigenvalue increases much faster than the smallest one, so that the condition number is also increasing. The measure P was taken to be the Lebesgue measure. This example shows, that in general the condition number of the Fisher information matrix is not a decreasing function of the final time T.

5.2 Tools for Choosing a Final Time T

One of the most important questions arising in practice when dealing with parameter estimation problems for dynamical systems of type (1.1) refers to choosing an appropriate final time T where

T	$\lambda_{\min}(T)$	$\lambda_{\max}(T)$	$\rho_F(T)$
10	0.4882	30.5071	62.4949
15	0.4944	64.2787	130.0073
20	0.4967	106.2849	215.0086
25	0.4977	156.5763	314.5862
30	0.4983	212.6548	426.7402

Table 2: Values for $\lambda_{\min}(T)$, $\lambda_{\max}(T)$ and the condition number of the Fisher information matrix corresponding to the model (5.4) for increasing values of the final time T.

to stop taking measurements. This is extremely important for applications where the costs of running the experiments per unit time may be high. For the choice of T we impose the following two criteria which are of heuristical nature:

1. Taking intervals $[0, T_1]$ with $T_1 < T$ when sampling data gives considerably less accurate estimates compared to the estimates we obtain when sampling data from $[0, T]$.
2. The improvement in the parameter estimates is negligible when increasing the sampling time beyond T (i.e., data outside $[0, T]$ is essentially irrelevant for the estimation of θ).

In practice, the experimentalists commonly use prior knowledge to choose T. In numerical simulations, T is taken sufficiently large so that the solution of (1.1) is already close to the steady state. Our approach here is to use a combination of mathematical tools which will help us determine a near optimal final time. To be more specific, we illustrate in the following how the evolution (as functions of final time T) of the traditional and generalized sensitivity functions, as well as the evolution of the condition number of the Fisher information matrix, and of the correlation coefficients can assist us in answering this important question.

5.2.1 Use of Sensitivity Functions

The logistic curve plotted in Figure 1(a), for the true parameter vector $\theta_0 = (17.5, 0.7, 0.1)$ nears its steady state at about $t = 20$ (or slightly earlier). If we look at the traditional sensitivity functions (Figure 1(b)), we observe that the TSFs with respect to r and x_0 are basically compactly supported in the interval $[0, 20]$ and essentially vanish afterwards. Also, the TSF with respect to K increases from zero and reaches one (the steady state) before $t = 20$. The behavior of the GSF curves with respect to the final time T leads to the same conclusion (Figure 2(a), (b) and (c)). We recall that, as a consequence of being cumulative functions, the shapes of the GSF curves stabilize for final time values around $T = 20$ and not before that. The time derivatives of the GSF curves (Figure 2(d)) support the same conclusion. So for the logistic model with nominal parameter vector $\theta_0 = (17.5, 0.7, 0.1)$, the interval $[0, 20]$ appears to be the optimal interval from which to sample data. It would not appear wise (at least intuitively) to sample data from a smaller interval, since that would eliminate data points where the TSF and the GSF curves exhibit a transient behavior, and therefore the model is still sensitive with respect to the parameters we wish to estimate.

A simple way to implement mathematically the previous visual analysis is to look at the TSF curves individually and to locate on each of them the time point after which there is no transitional behavior. In other words, we need to find T_K, T_r and T_{x_0} given by (in order to simplify notation we do not indicate the dependence of the functions on the nominal parameter)

$$\begin{aligned} T_K &= \max\big\{t : \ |s_K(u) - s_K(v)| > \varepsilon \text{ for some } u, v \in [t, t+h]\big\}, \\ T_r &= \max\big\{t : \ |s_r(u) - s_r(v)| > \varepsilon \text{ for some } u, v \in [t, t+h]\big\}, \\ T_{x_0} &= \max\big\{t : \ |s_{x_0}(u) - s_{x_0}(v)| > \varepsilon \text{ for some } u, v \in [t, t+h]\big\}, \end{aligned} \tag{5.5}$$

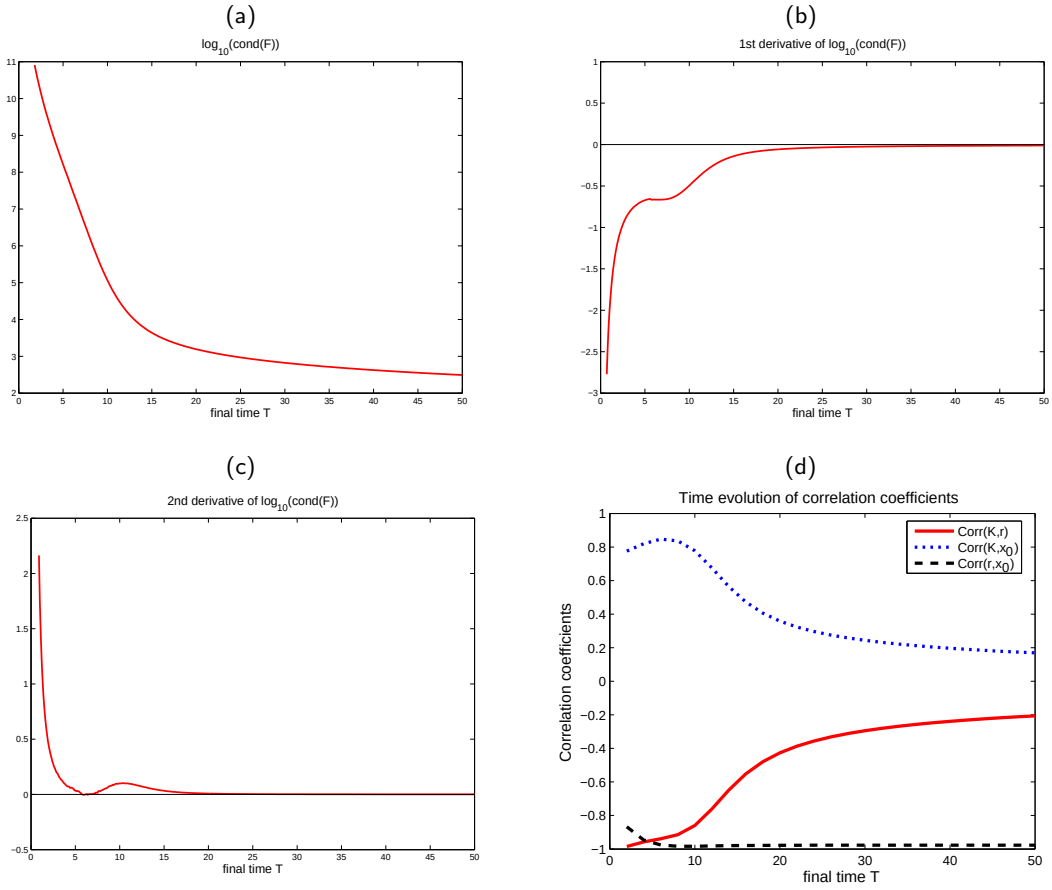

Figure 4: $\log_{10}$ of the condition number of the FIM (panel (a)); first and second order derivatives of $\log_{10}$ of the condition number of the FIM (panels (b) and (c)); correlation coefficients (panel (d)). All figures are plotted for the *nominal parameter vector* $\theta_0 = (17.5, 0.7, 0.1)$.

where s_K, s_r and s_{x_0} are the traditional sensitivity functions (3.1) with respect to the corresponding parameters, ε is a small tolerance, and $h > 0$ is the length of the time window used. The final time T we seek, when all three TSF no longer have a significant amount of change, is simply the maximum of the previous times, i.e.,

$$T = \max\big(T_K, T_r, T_{x_0}\big).$$

5.2.2 Condition Number of the Fisher Information Matrix

In Figure 4(a) – (c), we plotted the decimal logarithm of the condition number of the FIM, along with its first and second time derivatives as functions of the final time T. The Fisher information matrix considered here is the one defined in (3.7) with P being the Lebesgue measure. As one can see, the condition number of the FIM is huge for small values of T, and decreases almost exponentially on the interval $[0, 15]$. It decreases continually after that, but at a much slower rate. To have a better idea about its rate of change and its curvature, we look at the first and second derivatives. By visually investigating these three curves, and looking for a common point after which there is no significant change, we conclude again that an appropriate stopping time

is approximately $T = 20$.

We note that the behavior of the condition number is in agreement with our theoretical results (see Theorem 5.1) and also with our insight about the logistic model. It is perfectly reasonable for the condition number of the FIM to decrease as T gets larger. Intuitively, if T is larger, we have more information about the model output, so we have a better chance to estimate the model parameters more accurately. A small condition number for the FIM is desirable, since it yields a "smaller" dispersion matrix (the inverse of the FIM) and potentially smaller standard errors for our estimates.

The evolution of the condition number of the FIM is interesting but not surprising. The very large values and the sharp decrease in the interval $[0, 7]$ suggest that we cannot estimate accurately all three parameters (in particular K) if we sample data only from $[0, 7]$, regardless of the number of points sampled there. This is the direct consequence of the sensitivity with respect to K (see Figure 1(b)), which shows that the interval $[0, 7]$ does not carry any information with respect to K. Another interesting feature of the condition number is its continuous improvement after $T = 20$. This doesn't necessarily mean that larger final times T are better than $T = 20$. This is an artifact, again due to the sensitivity with respect to K, which reaches a steady state of one instead of zero (the other two sensitivities are essentially zero after $T = 20$). This results in updating the FIM with rank one matrices with only one nonzero entry, therefore slightly improving its condition number.

5.2.3 Correlation Coefficients

It is also of interest to analyze the evolution of the correlation coefficients for our problem (see Figure 4(d)) as functions of the final time T (recall $\text{Corr}(\hat{\Theta}_i, \hat{\Theta}_j) = \frac{(\text{Cov}\,\hat{\Theta})_{ij}}{\sigma_{\hat{\Theta}_i}\sigma_{\hat{\Theta}_j}}$). It is well known that in general it is more difficult to estimate parameters which are strongly correlated than those which are moderately correlated. By looking at the evolution of the correlation coefficients, and comparing their rates of change, we can conclude that $T = 20$ is a reasonable choice.

In conclusion, the graphs of the TSFs, GSFs, the condition number of the FIM, and the correlation coefficients corresponding to true parameter $\theta_0 = (17.5, 0.7, 0.1)$ suggest in agreement that most of the relevant information is found in $[0, 20]$ and little is left outside this interval.

5.3 Use of Prior Knowledge/2D TSF and GSF

In the previous subsection we presented mathematical tools which can help us identify an optimal final time *when the true parameter θ_0 is known*. However, it is very important to note that in practice the true value θ_0 *is not known a priori*. We face a conundrum here, since the mathematical tools we want to use to indicate an optimal stopping time (and later an optimal sampling distribution), depend in fact on the *unknown true value* of the parameters we wish to estimate. A generally accepted way to alleviate this predicament is to *use prior knowledge*, typically credited to experimentalists' expertise.

In optimal design [12] (as well as in general optimization algorithms), one often makes the assumption that a good initial guess for θ_0 is available, and locally optimal sampling distributions are found based on this initial guess. Our assumption here will be that we have a range of reasonable values (available from prior knowledge) for the components of θ_0. We note that this is a much weaker assumption than having a good initial guess for θ_0. In consequence, instead of using the 1D version (corresponding to θ_0) of the TSF, GSF and the condition number, we need to use the 2D versions of these functions, with double dependence on the final time T and corresponding parameters θ_k varying in some *a priori* given window around the true value (see Figure 5).

A natural question to address now is: How do we choose a stopping time T in the more general case of the 2D maps plotted in Figure 5? To answer this question, we use again the

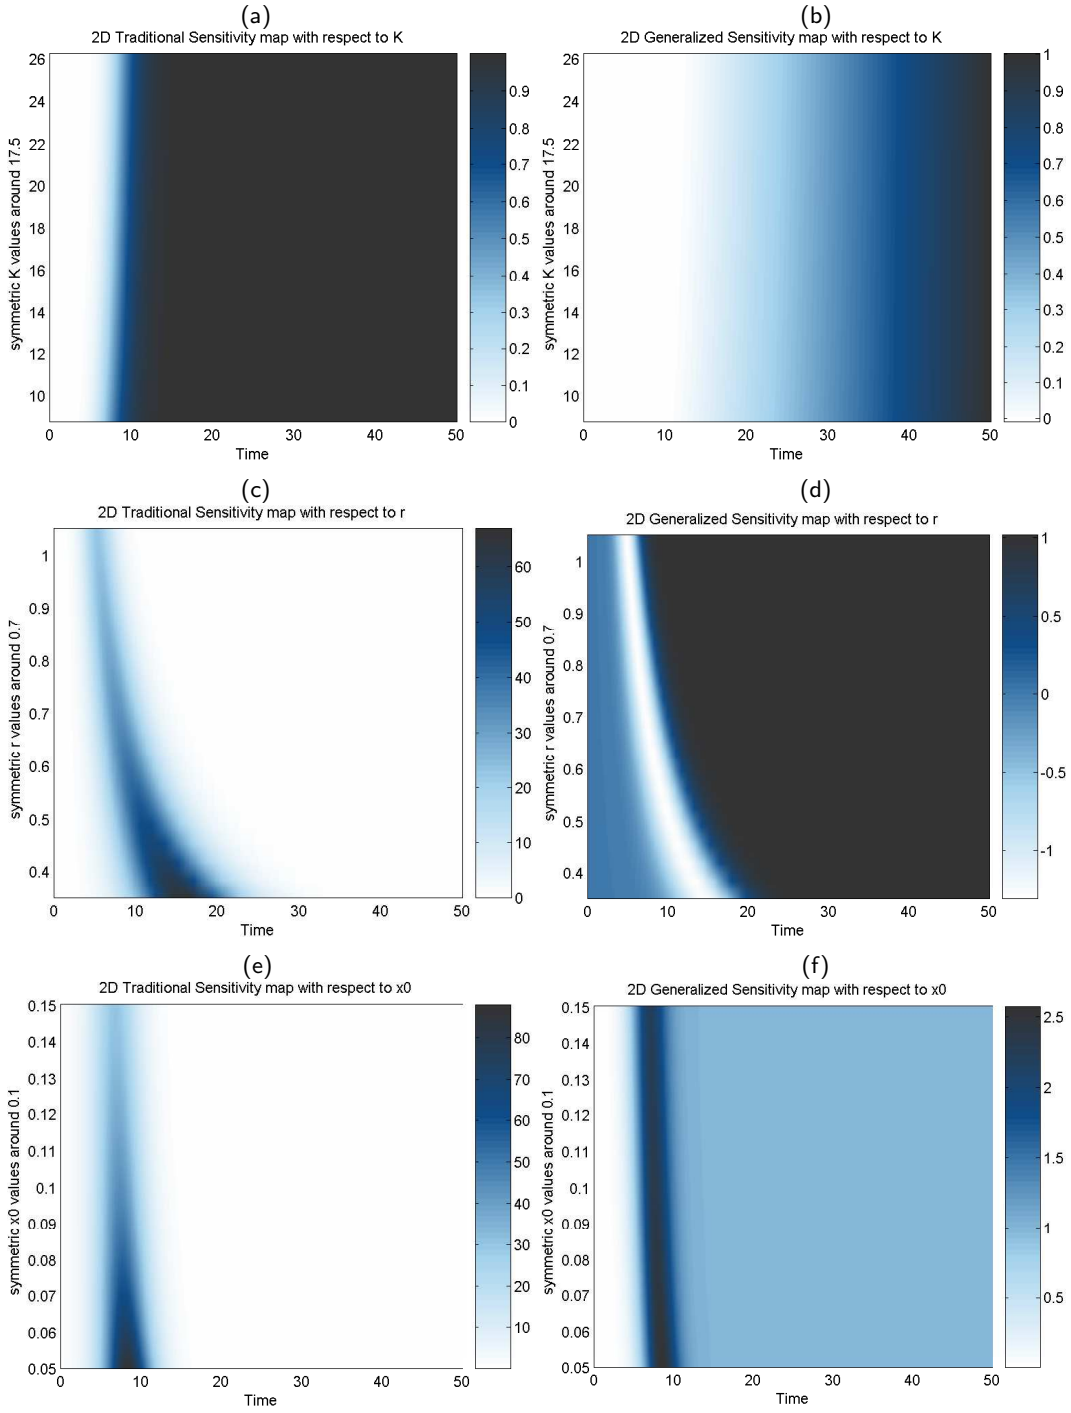

Figure 5: (a) and (b): 2D TSF and GSF maps for $T = 50$ and K around 17.5 (true K); (c) and (d): 2D TSF and GSF maps for $T = 50$ and r around 0.7 (true r); (e) and (f): 2D TSF and GSF maps for $T = 50$ and x_0 around 0.1 (true x_0).

"worst-case scenario" approach, but this time with respect to the 2D maps. More specifically, we choose a minimum T after which both the TSFs and the GSFs do not exhibit any significant transient behavior, both as a functions of time and the corresponding parameters.

For example, in Figure 5,(a) and (b), we plotted the 2D TSF and GSF maps corresponding to $T = 50$ and K varying in the parameter window $[K_0 - \delta K_0, K_0 + \delta K_0]$ with $\delta = 0.5$. We see that for smaller values of K, the solution is insensitive to changes in K until around $t = 7$, whereas for larger values, the change occurs slightly higher at around $t = 10$. This change is more obvious when we consider the parameter r. In Figure 5(c) and (d), we consider the TSF and the GSF maps corresponding to r over the range $[r_0 - \delta r_0, r_0 + \delta r_0]$, again with $T = 50$ and $\delta = 0.5$. Now we see that for the lower values of r, the solution is sensitive to changes in r for a longer time period than when considering larger values of r. Finally, in Figure 5(e) and (f), we plotted the TSF and the GSF maps corresponding to x_0. We notice small changes when x_0 varies over $[x_0 - \delta x_0, x_0 + \delta x_0]$. Next, we apply the same idea of considering the maximum of the individual final times, i.e.,

$$T_{2D} = \max\Big(\max_{K\in D_K} T_K(K), \max_{r\in D_r} T_r(r), \max_{x_0\in D_{x_0}} T_{x_0}(x_0)\Big),$$

where D_K, D_r, D_{x_0} refer to the corresponding parameter window centered around the true values, and T_K, T_r, and T_{x_0} are defined as in (5.5). The resulting optimal final time points T_{2D} can be seen in Table 3 for various δ values. As expected, we see that better prior knowledge (as

δ	0	0.25	0.4	0.5
T_{2D}	20	21	27	33

Table 3: Optimal final time T_{2D}, for various values of δ.

reflected by the decreasing magnitude of δ, which controls the width of the parameter window) yields smaller values for the optimal final time T_{2D}, and a closer estimate to the final time value corresponding to the true value parameter θ_0 with no uncertainty.

5.4 Numerical Simulations

In order to verify the analysis carried out in the previous subsections, we perform numerical simulations which demonstrate that the TSFs used in conjunction with the GSFs and the condition number of the Fisher information matrix, are efficient tools in determining an appropriate final time.

To illustrate our ideas, we solve the inverse problem of estimating K, r and x_0 from noisy data using ordinary least squares repeatedly, with data sampled from gradually increasing time intervals $[0, T]$. Our goal is to investigate the evolution of the corresponding standard errors as functions of T, and to identify a final time based on when the standard errors with respect to K, r and x_0 reach a steady state or exhibit negligible variation. We allow T to increase by one unit at a time, starting from $T = 10$ until $T = 50$, a value that is sufficiently high for the purposes of these simulations. For consistency reasons, the data sets have a hierarchical structure, i.e., the data set from $[0, T + 1]$ contains all the data from the interval $[0, T]$ plus an additional point sampled at $T + 1$. We do not consider starting values smaller than $T = 10$, since the resulting data do not have sufficient information about our parameters as reflected by the very large values of the standard errors and of the condition number. We estimate our parameters from synthetic noisy data obtained by adding Gaussian noise with zero mean and variance $\sigma_0^2 = 0.16$ to the deterministic solution of the logistic model corresponding to the true value $\theta_0 = (17.5, 0.7, 0.1)$. The optimization algorithm we use to solve the least squares problems is

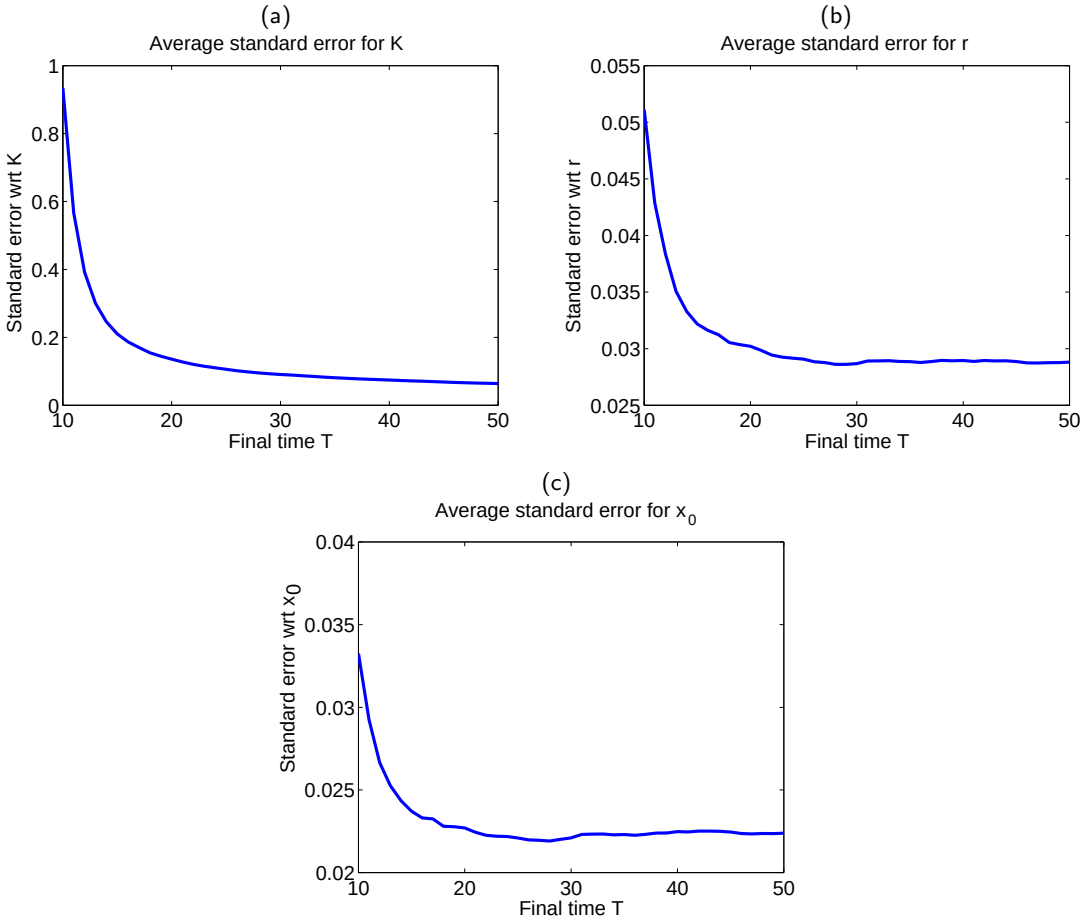

Figure 6: Average standard errors as functions of the final time T, SE_K (panel (a)), SE_r (panel (b)), and SE_{x_0} (panel (c)). Averages were obtained over 100 sets of noisy data with $\sigma_0^2 = 0.16$.

based on the Nelder-Mead method and implemented by the MATLAB function fminsearch. We use the initial guess $\theta^0 = 1.4 \cdot \theta_0 = (24.5, 0.98, 0.14)$ for our optimization routines. In order to avoid potential misleading results caused by using only a particular data set, we base our conclusions on the mean standard errors, obtained by solving the least squares problem 100 times with 100 different data sets (for each final time T) and averaging the results.

We plot the evolution of the average standard errors in Figure 6 and we note that their magnitudes remain consistently low and exhibit very little variation after $T = 25$. However, in all three plots, the sharpest decrease in magnitude is clearly attained by $T = 20$. This suggests that a minimum sampling interval should be $[0, 20]$ although a more conservative choice is $[0, 25]$. The same graphs suggest that adding additional data points sampled beyond T=25 will contribute very little to the goal of accurately estimating K, r and x_0. Thus the results of these numerical simulations confirm that the mathematical tools presented in the previous subsection (TSFs, GSFs, condition number of the FIM and correlation coefficients) are efficient tools in determining an appropriate final time T. We attribute the slight mismatch ($T = 20$ returned by mathematical tools vs. $T \approx 25$ returned by numerical simulations) mainly to the following two factors:

- the value $T = 20$ was returned by deterministic mathematical tools *based on the particular nominal vector* $\theta_0 = (17.5, 0.7, 0.1)$, and
- the uncertainty introduced by using relatively noisy data ($\sigma_0^2 = 0.16$) for numerical simulations.

We conclude that $T = 25$ should be a sufficient final time in order to guarantee accurate estimation of our model parameters. However, depending on the cost of collecting data, it may be necessary to be less conservative with the final time choice, and a final time of $T = 20$ should be reasonable as well. Finally, instead of basing our choice of the final time on the 1D version of the sensitivity functions or the FIM condition number corresponding to a particular value of θ_0, we may use the corresponding 2D maps and the results given in Table 3.

6 Choosing a Sampling Distribution in $[0, T]$

As we have seen in Subsection 3.2, the study of the sensitivity of the estimator $\hat{\Theta}$ with respect to model parameters θ naturally led to the definition of the Fisher information matrix and the generalized sensitivity functions (3.10). In the previous section, we have shown how these mathematical tools can be used to find an appropriate final time for data collection in a parameter estimation problem.

Having chosen a final time T, we find the next step in our investigation is to answer the following questions:

1. How do we estimate *a priori* the sample size, i.e., the number n of points sampled in $[0, T]$, needed to achieve certain target levels for the standard errors?
2. How to choose an optimal sampling distribution in $[0, T]$, once the sample size n and the final time T are determined?

In the Appendix, Section A.3, we show that in case of general sampling distributions the weighted least squares estimator $\hat{\Theta}$ is unbiased, i.e., $E\big(\hat{\Theta}\big) = \theta_0$, with a covariance matrix Σ_0 which is the inverse of the Fisher information matrix (A.23). Our strategy in answering the first question above is to relate the standard errors for discrete, uniform sampling distributions in $[0, T]$ to those for the continuous case. This allows us to estimate the size of *discrete uniform sampling distributions* necessary to attain given targets for the magnitudes of the standard errors without re-computing the covariance matrix for each specific uniform discrete sampling distribution. We consider this approach in detail in Subsection 6.1.

There are several approaches one can take when trying to answer the second question. A *first approach* is to consider graphical tools like the traditional sensitivity functions (3.1) and the generalized sensitivity functions (3.10) or their first and second time derivatives (see Figures 1 and 2), whose monotonicity indicates regions of high information content with respect to model parameters from which to sample data. Sampling additional data points in these regions yields typically more accurate estimates as illustrated by the magnitude of the standard errors. A *second approach* is to use an optimal design criterion (for example the D-, E- or c-optimal designs, or the one we use in Subsection 6.3), which will return an optimal sampling distribution in $[0, T]$. In the next Subsection 6.2 we discuss the first approaches in great detail, whereas the second approach is discussed in Subsection 6.3.

6.1 Choosing the Size of a Sampling Distribution

In this subsection we address the question: Given a final time T for our experiment, can we estimate *a priori the size of a uniform sampling distribution* in $[0, T]$ (that is the number of uniformly spaced observations to take in $[0, T]$) which would yield estimates with a desired

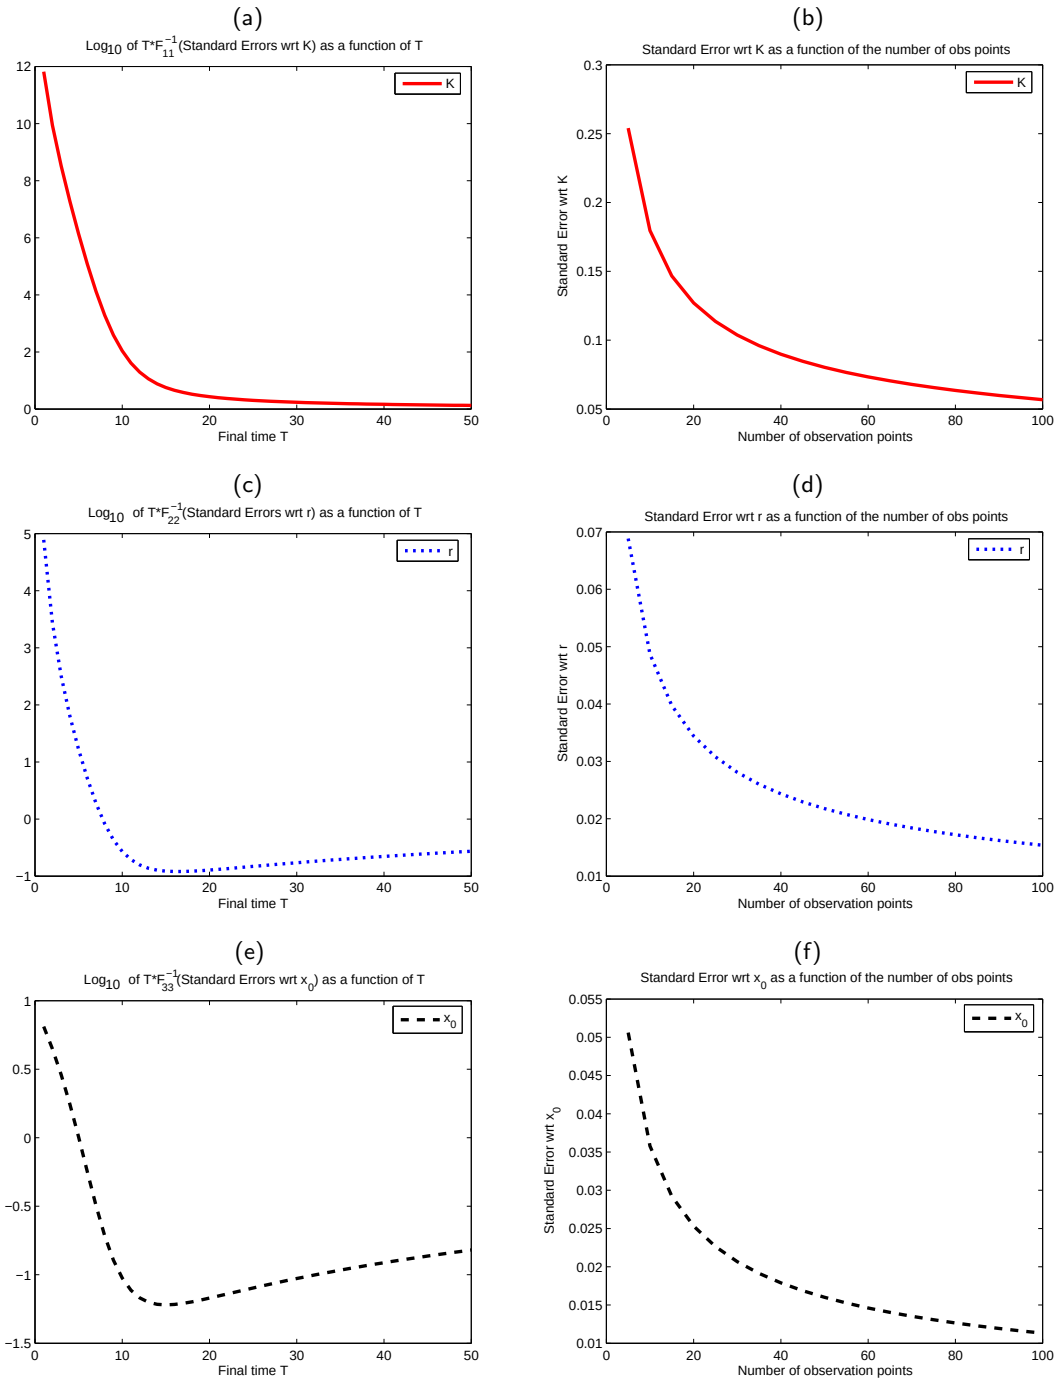

Figure 7: (a), (c) and (e): behavior of $\mathrm{diag}(TF^{-1})$ (the standard errors, up to a multiplicative constant) as a function of the final time T; (b), (d) and (f): behavior of the standard errors as a function of the number of observations for fixed $T = 25$.

accuracy? In order to answer this question we we set $\Delta t = T/n$ and introduce the uniform

meshes $t_i^{(n)} = i\Delta t$, $i = 1, \ldots, n$, $n = 1, 2, \ldots$, for our discrete sampling procedures.

According to (A.5) in Subsection A.1.2, we see that the standard errors for the discrete case are given by

$$SE_k^{(n)} = \sqrt{\left(\left(\chi_w^{(n)}(\theta_0)\right)^{\mathsf{T}} \chi_w^{(n)}(\theta_0)\right)^{-1}_{kk}}, \quad k = 1, \ldots, p,$$

where

$$\left(\chi_w^{(n)}(\theta_0)\right)^{\mathsf{T}} \chi_w^{(n)}(\theta_0) = \sum_{i=1}^{n} \frac{1}{\sigma^2(t_i^{(n)})} \nabla_\theta^{\mathsf{T}} f(t_i^{(n)}, \theta_0) \nabla_\theta f(t_i^{(n)}, \theta_0) := F^{(n)}(T, \theta_0).$$

From Remark 3.1, b), we see that

$$\lim_{n\to\infty} \frac{T}{n} F^{(n)}(T, \theta_0) = F_C(T, \theta_0),$$

where F_C is the continuous time Fisher information matrix

$$F_C(T, \theta_0) \equiv \int_0^T \frac{1}{\sigma^2(t)} \nabla_\theta^{\mathsf{T}} f(t, \theta_0) \nabla_\theta f(t, \theta_0)\, dt, \tag{6.1}$$

i.e., the Fisher information matrix with P being taken as Lebesgue measure on $[0, T]$. For n sufficiently large we can set $F^{(n)}(T, \theta_0) \approx (n/T) F_C(T, \theta_0)$ and obtain from (A.4)

$$\mathrm{Cov}\left(\hat{\Theta}^{(n)}\right) \approx \frac{T}{n} F_C(T, \theta_0)^{-1}, \tag{6.2}$$

which implies

$$SE_k^{(n)} \approx \sqrt{\frac{T}{n}\left(F_C(T, \theta_0)^{-1}\right)_{kk}}, \quad k = 1, \ldots, p. \tag{6.3}$$

Equations (6.2) and (6.3) relate the covariance matrix and the standard errors we would obtain with n uniformly distributed time observations in $[0, T]$ to the inverse of the continuous FIM (6.1). We can use expression (6.3) to obtain *a priori* estimates on the number of sample points n and on the final time T needed to obtain standard errors of desired magnitudes. For example if the final time T is fixed, and if we have prior knowledge on the noise level (we do not need to know the variances σ_i^2 exactly, we just need an upper bound for them), we can use (6.3) to evaluate the minimum number n of observations we have to take (conservative choice) in order to achieve a given accuracy for our parameters estimates (see Figure 7(b), (d) and (f)). Similarly, when n is fixed we can study the behavior of the standard errors as functions of the final time which can also help us choose the appropriate final time T (see Figure 7(a), (c) and (e)).

6.2 Sensitivity Functions as Design Tools – Numerical Experiments

In the previous subsection, we related the discrete version of the Fisher information matrix to its continuous counterpart and we derived formula (6.3) which helps us to *estimate a priori* the number of time observations in a uniform sampling distribution that allows us to achieve a certain targeted accuracy for our parameters.

In this subsection, we present numerical experiments which illustrate the utility of the traditional and generalized sensitivity functions to parameter estimation problems. The approach we consider here is motivated by the analysis carried out in Section 3 where we showed that the graphs of the sensitivity functions indicate regions of high information content with respect to the parameters to be estimated. As we shall see in the following, sampling additional data points in these regions yield more accurate estimates as illustrated by the magnitude of the standard

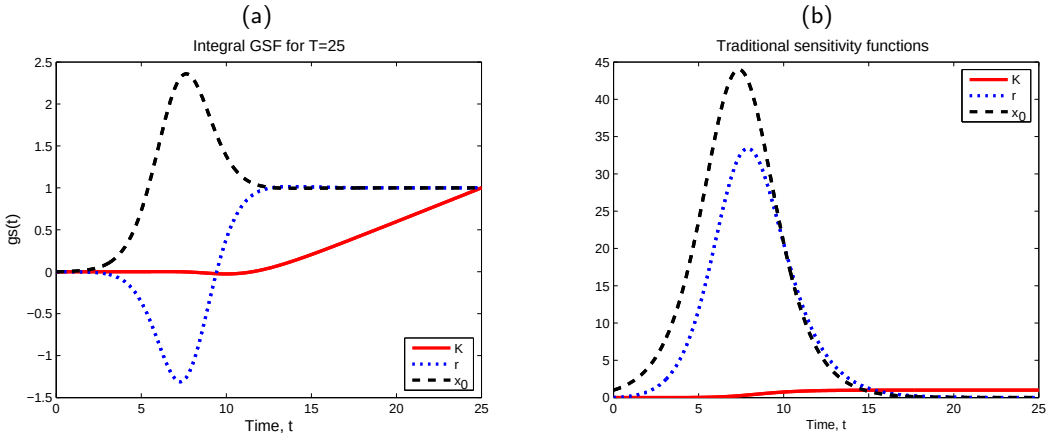

Figure 8: Generalized sensitivity functions (panel (a)) and traditional sensitivity functions (panel (b)) plotted for the *true parameter vector* $\theta_0 = (17.5, 0.7, 0.1)$ in the logistic model.

errors. Thus, we advocate the use of the traditional and generalized sensitivity functions as optimal design tools in parameter estimation problems.

Before we begin our numerical investigations, we revisit shortly the graphs of the generalized sensitivity functions and of the traditional sensitivity functions. For convenience, we plot them side by side in Figure 8 for the time interval $[0, 25]$. Based on the evolution of the generalized sensitivity functions, we distinguish two well defined regions approximately given by the intervals $R_1 = [0, 12.5]$ and $R_2 = [12.5, 25]$. On R_1 the GSFs with respect to r and x_0 exhibit a large transitional behavior, whereas the GSF with respect to K is basically zero. On R_2 the GSF with respect to K shows a sharp steady increase from zero to one, whereas the GSFs with respect to r and x_0 show no transitional behavior and are almost constant equal to one due to a phenomenon which we previously called the "forced-to-one" artifact [3, 8]. This is in total agreement with the information provided by the time derivatives of GSFs with respect to r and x_0, which essentially are supported in R_1 and are almost vanishing in R_2, and the time derivatives of the GSF with respect to K, which is zero in R_1 and supported in R_2 (see Figure 2, (d)).

The TSFs provide insight consistent with the information given by the GSFs, although the time regions where the TSFs are nonzero are typically larger and include the ones specified by the GSFs (see Figure 8). For example the TSFs with respect to r and x_0 are supported in the interval $[0, 20]$ which includes R_1, and the TSF with respect to K is supported in $[7.5, 25]$ which includes R_2. This inclusion is quite intuitive. Recall that the TSFs give the sensitivity of the model output with respect to model parameters whereas the GSFs give the sensitivity of the estimates with respect to model parameters. Therefore, it makes sense that the GSF regions should be included in the TSF regions, a fact which can also be argued by using the definition (3.10) of the GSFs. Aside from this inclusion, of most importance is the fact that the region of maximum sensitivity, i.e., in our case approximately the interval $[5, 10]$, is commonly specified by both the GSFs and the TSFs.

Next, we present numerical experiments which help us to illustrate that in order to obtain more accurate estimates, the best way to add extra observations to an existing distribution is to sample them in regions of high information content specified by the sensitivity functions. To discuss this, we consider the problem of estimating the parameters K, r and x_0 using ordinary least squares and noisy data sampled on various distributions in the interval $[0, 25]$. The noisy data is obtained by adding Gaussian noise with zero mean and constant variance $\sigma_0^2 = 0.16$ to the solution of the logistic equation (we use the MATLAB-function randn). We solve the

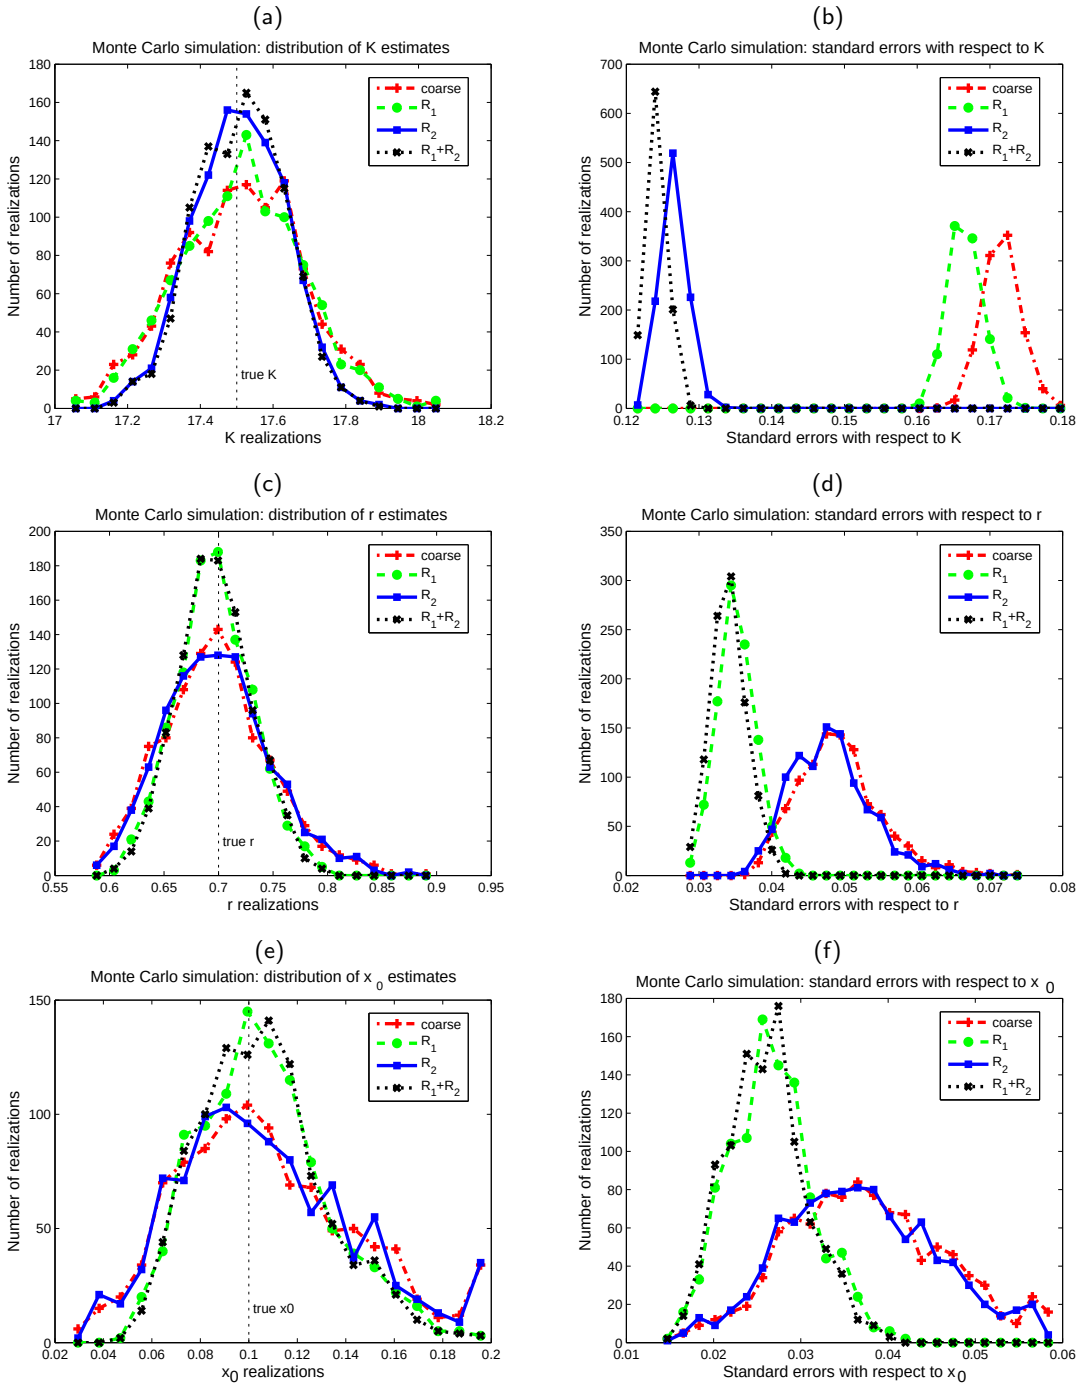

Figure 9: Monte Carlo simulations with 1000 iterations: (a), (c), (e) distribution of the estimates for K, r, x_0; (b), (d), (f) distribution of the corresponding standard errors for the estimates of K, r, x_0.

nonlinear least squares problems which arise by using the specially designed MATLAB-function lsqnonlin in the Optimization Toolbox. In order to compare between various sampling strategies, we use the same initial guess and the same constraint options (upper and lower bound, stopping criteria, etc.) for the optimization algorithm throughout the numerical simulations. To avoid committing an inverse crime, we generate the noisy data by computing the solution of the logistic equation numerically (using the MATLAB-function ode45), whereas we solve the inverse problem by utilizing the closed form solution of the logistic model. Finally, to quantify uncertainty in the results obtained with different sampling strategies, we follow standard statistical practice and compare the magnitudes of the standard errors. However, since the standard errors are defined in terms of the true parameter vector θ_0 which is unknown (see equation (2.7)), in practice one replaces θ_0 with the estimate $\hat{\theta}_0$. As a consequence, we obtain different standard errors for different noisy data sets. This leads us to consider Monte Carlo (MC) simulations with a large number of iterations for a given sampling strategy and then compare the different sampling strategies based on the average and on the distribution of the standard errors obtained.

In our first numerical experiment, we compare between various strategies to add extra points to an existing (*coarse*) distribution and we analyze the results by comparing the magnitude of the average MC standard errors. Given $a, d \in \mathbb{R}$ and $b = kd$ for some $k \in \mathbb{N}$, we use the MATLAB notation $[a : d : b] = \{a, a + d, \ldots, a + kd\}$. Then the sampling distributions S_{coarse}, S_{R_1}, S_{R_2} and $S_{R_1 \cup R_2}$ are given by

$$S_{\text{coarse}} = [0 : 2.5 : 25], \quad S_{R_1} = S_{\text{coarse}} \cup [1.25 : 2.5 : 11.25],$$
$$S_{R_2} = S_{\text{coarse}} \cup [13.75 : 2.5 : 23.75], \quad S_{R_1 \cup R_2} = S_{R_1} \cup S_{R_2} = [0 : 1.25 : 25].$$

Sampling Distribution	S_{coarse}	S_{R_1}	S_{R_2}	$S_{R_1 \cup R_2}$
Number of points	11	16	16	21
SE_K	0.1715	0.1665	0.1265	0.1241
SE_r	0.0492	0.0351	0.0481	0.0341
SE_{x_0}	0.0375	0.0266	0.0369	0.0260

Table 4: Average standard errors corresponding to four sampling strategies obtained with Monte Carlo simulation for 1000 iterations.

As noted we use a Monte Carlo procedure, and we solve the inverse problem repeatedly 1000 times for a given sampling distribution, each time with a different noisy data set, but keeping the same characteristics and optimization options as discussed above. In Figure 9 we plot a digitized version of the empirical distributions obtained for the estimates of K, r and x_0 along with the corresponding distributions of the standard errors. More precisely, we plot the number of realizations (out of our sample of 1000) against values of the realization in the left panels. Different curves in a given plot correspond to different additional (to a basic coarse sampling grid) sampling regions (R_1, R_2, etc.) in each of the 1000 MC runs. The average Monte Carlo standard errors obtained are displayed in Table 4.

As one can see from this table, the average SE_K obtained with the coarse sampling distribution S_{coarse} is 0.1715. When we sample 5 additional data points in R_1, i.e., we use the sampling distribution S_{R_1} the average SE_K decreases from 0.1715 to 0.1665 which represent only a 2.92% increase in accuracy. A totally different scenario occurs if we sample the 5 extra points in R_2, i.e., we take the sampling distribution S_{R_2}. Note that according to the GSF with respect to K, R_2 is a region of high information content for K. Now the average SE_K decreases from 0.1715 to 0.1265, which means a 26.24% increase in accuracy. If in addition to the 5 points sampled

in R_2, we also add the 5 points in R_1 in order to get the sampling distribution $S_{R_1 \cup R_2}$ the qualitative gain in accuracy is minor (only 27.64% increase in accuracy compared to 26.24%).

We observe a similar behavior if we consider the average MC standard errors with respect to r and x_0, but this time with the regions R_1 and R_2 interchanged. We see an important drop in magnitudes, if we compare the standard errors obtained with the sampling distributions S_{coarse} and S_{R_1} (28.66% and 29.07% improvement in accuracy for r and x_0, respectively) but only minor gains when we compare S_{coarse} to S_{R_2} or S_{R_1} to $S_{R_1 \cup R_2}$. The conclusions presented above are graphically supported by the plots in Figure 9. As one can see, the distributions of the standard errors (as well as those of the parameter estimates) are grouped two by two showing important increases in the accuracy for a parameter (average *SE* considerably shifted to the left and a narrower support of the curves) when extra data points are taken in the regions specified by the GSF, and only insignificant gains when sampled outside.

Based on the evolution of the generalized sensitivity functions, and also supported by numerical results, we have shown that the interval $R_1 = [0, 12.5]$ includes the region of highest information content with respect to r and x_0, and $R_2 = [12.5, 25]$ contains the region of highest information content with respect to K. However, unlike the GSF with respect to K which shows a steady increase from 0 to 1 throughout the region R_2, the GSFs with respect to r and x_0 change monotonically in R_1, and exhibit both, regions of increase and decrease (see Figure 8). Our next goal is to verify numerically the analytical insight stated at the end of the Subsection 3.2, namely that the regions of the steepest increase and decrease on the graphs of the GSF curves (in our case the interval $[5, 10]$, the maximum sensitivity region) correspond to regions of high information content with respect to the parameters to be estimated. In particular we will illustrate that region with steepest increase is not necessarily the one which carries the maximum information content (as argued by Thomaseth and Cobelli [34]) and that the region with steepest decrease may carry relevant information as well. As our results will demonstrate, the sensitivity functions are efficient tools which can help us design enhanced sampling strategies for adding additional data points to a coarse distribution.

In our second numerical experiment, we systematically add eight extra points in the intervals $[2.5, 7.5]$ and $[7.5, 12.5]$, respectively $[5, 10]$, to the distribution S_{coarse}, so that the sampling distribution in the intervals $[2.5, 7.5]$ and $[7.5, 12.5]$, respectively $[5, 10]$, is again uniform. This gives the sampling distributions

$$S_{[2.5,7.5]} = S_{\text{coarse}} \cup [2.5 : 0.5 : 7.5], \quad S_{[5,10]} = S_{\text{coarse}} \cup [5 : 0.5 : 10],$$
$$S_{[7.5,12.5]} = S_{\text{coarse}} \cup [7.5 : 0.5 : 12.5] \quad S_{\text{union}} = S_{\text{coarse}} \cup [2.5 : 0.5 : 12.5].$$

The intervals where we add additional points are of equal length and correspond approximately to regions of steepest increase, steepest decrease and maximal magnitudes of the GSFs and the TSFs with respect to r and x_0. As with our first numerical experiment, we use a Monte Carlo procedure and we estimate the parameters K, r and x_0 repeatedly 1000 times, each time with a different data set. The average MC standard errors are given in Table 5. We also considered a sampling distribution $S^*_{[5,10]}$ in which two points are removed for $S_{[5,10]}$ (see details below).

Sampling Distribution	S_{coarse}	$S_{[2.5,7.5]}$	$S_{[7.5,12.5]}$	$S_{[5,10]}$	$S^*_{[5,10]}$	S_{union}
Number of points	11	19	19	19	17	27
SE_K	0.1715	0.1684	0.1593	0.1688	0.1963	0.1538
SE_r	0.0492	0.0338	0.0373	0.0289	0.0291	0.0241
SE_{x_0}	0.0375	0.0226	0.0306	0.0216	0.0217	0.0175

Table 5: Average Monte Carlo standard errors obtained with 1000 iterations.

Since the sampling distributions corresponding to the intervals above do not contain extra points sampled in R_2 (the region of highest information content with respect to K), as expected, the corresponding standard errors for K in Table 5 do not improve considerably. However, the standard errors with respect to r and x_0 displayed in the same table reveal significant reduction in the uncertainty of estimation for the two parameters. We note that the intervals $[2.5, 7.5]$ and $[7.5, 12.5]$ correspond to the steepest increase regions of the GSF with respect to x_0 and r respectively. The same intervals in reversed order correspond to the steepest decrease regions of the GSFs with respect to these two parameters. We observe that by adding the same number of eight extra points in the intervals $[2.5, 7.5]$ and $[7.5, 12.5]$ the magnitude of the standard errors with respect to r and x_0 decreases considerably, but *the best performance* is achieved when adding these extra points in $[5, 10]$, the interval of maximum sensitivity, as illustrated by the graphs of the GSFs and TSFs. The corresponding distribution $S_{[5,10]}$ yields a 41.41% reduction in uncertainty for r and 42.28% for x_0.

An interesting fact to report is that the steepest decrease regions of the GSF contain a great deal of information content with respect to the two parameters, and in the case of r, the steepest decrease region outperforms the steepest increase one (see the standard errors corresponding to the sampling distributions $S_{[2.5,7.5]}$ and $S_{[7.5,12.5]}$). Therefore, we suggest that the steepest decrease regions are equally valuable as those of steepest increase, and arguably they should be considered when designing optimal sampling strategies. The regions of steepest increase and decrease of the GSF curves, with maxima or minima exceeding 1 in absolute value, are simply artifacts of high correlation between parameters. In absence of correlation, these regions will yield a single steepest increase region as in Figure 3(a).

Another interesting fact to observe in Table 5 is that the distribution $S^*_{[5,10]}$ (obtained by discarding two points from $S_{[5,10]}$) which in spite of the fact it consists of only 17 sampling points, yields very close performance to $S_{[5,10]}$ for r and x_0. The two points of S_{coarse} discarded in forming $S^*_{[5,10]}$ were discarded based on the relative magnitudes of the $\varepsilon(t)\nabla^2_{\theta\theta} f(t, \theta_0)$ and $\nabla^{\mathsf{T}}_{\theta} f(t, \theta_0)\nabla_{\theta} f(t, \theta_0)$ terms in equation (3.4). (Note that the factor $\nabla^2_{\theta\theta} f(t, \theta_0)$ amplifies the noise $\varepsilon(t)$, which for the purpose on this example was assumed to be normally distributed). The design of this distribution is an example which shows the power of sensitivity functions in optimal design and their potential applicability in high cost data point collection experiments. With only 17 sampling points $S^*_{[5,10]}$ delivers the same performance as the sampling distribution S_{union} with 27 points when estimating r and x_0 (see Figure 7(d) and (f)).

6.3 Using Optimal Design Criteria

In this subsection we advocate a new optimal design criterion which is based on the idea of finding a measure P of the form (2.4), i.e.,

$$P = \sum_{i=1}^{n} \delta_{t_i},$$

which provides the best discrete approximation for the FIM F_C obtained with $P = \mu$, where μ is the Lebesgue measure on $[0, T]$. The FIM corresponding to P is denoted by F_D. In other words, we want to find the optimal time distribution $t_1 < t_2 < \cdots < t_n$ in $[0, T]$ such that

$$F^n_D = \sum_{j=1}^{n} \frac{1}{\sigma^2(t_j)} \nabla_{\theta} f(t_j, \theta_0)^{\mathsf{T}} \nabla_{\theta} f(t_j, \theta_0) \tag{6.4}$$

best approximates

$$F_C = \int_0^T \frac{1}{\sigma^2(s)} \nabla_{\theta} f(s, \theta_0)^{\mathsf{T}} \nabla_{\theta} f(s, \theta_0)\, ds. \tag{6.5}$$

This design criterion is motivated by the intuitive notion that the best case scenario we can possible have when estimating the parameter θ_0 is the one when we have available all the measurements at all times in the interval $[0, T]$. The natural approach when we have a limited number of observations is therefore to find the sampling distribution which gives the best discrete approximation for F_C. By approximation we mean approximation with respect to the Frobenius norm of matrices, i.e., we use the design functional

$$\mathcal{J}_1(F_D^n) = \sum_{i=1}^{p}\sum_{j=1}^{p} \left|(F_C)_{ij} - (F_D^n)_{ij}\right|^2, \tag{6.6}$$

For comparison we also use the design functional

$$\mathcal{J}_2(F) = \sum_{i=1}^{p} \frac{1}{\theta_{0,i}^2}\left(F^{-1}\right)_{ii}. \tag{6.7}$$

This functional is the sum, over all components of the parameter vector, of the squared normalized standard errors for the components of the parameter vector and can be defined for general probability measures P.

Numerical Simulations

We consider two inverse problems:

1. We assume that we know the true parameter value θ_0, and seek to optimize the placement of time points for data collection. As an optimization criterion we will first minimize J_1. In order to optimize the placement of the time points t_i, we choose n positive numbers μ_i, $i = 1, \ldots, n$, such that $\sum_{i=1}^{n} \mu_i = T$, where T is the final time. Then our time points are established as $t_1 = 0$ and $t_i = t_{i-1} + \mu_{i-1}, i = 2, \ldots, n+1$. We initially chose $\mu_i = T/n$, creating a uniformly spaced time distribution, t_{unif}. Then we use fmincon, a constrained gradient-based optimization algorithm in MATLAB, to return an optimal μ-vector, which in turn produces an optimal t-vector, denoted by t_{opt}. We note that this will produce an optimal mesh among those meshes containing both 0 and T; however, this mesh may not be optimal among all possible sampling distributions including those in which 0 or T or both may be missing.
2. We then considered a parameter estimation problem which compares the effectiveness of t_{unif} versus t_{opt}. We created noisy data $y(t)$ by generating the logistic curve using the true parameter vector θ_0 and then adding Gaussian noise with $\sigma_0^2 = 0.16$. Then we solve two parameter estimation problems (one using only the points in t_{unif}, and the other using only the points in t_{opt}) to determine which grid best recovers θ. We use the least squares functional
$$J(\theta) = \sum_{i=1}^{n+1} \left|y(t_i) - f(t_i, \theta)\right|^2 \tag{6.8}$$
to determine the optimal θ in each case. We also considered the corresponding standard errors.

For these simulations we use $n = 10$, $\theta_0 = (17.5, 0.7, 0.1)$ and $T = 25$. We run the first inverse problem one time for the cost functional (6.6) to obtain an optimal sampling distribution. We then run the second inverse problem (for (6.8)) 1000 times (using an initial estimate $\theta^0 = 1.4 \cdot \theta_0$), and compute the average $\hat{\theta}^n$ and SE-values over the total number of simulations. The resulting standard errors for each parameter can be seen in Figure 10(a) – (c). Then in Figure 10(d), we depict the logistic curve plotted using the true θ_0, along with the time point location for both

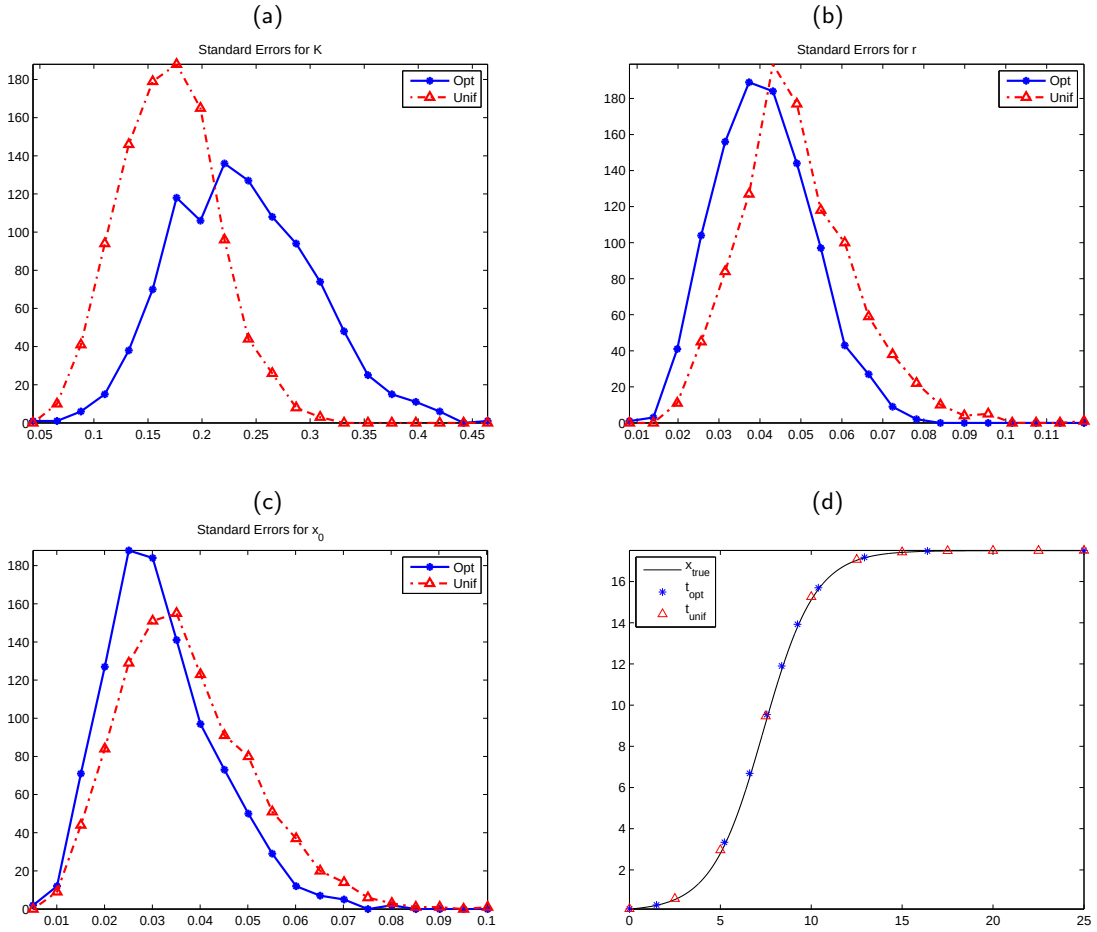

Figure 10: Comparison of optimal (using $\mathcal{J}_1$) versus uniform sampling time points in terms of SE_K (panel (a)), SE_r (panel (b)), SE_{x_0} (panel (c)) and distribution of time points, plotted on the logistic curve (panel (d)).

t_{unif} and t_{opt}. The average θ value over these 1000 runs is $(17.523, 0.69728, 0.10743)$, and the average standard errors are approximately $(0.24, 0.042, 0.033)$.

Next, we run the first inverse problem one time for the cost functional (6.7). We then run the second inverse problem 1000 times (using an initial estimate $\theta^0 = 1.4 \cdot \theta_0$), and compute the average $\hat{\theta}^n$ and SE values over the total number of simulations. The resulting standard errors for each parameter can be seen in Figure 11(a) – (c). Then in Figure 11(d), we plot the logistic curve using the true θ_0, along with the time point locations for each of t_{unif} and t_{opt}. The average θ value over these 1000 runs is $(17.5062, 0.6992, 0.1045)$, and the average standard errors are approximately $(0.23, 0.037, 0.027)$.

Overall, it appear that for this example the second cost functional (6.7), which considers the standard errors as opposed to the continuous Fisher information matrix, performs slightly better (although both functionals perform adequately).

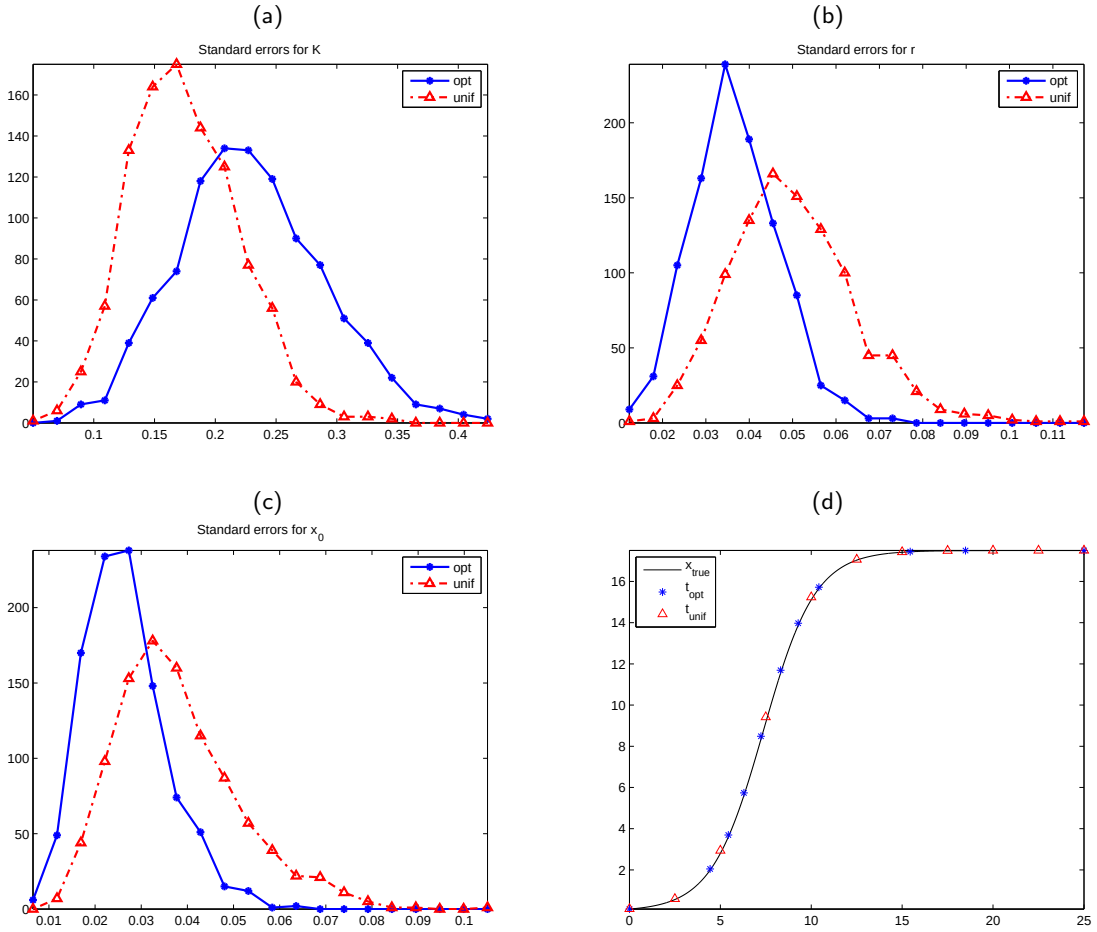

Figure 11: Comparison of optimal (using J_2) versus uniform time points in terms of SE_K (panel (a)), SE_r (panel (b)), SE_{x_0} (panel (c)) and distribution of time points, plotted on the logistic curve (panel (d)).

Appendix

A.1 Asymptotic Properties of the Least Squares Estimator – Standard Theory

In this section we summarize results from the literature [7, 16, 24, 33] on asymptotic statistical properties of least squares estimators in the case of finite discrete time sampling data.

A.1.1 The Constant Variance Case

As in the standard ordinary linear squares problem we assume the measurement process is described by the statistical model

$$Y_i = f(t_i, \theta_0) + \mathcal{E}_i, \quad i = 1, \ldots, n, \tag{A.1}$$

where $0 \leq t_1 < \cdots < t_n \leq T$ are the sampling times for the measurements, θ_0 is the nominal (true) parameter vector and $\mathcal{E}_i$ is a random variable satisfying

$$\mathcal{E}_i \sim \mathcal{N}\big(0, \sigma_0^2\big), \quad i = 1, \ldots, n,$$

with *constant* but unknown variance σ_0^2. Let $\hat{\Theta}^{(n)}$ denote the least squares estimator obtained with the observations (A.1) and define the $p \times p$ matrix Ω by

$$\Omega = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \nabla_\theta^{\mathsf{T}} f(t_i, \theta_0) \nabla_\theta f(t_i, \theta_0) = \lim_{n\to\infty} \frac{1}{n} \big(\chi(\theta_0)\big)^{\mathsf{T}} \chi(\theta_0), \tag{A.2}$$

where $\chi(\theta)$ represents the $n \times p$ sensitivity matrix

$$\chi(\theta) = \Big(\frac{\partial f(t_i, \theta)}{\partial \theta_j}\Big)_{i=1,\ldots,n,\, j=1,\ldots,p} = \begin{pmatrix} \nabla_\theta f(t_1, \theta) \\ \vdots \\ \nabla_\theta f(t_n, \theta) \end{pmatrix}.$$

For brevity we suppress notation indicating the dependence of χ on $t_1, \ldots, t_n$. We assume that Ω *exists* and is *nonsingular*. From the asymptotic theories of statistical analysis for least squares estimation, we have that under certain assumptions (see 12.22 and assumptions A1 – A9 in [33, Chapter 12]) the relation

$$\sqrt{n}\big(\hat{\Theta}^{(n)} - \theta_0\big) \sim \mathcal{N}_p\big(0, \sigma_0^2 \Omega^{-1}\big) \tag{A.3}$$

holds asymptotically as $n \to \infty$.

Relation (A.3) is essential when we define the *standard errors* which are the traditional tools in quantifying the accuracy of the parameter estimation. When n is sufficiently large, we can drop the limit sign in (A.2) and use the approximation

$$\Omega \approx \frac{1}{n} \chi(\theta_0)^{\mathsf{T}} \chi(\theta_0)$$

in (A.3) to obtain

$$\sqrt{n}\big(\hat{\Theta}^{(n)} - \theta_0\big) \sim \mathcal{N}_p\left(0, \sigma_0^2 n \big(\chi^{\mathsf{T}}(\theta_0)\chi(\theta_0)\big)^{-1}\right).$$

From the relation above we have

$$E\Big(\sqrt{n}\big(\hat{\Theta}^{(n)} - \theta_0\big)\Big) \approx 0,$$

which implies that that the estimator $\hat{\Theta}^{(n)}$ is asymptotically unbiased when $n \to \infty$, i.e.,

$$E\big(\hat{\Theta}^{(n)}\big) \approx \theta_0$$

and

$$\mathrm{Cov}\big(\sqrt{n}(\hat{\Theta}^{(n)} - \theta_0)\big) \approx \sigma_0^2 n \big(\chi^{\mathsf{T}}(\theta_0)\chi(\theta_0)\big)^{-1}$$

for n large, which gives an approximation for the covariance matrix of $\hat{\Theta}^{(n)}$

$$\mathrm{Cov}\big(\hat{\Theta}^{(n)}\big) \approx \sigma_0^2 \big(\chi^{\mathsf{T}}(\theta_0)\chi(\theta_0)\big)^{-1} := \Sigma_0.$$

The standard errors are simply the square roots of the diagonal elements of the covariance matrix Σ_0, i.e.,

$$SE_k = \sqrt{(\Sigma_0)_{kk}} = \sqrt{\sigma_0^2 (\chi^{\mathsf{T}}\chi)^{-1}_{kk}}, \qquad k = 1, \ldots, p.$$

A.1.2 The Time Dependent Variance Case

Assume now that we have noise with time dependent variance, i.e., $\mathcal{E}_i \sim \mathcal{N}\big(0, \sigma(t_i)\big)$ for $i = 1, \ldots, n$. The relation (A.3) becomes now

$$\sqrt{n}\big(\hat{\Theta}^{(n)} - \theta_0\big) \sim \mathcal{N}_p\big(0, \Omega_w^{-1}\big),$$

which holds asymptotically when $n \to \infty$. Here Ω_w is the weighted $p \times p$ matrix

$$\Omega_w = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \frac{1}{\sigma^2(t_i)} \nabla_\theta^\mathsf{T} f(t_i, \theta_0) \nabla_\theta f(t_i, \theta_0) = \lim_{n\to\infty} \frac{1}{n} \chi_w^\mathsf{T}(\theta_0) \chi_w(\theta_0),$$

assumed to exist and to be nonsingular, and $\chi_w(\theta)$ represents the weighted $n \times p$ sensitivity matrix

$$\chi_w(\theta) = \left(\frac{\partial f_w(t_i, \theta)}{\partial \theta_j} \right)_{i=1,\ldots,n,\ j=1,\ldots,p} = \begin{pmatrix} \nabla_\theta f_w(t_1, \theta) \\ \vdots \\ \nabla_\theta f_w(t_n, \theta) \end{pmatrix},$$

where we have set $f_w(t_i, \theta) = f(t_i, \theta)/\sigma(t_i)$. Using the same steps as in the constant variance case we obtain that the estimator $\hat{\Theta}^{(n)}$ is asymptotically unbiased when $n \to \infty$, i.e.,

$$E\big(\hat{\Theta}^{(n)}\big) \approx \theta_0,$$

with covariance matrix $\mathrm{Cov}\big(\hat{\Theta}^{(n)}\big)$ approximated by

$$\mathrm{Cov}\big(\hat{\Theta}^{(n)}\big) \approx \big(\chi_w^\mathsf{T}(\theta_0)\chi_w(\theta_0)\big)^{-1}. \tag{A.4}$$

for n large. Likewise, the standard errors in this case are given by

$$SE_k = \sqrt{\big(\chi_w^\mathsf{T}(\theta_0)\chi_w(\theta_0)\big)^{-1}_{kk}}, \quad k = 1, \ldots, p. \tag{A.5}$$

A.2 Linear Regression – Distributed Measurements

In this subsection we consider the case where the output model (1.2) is linear in θ, i.e., we have

$$\eta(t) = G(t)\theta, \quad 0 \le t \le T,\ \theta \in \mathbb{R}^p, \tag{A.6}$$

where

$$G(t) = \big(g_1(t), \ldots, g_p(t)\big) \in \mathbb{R}^{1\times p}, \quad 0 \le t \le T,$$

with functions $g_i : [0, T] \to \mathbb{R}$, $i = 1, \ldots, p$. We assume that the measurements are given by

$$y(t, \theta_0) = G(t)\theta_0 + \epsilon(t), \quad 0 \le t \le T, \tag{A.7}$$

where θ_0 is the true or nominal parameter vector and $\epsilon(\cdot)$ is a realization of the noise process $\mathcal{E}(\cdot)$ which is assumed to satisfy (2.2). The parameter estimation problem would require one to solve the equation $G(t)\theta = y(t, \theta_0)$, which in general does not have a solution because of the noise term in $y(\cdot, \theta_0)$. Therefore we consider the residual $y(\cdot, \theta_0) - G(\cdot)\theta$ and attempt to find the estimate $\hat{\theta} = \hat{\theta}(\theta_0)$ as an element in $\mathbb{R}^p$ which minimizes a weighted means square of this residual, i.e.,

$$\hat{\theta} = \operatorname*{argmin}_{\theta\in\mathbb{R}^p} \int_0^T \frac{1}{\sigma^2(t)} \big(y(t, \theta_0) - G(t)\theta\big)^2 dP(t). \tag{A.8}$$

We introduce the weighted L^2-space $L_w^2 = L_w^2(0, T)$ of all P-measurable functions $u : [0, T] \to \mathbb{R}$ such that u/σ is square integrable on $[0, T]$ with respect to the measure P and define the operator

$\mathcal{G}:\mathbb{R}^p \to L^2_w$ by the right side of (A.6) assuming that $g_i \in L^2_w$, $i=1,\dots,p$. Then we can write (A.8) as

$$\hat{\theta} = \operatorname*{argmin}_{\theta\in\mathbb{R}^p} \big\| y(\cdot,\theta_0) - \mathcal{G}\theta \big\|^2_{L^2_w}. \tag{A.9}$$

This is a standard minimum norm problem [28] in the Hilbert space L^2_w which always has solutions $\hat{\theta}$. These solutions satisfy the normal equations

$$\mathcal{G}^*\mathcal{G}\,\hat{\theta} = \mathcal{G}^* y(\cdot,\theta_0), \tag{A.10}$$

where $\mathcal{G}^* : L^2_w \to \mathbb{R}^p$ is the adjoint operator of $\mathcal{G}$ defined by $\langle \mathcal{G}\theta, u\rangle_{L^2_w} = \langle \theta, \mathcal{G}^* u\rangle_{\mathbb{R}^p}$ for all $\theta\in\mathbb{R}^p$ and all $u\in L^2_w$. By the general theory of *generalized* or *pseudo inverses* and minimum norm problems in Hilbert spaces [28, Chapter 6] one can under appropriate conditions obtain a unique $\hat{u}\in \operatorname{range}\mathcal{G}$ such that $\|y(\cdot,\theta_0)-\hat{u}\|_{L^2_w} = \min_{u\in L^2_w}\|y(\cdot,\theta_0)-u\|_{L^2_w}$. In order to guarantee existence of a unique $\hat{\theta}$ which solves (A.9), i.e., a unique $\hat{\theta}$ with $\hat{u}=\mathcal{G}\hat{\theta}$, we need injectivity of the linear map $\mathcal{G}$ or equivalently surjectivity of the adjoint map $\mathcal{G}^*$ (i.e., $\operatorname{range}\mathcal{G}^* = \mathbb{R}^p$) (guaranteed, of course, if the functions $\{g_i\}$ are linearly independent in L^2_w). This is in turn equivalent to the fact that the linear map $\mathcal{G}^*\mathcal{G}:\mathbb{R}^p\to\mathbb{R}^p$ is bijective and in this case we obtain from (A.10) the unique solution

$$\hat{\theta}(\theta_0) = \big(\mathcal{G}^*\mathcal{G}\big)^{-1}\mathcal{G}^* y(\cdot,\theta_0). \tag{A.11}$$

Without bijectivity, one often can obtain a unique solution with minimum norm (e.g., one typically minimizes over a compact set $\mathcal{AP}$ of admissible parameters θ thereby guaranteeing that $\operatorname{range}\mathcal{G}$ is closed in L^2_w – see [28]) by using the pseudoinverse $\mathcal{G}^\dagger$ in (A.11) in place of $\big(\mathcal{G}^*\mathcal{G}\big)^{-1}\mathcal{G}^*$. This also reveals that *identifiability* of the parameters θ, which in the linear case is a global property, is equivalent to $\operatorname{range}\mathcal{G}^* = \mathbb{R}^p$ or, equivalently, to invertibility of $\mathcal{G}^*\mathcal{G}$.

Using the representation $\big(\mathcal{G}\theta\big)(t) = G(t)\theta$ one can easily see that

$$\mathcal{G}^* u = \big(\langle g_1,u\rangle_{L^2_w},\dots,\langle g_p,u\rangle_{L^2_w}\big)^{\mathsf{T}}, \quad u\in L^2_w.$$

Taking $u=\mathcal{G}\theta = g_1\theta_1+\cdots+g_p\theta_p$ we obtain $\langle g_i,u\rangle_{L^2_w} = \big(\langle g_i,g_1\rangle_{L^2_w},\dots,\langle g_i,g_p\rangle_{L^2_w}\big)\theta$. From this we see that the matrix representation of the operator $\mathcal{G}^*\mathcal{G}$ is given by the $p\times p$-matrix

$$\big(\langle g_i,g_j\rangle_{L^2_w}\big)_{i,j=1,\dots,p} = \left(\int_0^T \frac{1}{\sigma^2(t)} g_i(t)g_j(t)\,dP(t)\right)_{i,j=1,\dots,p}. \tag{A.12}$$

This matrix is precisely the weighted Fisher information matrix for the linear case we are considering in this section (note that $\nabla_\theta G(t)\theta = G(t) = \big(g_1(t),\dots,g_p(t)\big)$, $0\le t\le T$, $\theta\in\mathbb{R}^p$).

From a statistical point of view, equation (A.11) relates the realizations $\hat{\theta}$ of the weighted least squares estimator $\hat{\Theta}$ to the measurements $y(\cdot,\theta_0)$ which is a realization of the stochastic process $Y(\cdot,\theta_0)$ given by

$$Y(t,\theta_0) = G(t)\theta_0 + \mathcal{E}(t), \quad 0\le t\le T. \tag{A.13}$$

An important question to be investigated is how the statistical properties of the estimator $\hat{\Theta}$ are related to the statistical properties of Y. The random variable equivalent of equation (A.11) is

$$\hat{\Theta}(\theta_0) = \big(\mathcal{G}^*\mathcal{G}\big)^{-1}\mathcal{G}^* Y(\cdot,\theta_0). \tag{A.14}$$

From this equation we obtain for the expected value

$$\begin{aligned} E\big(\hat{\Theta}(\theta_0)\big) &= E\big((\mathcal{G}^*\mathcal{G})^{-1}\mathcal{G}^* Y(\cdot,\theta_0)\big) = (\mathcal{G}^*\mathcal{G})^{-1}E\big(\mathcal{G}^* Y\big) \\ &= (\mathcal{G}^*\mathcal{G})^{-1}\mathcal{G}^*\mathcal{G}\theta_0 = \theta_0, \end{aligned} \tag{A.15}$$

where we have used $E\big(\mathcal{G}^*Y\big) = \mathcal{G}^*\mathcal{G}\theta_0$, which is a consequence of (A.13). Equation (A.15) reveals that the estimator $\hat{\Theta}(\theta_0)$ is unbiased for any $\theta_0 \in \mathbb{R}^p$. For the covariance matrix we obtain

$$\begin{aligned}\mathrm{Cov}(\hat{\Theta}(\theta_0)) &= \mathrm{Cov}\big(\big(\mathcal{G}^*\mathcal{G}\big)^{-1}\mathcal{G}^*Y(\cdot,\theta_0)\big) = \big(\mathcal{G}^*\mathcal{G}\big)^{-1}\,\mathrm{Cov}\big(\mathcal{G}^*Y(\cdot,\theta_0)\big)(\mathcal{G}^*\mathcal{G})^{-1} \\ &= \big(\mathcal{G}^*\mathcal{G}\big)^{-1}(\mathcal{G}^*\mathcal{G})(\mathcal{G}^*\mathcal{G})^{-1} = \big(\mathcal{G}^*\mathcal{G}\big)^{-1}, \quad \theta_0 \in \mathbb{R}^p.\end{aligned} \tag{A.16}$$

In this equation we used the fact that

$$\mathrm{Cov}\big(\mathcal{G}^*Y(\cdot,\theta_0)\big) = \mathcal{G}^*\mathcal{G}. \tag{A.17}$$

Since the proof of this equation is not entirely trivial, we present it in some detail below.

The (ij)-entry of the covariance matrix is

$$\Big(\mathrm{Cov}\big(G_w^*Y_w\big)\Big)_{ij} = E\Big(\int_0^T \frac{1}{\sigma^2(t)} g_i(t)\mathcal{E}(t)\,dP(t)\int_0^T \frac{1}{\sigma^2(s)} g_j(s)\mathcal{E}(s)\,dP(s)\Big).$$

In order to establish (A.17) we must argue that

$$E\Big(\int_0^T \frac{1}{\sigma^2(t)} g_i(t)\mathcal{E}(t)\,dP(t)\int_0^T \frac{1}{\sigma^2(s)} g_j(s)\mathcal{E}(s)\,dP(s)\Big) = \int_0^T \frac{1}{\sigma^2(t)} g_i(t)g_j(t)\,dP(t). \tag{A.18}$$

We consider two distinct cases of interest: (i) the case where the measure P is absolutely continuous (in view of the weighings $\frac{1}{\sigma(t)}$ we can without loss of generality simple choose P to be Lebesgue measure in this case) and (ii) P is a discrete measure with a finite number of atoms. We first consider the argument for P being Lebesgue measure.

From the independency assumption on $\mathcal{E}(t)$ we have that (see (2.2))

$$E\big(\mathcal{E}(t)\mathcal{E}(s)\big) = \int_\Omega \mathcal{E}(t)(\omega)\mathcal{E}(s)(\omega)\,d\omega = \sigma(t)\sigma(s)\delta(t-s).$$

Multiplying this equation above by $g_i(t)/\sigma^2(t)$, integrating with respect to t and changing the order of integration we obtain

$$E\Big(\int_0^T \frac{1}{\sigma^2(t)} g_i(t)\mathcal{E}(t)\,dt\mathcal{E}(s)\Big) = g_i(s).$$

If we multiply now this equation with $g_j(s)/\sigma^2(s)$, and repeat the previous steps one more time (of course, now integrating with respect to s), we obtain

$$E\Big(\int_0^T \frac{1}{\sigma^2(t)} g_i(t)\mathcal{E}(t)\,dt\int_0^T \frac{1}{\sigma^2(s)} g_j(s)\mathcal{E}(s)\,ds\Big) = \int_0^T \frac{1}{\sigma^2(s)} g_i(s)g_j(s)\,ds,$$

which proves (A.18) in the case of Lebesgue measure.

If instead of the Lebesgue measure we consider a discrete measure

$$P = \sum_{i=1}^N \delta_{t_i} \tag{A.19}$$

with atoms at $t_1 < t_2 < \cdots < t_n$ in $[0,T]$, from the left side of (A.18), we have

$$\begin{aligned} E\Big(\sum_{k=1}^{N} \frac{1}{\sigma^2(t_k)} g_i(t_k)\mathcal{E}(t_k) \sum_{l=1}^{N} \frac{1}{\sigma^2(t_l)} g_j(t_l)\mathcal{E}(t_l)\Big) & \\ = \sum_{k,l=1}^{N} \frac{1}{\sigma^2(t_k)} g_i(t_k) \frac{1}{\sigma^2(t_l)} g_j(t_l) E\big(\mathcal{E}(t_k)\mathcal{E}(t_l)\big) & \\ = \sum_{k,l=1}^{N} \frac{1}{\sigma^2(t_k)} g_i(t_k) \frac{1}{\sigma^2(t_l)} g_j(t_l) \sigma(t_k)\sigma(t_l)\delta_{kl} & \\ = \sum_{k=1}^{N} \frac{1}{\sigma^2(t_k)} g_i(t_k) g_j(t_k), & \end{aligned}$$

which is (A.18) for the discrete case.

The expected value and the covariance matrix of the parameter estimates $\hat{\theta}$ are important because they allow us to compute the confidence intervals for individual model parameters.

The expected value and the covariance matrix for the discrete time linear regression case determined by (A.19) is readily obtained from the continuous time case, if we observe that (see (A.12))

$$\mathcal{G}^*\mathcal{G} = \left(\sum_{k=1}^{N} \frac{1}{\sigma^2(t_k)} g_i(t_k) g_j(t_k)\right)_{i,j=1,\ldots,p}. \tag{A.20}$$

Moreover, the generalized weighted least square problem (A.8) becomes now a weighted least squares problem, i.e.,

$$\hat{\theta} = \operatorname*{argmin}_{\theta} \sum_{i=1}^{N} \frac{1}{\sigma^2(t_i)} \big(y(t_i) - (G\theta)(t_i)\big)^2,$$

whose solution $\hat{\theta}$ satisfies the corresponding discrete normal equations and has statistical properties which are the discrete equivalent of (A.15) and (A.16). Finally, we also note that for the linear regression case, the expected value (A.15) and the covariance matrix (A.16) are *global* with respect to θ_0, that is they do not depend on the particular choice of the true parameter vector θ_0.

We note that from the above discussions one may readily discern obvious connections between the needed well-conditioning of the discrete Fisher matrix (A.20), the requirements for $\mathcal{G}^*\mathcal{G}$ having full rank in the discussions for (A.11), the pseudoinverse theory for minimum norm solutions of least squares problems in Hilbert spaces presented in [28], and the asymptotic theory conditions involving nonsingularity of Ω in (A.2).

A.3 Nonlinear Regression – Distributed Measurements

Assume now that we have a single output nonlinear regression model

$$y(t) = f(t,\theta) + \varepsilon(t), \quad t \in [0,T],$$

with distributed time measurements in the interval $[0,T]$ where $y(t)$ and $\varepsilon(t)$ are realizations of the corresponding statistical model processes $Y(t)$ and $\mathcal{E}(t)$. The process $\mathcal{E}(t)$ is a noisy process which describes the measurement error at time t, and is assumed to have the same properties as in the linear case. We shall apply the results from Subsection A.2 to a linearized problem, thus obtaining only approximate statistical properties of the estimators in the nonlinear problem.

The residual function in this case is $y(t) - f(t,\theta)$, and the weighted least squares problem becomes

$$\hat{\theta} = \operatorname*{argmin}_{\theta} \int_0^T \frac{1}{\sigma^2(t)} \big(y(t) - f(t,\theta)\big)^2 \, dP(t). \tag{A.21}$$

Again, P is a measure (either continuous or discrete) which gives the distribution of the measurements in the interval $[0,T]$.

Similar to the analysis carried out for the linear case, in the following we will be interested in the statistical properties (i.e., mean and covariance matrix) of the estimator $\hat{\Theta}$. For the nonlinear regression problem we do not have the direct linear relationship (A.16) between the data and the estimated model parameters, and consequently we can no longer use the formulas (A.15) and (A.16). However, we will use a linearization approach to obtain approximations for the statistical properties of $\hat{\Theta}$. In a small neighborhood of θ_0, we have the linear Taylor expansion

$$f(t,\theta) \approx f(t,\theta_0) + \nabla_\theta f(t,\theta_0)(\theta - \theta_0).$$

Substituting this into the cost functional (A.21) we obtain

$$J(y,\theta) = \int_0^T \frac{1}{\sigma^2(t)} \big(y(t) - f(t,\theta)\big)^2 \, dP(t) \approx \int_0^T \frac{1}{\sigma^2(t)} \big(y(t) - f(t,\theta_0) - \nabla_\theta f(t,\theta_0)(\theta - \theta_0)\big)^2 \, dP(t).$$

We set $z(t) = y(t) - f(t,\theta_0)$ $(= \varepsilon(t))$, $\beta = \theta - \theta_0$ and obtain

$$J(y,\theta) \approx \tilde{J}(y,\beta) := \int_0^T \frac{1}{\sigma^2(t)} \big(z(t) - \nabla_\theta f(t,\theta_0)\beta\big)^2 \, dP(t).$$

Let $\hat{\beta}$ be a minimizer of $\tilde{J}(y,\beta)$. Then $\tilde{\theta} = \hat{\beta} + \theta_0$ is (hopefully) a good approximation for $\hat{\theta} = \operatorname{argmin} J(y,\theta)$. The big advantage of using the linearization of f around θ_0 is that it leads to a linear least squares problem for β. The solution $\hat{\beta}$ is given by the normal equations (A.10),

$$\hat{\beta} = \big(G_w^* G_w\big)^{-1} G_w^* z, \tag{A.22}$$

where the operator G_w now is given by

$$G_w = \nabla_\theta f_w(\cdot,\theta_0),$$

i.e., we have $g_j(t) = \big(\partial f/\partial\theta_j\big)(t,\theta_0)$, $j = 1,\dots,p$. As in Section A.2 we tacitly also assume here that $\operatorname{rank} G = p$, which implies that the $p \times p$ matrix $G_w^* G_w$ is invertible. From (A.22) we obtain

$$\hat{\theta} - \theta_0 \approx \big(G_w^* G_w\big)^{-1} G_w^* z.$$

The non-singular matrix $G_w^* G_w$ is the $p \times p$ weighted Fisher information matrix

$$F(T,\theta_0) = \int_0^T \frac{1}{\sigma^2(t)} \nabla_\theta{}^{\mathsf{T}} f(t,\theta_0) \nabla_\theta f(t,\theta_0) \, dP(t). \tag{A.23}$$

Applying the results obtained in Section A.2 for linear regression models with distributed measurements, we have the approximate relationships

$$E(\hat{\Theta}) = \theta_0,$$

and

$$\operatorname{Cov}(\hat{\Theta}) = F(T,\theta_0)^{-1}.$$

Thus we have

$$\hat{\Theta} \sim \mathcal{N}_p\big(\theta_0, F(T,\theta_0)^{-1}\big).$$

Acknowledgements: This research was supported in part by the U.S. Air Force Office of Scientific Research under grants AFOSR-FA9550-08-1-0147 and AFOSR-FA9550-09-1-0226 and in part by the National Institute of Allergy and Infectious Disease under grant NIAID 9R01AI071915-05.

References

[1] R. C. Aster, B. Borchers and C. H. Thurber, *Parameter Estimation and Inverse Problems,* Elsevier Academic Press, New York, NY, 2005.

[2] H. M. Adelman and R. T. Haftka, Sensitivity analysis of discrete structural systems, *A.I.A.A. Journal*, **24** (1986), 823 – 832.

[3] P. Bai, H. T. Banks, S. Dediu, A. Y. Govan, M. Last, A. Loyd, H. K. Nguyen, M. S. Olufsen, G. Rempala and B. D. Slenning, Stochastic and deterministic models for agricultural production networks, *Math. Biosci. and Engineering,* **4** (2007), 373 – 402.

[4] H. T. Banks and K. L. Bihari, Modeling and estimating uncertainty in parameter estimation, *Inverse Problems*, **17** (2001), 95 – 111.

[5] H. T. Banks and D. M. Bortz, Inverse problems for a class of measure dependent dynamical systems, CRSC-TR04-33, NCSU, September, 2004; *J. Inverse and Ill-posed Problems*, **13** (2005), 103 – 121.

[6] H. T. Banks, D. M. Bortz, G. A. Pinter and L. K. Potter, Modeling and imaging techniques with potential for application in bioterrorism, Chapter 6 in *Bioterrorism: Mathematical Modeling Applications in Homeland Security*, (H.T. Banks and C. Castillo-Chavez, eds.), Frontiers in Applied Mathematics **FR28**, SIAM, Philadelphia, 2003, pp. 129–154.

[7] H. T. Banks, M. Davidian, J. R. Samuels Jr., and K. L. Sutton, An inverse problem statistical methodology summary, CRSC-TR08-01, NCSU, January, 2008; Chapter 11 in *Statistical Estimation Approaches in Epidemiology*, (edited by Gerardo Chowell, Mac Hyman, Nick Hengartner, Luis M.A Bettencourt and Carlos Castillo-Chavez), Springer, Berlin Heidelberg New York, 2009, pp. 249–302.

[8] H. T. Banks, S. Dediu and S.L. Ernstberger, Sensitivity functions and their uses in inverse problems, *J. Inv. Ill-posed Problems*, **15** (2007), 683 – 708.

[9] H. T. Banks, S. Dediu and H. K. Nguyen, Sensitivity of dynamical systems to parameters in a convex subset of a topological vector space, *Math. Biosci. and Engineering*, **4** (2007), 403 – 430.

[10] H. T. Banks, S. L. Ernstberger and S. L. Grove, Standard errors and confidence intervals in inverse problems: sensitivity and associated pitfalls, *J. Inv. Ill-posed Problems*, **15**, (2006), 1 – 18.

[11] J. J. Batzel, F. Kappel, D. Schneditz and H. T. Tran, *Cardiovascular and Respiratory Systems: Modeling, Analysis and Control*, Frontiers in Applied Mathematics **FR34**, SIAM, Philadelphia, 2006.

[12] M. P. F. Berger and W. K. Wong (Editors), *Applied Optimal Designs*, John Wiley & Sons, Chichester, UK, 2005.

[13] P. Billingsley, *Convergence of Probability Measures*, John Wiley & Sons, New York, NY, 1968.

[14] G. Casella and R. L. Berger, *Statistical Inference,* Duxbury, California, 2002.

[15] J. B. Cruz, (Editor), *System Sensitivity Analysis*, Dowden, Hutchinson & Ross, Inc., Stroudsberg, PA, 1973.

[16] M. Davidian and D. Giltinan, *Nonlinear Models for Repeated Measurement Data*, Chapman & Hall, London, 1998.

[17] J. Dieudonné, *Foundations of Modern Analysis*, Academic Press, New York, 1960.

[18] M. Eslami, *Theory of Sensitivity in Dynamic Systems: An Introduction*, Springer-Verlag, Berlin, 1994.

[19] V. V. Fedorov, *Theory of Optimal Experiments*, Academic Press, New York and London, 1972.

[20] V. V. Fedorov and P. Hackel, *Model-Oriented Design of Experiments*, Springer-Verlag, New York, NY, 1997.

[21] P. M. Frank, *Introduction to System Sensitivity Theory*, Academic Press, New York, NY, 1978.

[22] A. R. Gallant, *Nonlinear Statistical Models*, John Wiley & Sons, Inc., New York, NY, 1987.

[23] P. J. Huber, *Robust Statistics*, John Wiley & Sons, Inc., New York, NY, 1981.

[24] R. I. Jennrich, Asymptotic properties of non-linear least squares estimators, *Ann. Math. Statist.*, **40** (1969), 633 – 643.

[25] F. Kappel and M. Munir, A New Approach to Optimal Design Problems, Proc. Internatl. Conf. Nonlinear Analysis and Optimization, October 6 – 10, 2008, Budva (Montenegro), to appear.

[26] M. Kleiber, H. Antunez, T. D. Hien and P. Kowalczyk, *Parameter Sensitivity in Nonlinear Mechanics: Theory and Finite Element Computations*, John Wiley & Sons, New York, NY, 1997.

[27] M. Kot, *Elements of Mathematical Ecology*, Cambridge University Press, Cambridge, 2001, p. 7 – 9.

[28] D. G. Luenberger, *Optimization by Vector Space Methods*, John Wiley & Sons, New York, NY, 1969.

[29] V. Paulauskas and A. Račkauskas, *Approximation Theory in the Central Limit Theorem*, Kluwer Academic Publ., Dordrecht, 1989.

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

[31] F. Pukelsheim, *Optimal Design of Experiments*, John Wiley & Sons, New York, NY, 1993.

[32] A. Saltelli, K. Chan and E. M. Scott, eds., *Sensitivity Analysis*, Wiley Series in Probability and Statistics, John Wiley & Sons, New York, NY, 2000.

[33] G. A. F. Seber and C. J. Wild, *Nonlinear Regression,* John Wiley & Sons, New York, NY, 1989.

[34] K. Thomaseth and C. Cobelli, Generalized sensitivity functions in physiological system identification, *Ann. Biomed. Eng.*, **27**(5) (1999), 607 – 616.