



FACTOR MAPPING AND METAMODELLING

Marco Ratto, Andrea Pagano and Peter Young



EUR 21878 EN - 2007

The mission of the JRC is to provide customer-driven scientific and technical support for the conception, development, implementation and monitoring of EU policies. As a service of the European Commission, the JRC functions as a reference centre of science and technology for the Union. Close to the policy-making process, it serves the common interest of the Member States, while being independent of special interests, whether private or national.

European Commission
Joint Research Centre

Contact information

Address: Joint Research Centre
via E. Fermi
21020 Ispra (VA) Italy
E-mail: marco.ratto@jrc.it
Tel.: +39 0332 78 5639
Fax: +39 0332 78 5733

<http://www.jrc.ec.europa.eu>

Legal Notice

Neither the European Commission nor any person acting on behalf of the Commission is responsible for the use which might be made of this publication.

A great deal of additional information on the European Union is available on the Internet. It can be accessed through the Europa server
<http://europa.eu/>

JRC37692

EUR 21878 EN
ISSN 1018-5593

Luxembourg: Office for Official Publications of the European Communities

© European Communities, 2007

Reproduction is authorised provided the source is acknowledged
Printed in Italy

Factor Mapping and Metamodelling

Marco Ratto, Andrea Pagano

European Commission

Joint Research Centre, Ispra (VA), ITALY

Peter Young

Centre for Research on Environmental Systems and Statistics,

Lancaster University, Lancaster, U.K.

and Centre for Resource and Environmental Studies,

Australian National University, Canberra

May, 2007

Abstract

In this work we present some techniques, within the realm of Global Sensitivity Analysis, which permit to address fundamental questions in term of model's understanding.

In particular we are interested in developing tools which allow to determine which factor (or group of factors) are most responsible for producing model outputs Y within or outside specified bounds ranking the importance of the various input factors in terms of their influence on the variation of Y .

On the other hand, we look for representing in a direct way (graphically, analytically, etc.) the relationship between input factors $\{X_1, \dots, X_k\}$ and output Y in order to get a better understanding of the model itself.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 2 |
| 2 | Monte Carlo Filtering (MCF) | 3 |
| 2.1 | Implementation of Monte Carlo Filtering | 3 |
| 2.2 | Pro's and con's | 5 |
| 2.3 | Examples | 6 |
| 2.3.1 | Stability analysis of a controlled chemical reactor | 6 |
| 2.3.2 | Stability analysis of a small macro-economic model | 12 |
| 2.3.3 | Mapping propagation of the infection in the simple infection dynamics model | 14 |
| 2.4 | Exercises | 17 |
| 2.5 | Solutions | 18 |
| 3 | Metamodelling and the High-Dimensional Model Representation | 28 |
| 3.1 | Estimating HDMR's and meta-models | 29 |
| 3.2 | Examples | 37 |
| 3.2.1 | A simple example | 37 |
| 3.2.2 | Another simple example | 41 |
| 3.3 | Exercises | 44 |
| 3.4 | Solutions | 44 |
| 4 | Conclusions | 49 |

Chapter 1

Introduction

In the present work we aim to address (as many modellers often do) questions such as: ‘Which factor or group of factors are most responsible for producing model outputs within or outside specified bounds? Which parameters determine uniqueness or instability or runaway conditions in a dynamic model?’ For example, if Y were a dose of contaminant, we might be interested in how much (how often) a threshold level for this contaminant is being exceeded; or Y might have to fulfil a set of constraints, based on the information available on observed systems. The latter situation is typical in calibration.

Another typical question is whether it is possible to represent in a direct way (graphically, analytically, etc.) the relationship between input factors and output $Y = f(X_1, \dots, X_k)$. Computing Y usually requires solving systems of non-linear differential equations and the relationship $f(\cdot)$ can only be evaluated numerically, its form remaining unknown. Sensitivity analysis techniques allow for ranking the importance of the various input factors in terms of influence on the variation of Y . In addition, some sort of direct representation of $Y = f(X_1, \dots, X_k)$ would make the model’s properties even more transparent.

Hence, we will discuss techniques that can help to provide answers to such questions. We assign all these methods to the Factor Mapping setting, in which specific points/portions of the model output realizations, or even the entire domain, are mapped backwards onto the space of the input factors.

Chapter 2

Monte Carlo Filtering (MCF)

Let us first consider the case where the analyst is interested in targeted portions (extreme values, ceilings, thresholds, etc.) of the space of Y -realizations. In this situation, it is natural to partition the model realizations into ‘good’ and ‘bad’. This leads very naturally to Monte Carlo Filtering (MCF), in which one runs a Monte Carlo experiment producing realizations of the output(s) of interest corresponding to different sampled points in the input factors’ space. Having done this, one ‘filters’ the realizations, i.e. the elements of the Monte Carlo sample that fall within the ‘good’ realization are flagged as ‘behavioural’, while the remaining ones are flagged as ‘non-behavioural’. Regionalized Sensitivity Analysis (RSA, see Young et al., 1978; Hornberger and Spear, 1981; Spear et al., 1994; Young et al. 1996; Young 1999 and references cited therein) is an MCF procedure that aims to identify which factors are most important in leading to realizations of Y that are either in the ‘behavioural’ or ‘non-behavioural’ regions. In typical cases, RSA can answer this question by examining, for each factor, the subsets corresponding to ‘behavioural’ and ‘non-behavioural’ realizations. It is intuitive that, if the two subsets are dissimilar to one another (as well as, one would expect, to the initial marginal distribution of the factor), then that factor is influential.

2.1 Implementation of Monte Carlo Filtering

In Monte Carlo filtering a multiparameter Monte Carlo simulation is performed, sampling model parameters (X_1, \dots, X_k) from prior ranges and propagating parameter values through the model. Then, based on a set of constraints targeting the desired characteristics, a categorization is defined for each MC model realization, as either within or outside the target region. The terms behavioural (B) or non-behavioural (\bar{B}) are current in the literature.

The $[B - \bar{B}]$ categorization is mapped back onto the input’s structural parameters, each of which is thus also partitioned into a B and \bar{B} subsample. Given a full set of N Monte Carlo runs, one obtains two subsets: $(X_i|B)$ of size n and $(X_i|\bar{B})$ of size \bar{n} , where $n + \bar{n} = N$. In general, the two sub-samples will come from different unknown probability density functions, $f_n(X_i|B)$ and $f_{\bar{n}}(X_i|\bar{B})$.

In order to identify the parameters that are most responsible for driving the model into the target behaviour, the distributions f_n and $f_{\bar{n}}$ are compared for each parameter. If for a given parameter X_i the two distributions are significantly different, then X_i is a key factor in driving the model’s behaviour and there will be clearly identifiable subsets of values in its predefined range that are more likely to fall under B than under \bar{B} . If the two distributions are not significantly different, then X_i is unimportant and any value in its predefined range

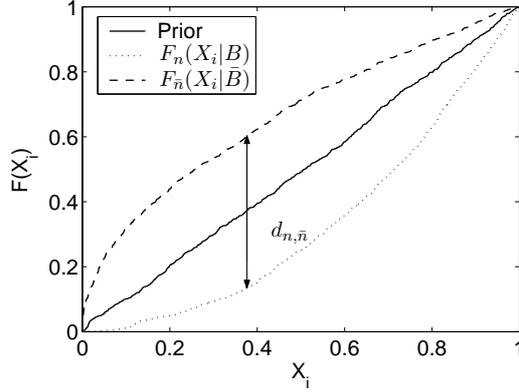


Figure 2.1: Graphical representation of the Smirnov test for an important parameter X_i . The vertical bar is the $d_{n,\bar{n}}$ statistic. Model realisations under the B category are more likely when X_i falls on the right of its pre-defined range.

is likely to fall into either B or \bar{B} .

This comparison can be made by applying standard statistical tests, such as the Smirnov two-sample test (two-sided version). In the Smirnov test the $d_{n,\bar{n}}$ statistic is defined for the cumulative distribution functions of X_i by

$$d_{n,\bar{n}}(X_i) = \sup \|F_n(X_i|B) - F_{\bar{n}}(X_i|\bar{B})\|$$

and the question answered by the test is: ‘At what significance level α does the computed value of $d_{n,\bar{n}}$ determine the rejection of the null hypothesis $f_n(X_i|B) = f_{\bar{n}}(X_i|\bar{B})$?’

The smaller α (or equivalently the larger $d_{n,\bar{n}}$), the more important the parameter is in driving the behaviour of the model. The procedure is exemplified in Figure 2.1 for a parameter X_i , uniformly distributed in the range $(0, 1)$ and displaying a significant difference between the B and \bar{B} subsets. In order to identify the portion of X_i values more likely to fall under B , the shape of the cumulative distribution $F_n(X_i|B)$ has to be examined. The latter is steeper on the right-hand side of the graph (or, equivalently, it has shifted downwards), indicating that X_i values on the right of its predefined range are more likely to produce a behavioural model realization.

The B and \bar{B} subsets can be further probed through bi-dimensional projections, in order to detect significant patterns. The standard procedure consists of computing the correlation coefficients ρ_{ij} between all parameters under the B or \bar{B} subsets, and plotting the bi-dimensional projections of the sample for the couples having $|\rho_{ij}|$ larger than a significance threshold. This usually makes it possible to ‘visualize’ relationships between parameters.

For example, let us consider a simple model given by the equation $Y = X_1 + X_2$, with $X_i \in (0, 1)$. Let us define the model’s target behaviour as $Y > 1$. Then, an MCF procedure can identify a significant negative correlation between X_1 and X_2 in the B subset, and the corresponding triangular pattern can be visualized through the projection of the B sample shown in Figure 2.2. From this pattern, one can deduce a constraint $X_1 + X_2 > 1$ to fulfil the target behaviour.

The same procedure can evidently be applied in more typical cases where the constraint on the factors is not evident from the form of the mathematical model, i.e. when the model is a computer code.

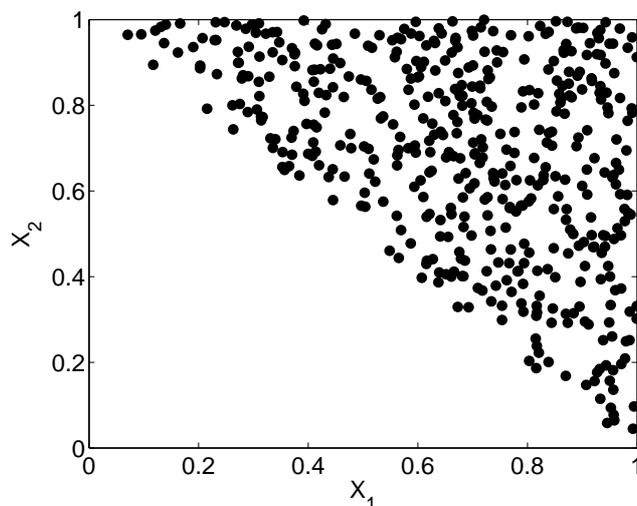


Figure 2.2: Bi-dimensional projection of the B sub-sample for the simple model $Y = X_1 + X_2$ (target behaviour $Y > 1$). The triangular pattern clearly indicates the relationship $X_1 + X_2 > 1$ for the target behaviour. The correlation coefficient between X_1 and X_2 in the MC sample is $\rho_{12} = -0.45$.

2.2 Pro's and con's

Similarly to variance based methods Regionalized Sensitivity Analysis has many global properties: (a) the whole range of values of the input factors is considered, and (b) all factors are varied at the same time. Smirnov analysis considers univariate marginal distributions and it relates not only to main effects of variance-based methods, but can also highlight certain types of interaction effects (see the Exercises below). Higher-order analysis can only be performed for two-way interactions that are detectable through correlation analysis, but no procedure is provided for more complex interaction structure. Spear et al. (1994), reviewing their experience with RSA, highlighted two key drawbacks:

1. A low success rate: practice has shown that the fraction of B is barely larger than 5% over the total simulations for large models (with number of factors $k > 20$), implying a lack of statistical power;
2. Difficulty in uncovering correlation and interaction structures of the B subset (see also Beck's review, 1987):
 - the Smirnov test is sufficient to ascertain whether a factor under analysis is important. However, it does not provide a necessary condition for importance, i.e. its non-significance does not ensure that a factor is non-influential;
 - many types of interaction structures induced by the classification are not detected by the univariate $d_{n,\bar{n}}$ statistic: e.g. factors combined as products or quotients may compensate (see Exercise 2.4.6 below, for $c = 0$);
 - the interaction structure is often far too complex for correlation analysis to be effective, i.e. bivariate correlation analysis is not revealing in many cases (see Saltelli et al., 2004, Example 2, pp. 159-161).

Such characteristics of RSA imply that no complete assessment can be performed with RSA, since for those factors proving unimportant in the Smirnov test, further inspection is needed (e.g. by applying other global SA tools) to verify that they are not involved in interactions. Only after this subsequent inspection can the relevance of an input factor be fully assessed. In order to address these limitations of RSA and to better understand the impact of uncertainty and interaction in the high-dimensional parameter spaces of models, Spear et al. (1994) developed the computer intensive Tree-Structured Density Estimation technique (TSDE), which allows for the characterization of complex interactions in that portion of the parameter space which gives rise to successful simulations. In TSDE, the B sub-sample is analysed by clustering regions of input factors characterized by high point density. This is based on a sequence of recursive binary splits of the B sample into two sub-domains (similarly to peaks and tails of histograms) of complementary characteristics:

- small regions of relatively high density;
- larger sparsely populated regions.

The TSDE procedure relies on the assumption that any non-random density pattern indicates an influence of input factors on the model output. Interesting applications of TSDE in environmental sciences can be found in Spear (1997), Grieb et al. (1999) and Pappenberger et al. (2006). In the latter reference it is shown how factor mapping can be used to identify areas of desirable and undesirable model behaviour, which aids the modelling process. Helton et al. (2006) also provide mapping techniques on scatterplots, based on the same assumptions as TSDE.

Our experience suggests that such extended RSA techniques for mapping B subsets can be revealing *when they work*, but, in spite of their higher coding and computational complexity, they may still be characterized by lack of statistical power in discriminating the significance of such density patterns in a robust manner. Hence, we limit our discussion here to the Smirnov test and correlation analysis, due to the simplicity of their implementation and their ease of interpretation. Later on we will present more powerful mapping techniques when discussing metamodelling.

2.3 Examples

We now show a few examples of the use of MCF to characterize the stability behaviour of dynamic models. We will give an example for a chemical reactor (continuous time model) and for a macro-economic model (discrete time model). Finally, we will also consider, in the light of MCF techniques, an infection dynamics model.

2.3.1 Stability analysis of a controlled chemical reactor

Let us consider a continuous stirred tank reactor (CSTR) where a first-order exothermic irreversible reaction $A \rightarrow B$ takes place in liquid phase.

The behaviour of the CSTR can be expressed in terms of mass and heat (enthalpy) balance equations. Mass balance reads as:

$$V \frac{dC_A}{dt} = Q(C_{A0} - C_A) - k(T)C_A \cdot V \quad (2.3.1)$$

where \bar{t} is time [s], V is the reactor volume [m³]; C_A is the concentration of A in the reactor and at its outlet [kmol/m³]; C_{A0} is the concentration of A at the inlet; Q is the volumetric flow rate [m³/s] at the input and output of the reactor, $k(T)$ is the kinetic ‘constant’ [1/s] of the first-order chemical reaction $A \rightarrow B$, which is expressed as a function of temperature T [K]:

$$k(T) = k_0 \exp\left(-\frac{E}{RT}\right),$$

where k_0 is the Arrhenius factor [1/s], E is the activation energy [kJ/kmol] and R is the gas constant [kJ/(kmol K)].

The mass balance equation (2.3.1) tells us that the rate of change of the amount of reactant A , given by the left-hand side term $V \frac{dC_A}{dt}$, equals the flow of A at the inlet, $Q \cdot C_{A0}$, minus the flow of A at the outlet, $Q \cdot C_A$, and minus the amount of A that is transformed into B per time unit, $k(T) \cdot C_A V$.

The heat balance reads as:

$$V \rho c_p \frac{dT}{dt} = Q \rho c_p (T_0 - T) - \Delta H_r k(T) C_A V - UA(T - T_c) \quad (2.3.2)$$

where ρ is the density of the reacting mixture [kg/m³], c_p is the specific heat of the mixture [kJ/(kg K)], T_0 is the temperature of liquid entering the reactor [K], T is the temperature of the liquid in the reactor and at its outlet, T_c is the temperature of the reactor’s coolant, $(-\Delta H_r)$ is the reaction enthalpy [kJ/kmol], U is the overall heat transfer coefficient between the inside of the reactor and the coolant [kJ/(s m² K)], A is the heat transfer area [m²].

The heat balance equation (2.3.2) tells us that the rate of change of enthalpy in the reactor, $V \rho c_p \frac{dT}{dt}$, equals the flow of enthalpy at the inlet, $Q \rho c_p T_0$, minus the flow of enthalpy at the outlet, $Q \rho c_p T$, plus the heat generated by the exothermic reaction, $(-\Delta H_r) k(T) C_A V$, minus the heat removed by the coolant, $UA(T - T_c)$.

The CSTR is controlled, in order to keep the temperature, and the associated quantity of product B , at the desired set-point T_s . The controlled variable is the temperature T and the manipulated variable is the coolant temperature T_c . The controller follows a standard proportional-integral (PI) design, implying the following control rule for the coolant temperature:

$$(T_c - T_{cs}) = -k_P(T - T_s) - k_I \int_0^t (T - T_s) d\bar{t} \quad (2.3.3)$$

where T_{cs} is the coolant temperature at the set-point, k_P is the proportional control gain and k_I is the integral control gain [1/s]. The control rule (2.3.3) tells us that the coolant ‘error’ $(T_c - T_{cs})$ is proportionally adjusted, with opposite sign, with respect to the temperature error $(T - T_s)$ and with respect to the integral of the temperature error. In other words, the coolant temperature will go down as the temperature error $(T - T_s)$ and its integral go up. The control gains k_P and k_I tell us by *how much* the coolant temperature is adjusted by the controller for a given temperature error and integral of the error: the higher the gains, the greater the change in the coolant. Finally, the integral action assures that the desired set-point is the unique steady state of the reactor.

Economists or econometricians will be familiar with Taylor rules: in monetary policy, for example, the controlled variable is inflation and the manipulated variable is the nominal interest rate. When the economy is ‘over-heated’ (i.e. with high inflation), the Central Bank *increases* the nominal interest rate to ‘cool’ the economy. In the case of the CSTR, the PI regulator *decreases* coolant temperature to cool an over-heated reactor.

The controlled CSTR is therefore described by the following system of differential equations, in dimensionless form:

$$\frac{d\xi}{dt} = -\xi + Da(\vartheta)(1 - \xi) \quad (2.3.4)$$

$$\frac{d\vartheta}{dt} = \vartheta_0 + N[\vartheta_{cs} - k_P(\vartheta - \vartheta_s) - k_I\tau\varphi] - (1 + N)\vartheta + Da(\vartheta)(1 - \xi) \quad (2.3.5)$$

$$\frac{d\varphi}{dt} = \vartheta - \vartheta_s \quad (2.3.6)$$

where (notation of Pellegrini and Biardi, 1990)

$$\begin{aligned} \xi &= \frac{C_{A0} - C_A}{C_{A0}} && \text{conversion of } A \\ \tau &= \frac{V}{Q} && \text{residence time [s]} \\ t &= \frac{\bar{t}}{\tau} && \text{dimensionless time} \\ N &= \frac{UA}{Q\rho c_p} && \text{dimensionless heat transfer coefficient} \\ \vartheta &= \frac{T - T_s}{\Delta T_a} && \text{dimensionless temperature} \\ \Delta T_a &= \frac{-\Delta H_r C_{A0}}{\rho c_p} && \text{adiabatic temperature rise [K]} \\ Da(\vartheta) &= k_0\tau \exp\left(-\frac{E}{R(T_s + \Delta T_a\vartheta)}\right) && \text{Damkoehler number} \\ \varphi &= \int_0^t (\vartheta - \vartheta_s) dt && \text{dimensionless integral of the error} \end{aligned}$$

The adiabatic temperature rise indicates the temperature increase that would be caused in the reactor if the entire amount of input A were converted into B under adiabatic conditions (i.e. without any heat exchange). The Damkoehler number indicates the average number of ‘reaction events’ that occur during the residence time.

The behaviour of this dynamical system, which can present complex dynamic features, from instability of the steady state to chaotic behaviour, has been intensively studied (see Pellegrini and Biardi, 1990; Giona and Paladino 1994; Paladino et al., 1995; Paladino and Ratto, 2000) .

Our aim in the present example is to study the stability conditions of the controlled CSTR. The local stability analysis of the reactor in the neighbourhood of the unique steady state $(\xi_s, \vartheta_s, \varphi_s)$ is performed by analysing the Jacobian:

$$J = \begin{pmatrix} -(1 + a_{21}) & a_{12} & 0 \\ -a_{21} & -(k_P N + N + 1) + a_{12} & -k_I\tau \\ 0 & 1 & 0 \end{pmatrix} \quad (2.3.7)$$

where

$$a_{12} = \xi_s \frac{E\Delta T_a}{R(T_s + \Delta T_a \vartheta_s)^2} > 0 \quad (2.3.8)$$

$$a_{21} = Da(\vartheta_s) > 0 \quad (2.3.9)$$

and

$$\xi_s = \frac{Da(\vartheta_s)}{1 + Da(\vartheta_s)}.$$

The steady state is stable if all three eigenvalues of J have negative real parts. This assures that, as the operating conditions are moved away from the set-point (i.e. the steady state), the reactor will return to steady state. Many authors have demonstrated that this system presents a Hopf bifurcation locus. At the Hopf locus the steady state becomes unstable and the dynamic behaviour of the reactor is characterized by persistent oscillations (limit cycle). This is, of course, unacceptable and must be avoided.

Although the Hopf bifurcation locus can be computed analytically (Giona and Paladino, 1994), here we analyse the stability conditions by applying the MCF techniques. This will allow us to confront the results of the MCF analysis with the analytic results. The problem can be formalized in the MCF framework as follows:

- the input factors are the control gains and the uncertain physico-chemical parameters of the CSTR model;
- the outputs are the eigenvalues of the Jacobian;
- the filtering criterion is:
 - behaviour B if all eigenvalues have negative real parts;
 - non-behaviour \bar{B} otherwise.

The nominal conditions of the CSTR are defined as follows:

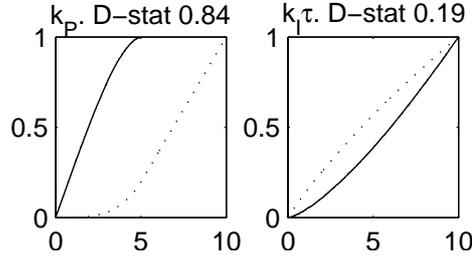
$$\begin{array}{ccccc} k_0 = 133600s^{-1} & E/R = 8000K & \tau = 3600s & \Delta T_a = 200K & T_0 = 298.42K \\ T_s = 430K & T_{cs} = 373.16 & \vartheta_s = 0 & \xi_s = 0.8 & N = 0.5 \end{array}$$

Such nominal conditions, depending on the values of the various physico-chemical parameters, are subject to a degree of uncertainty.

Let us first analyse the stability of this system under the nominal conditions, by varying only the control gains k_P and $k_I\tau$. As anticipated, this analysis can be performed analytically. We analyse here the Hopf bifurcation locus in the $(k_P, k_I\tau)$ plane using the MCF approach. The analysis requires the following steps:

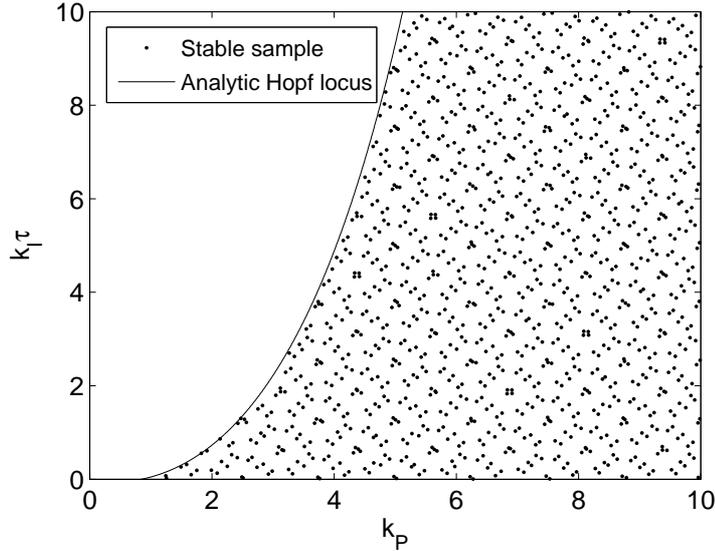
- Sample the control gains uniformly in the range $[0, 10]$;
- Compute the eigenvalues of the Jacobian;
- Check the stability condition;
 - the set of control gains providing stable eigenvalues (negative real parts) is the behavioural set;
 - the set of control gains providing unstable eigenvalues (non-negative real parts) is the non-behavioural set;

- Perform the Smirnov analysis;



Looking at the plot for k_P , we can see that the non-behavioural cumulative distribution (solid line) has a limit threshold at about $k_p^* = 5$, above which only stable solutions are present; this implies that sufficiently large values of the proportional control gain ($k_P > k_p^*$) are able to stabilize the reactor, whatever the value of the integral control gain. The latter gain, on the other hand, is more likely to produce a stable reactor for small values. However, the two cumulative distributions for $k_I\tau$ have the same support, i.e. both ranges of the stable and unstable sets span the entire support $[0, 10]$, implying that no clear threshold of stability can be identified.

- Plot the the behavioural sample on the $(k_P, k_I\tau)$ plane.



This shows the boundary of stability that exactly corresponds to the analytic solution (solid line). Note also that the patterns in the scatterplot are due to the Sobol' quasi-random sequences used for the example.

We now check the stability analysis for robustness, by allowing physico-chemical parameters to be uncertain. These uncertainties are given by normal distributions, with the following characteristics:

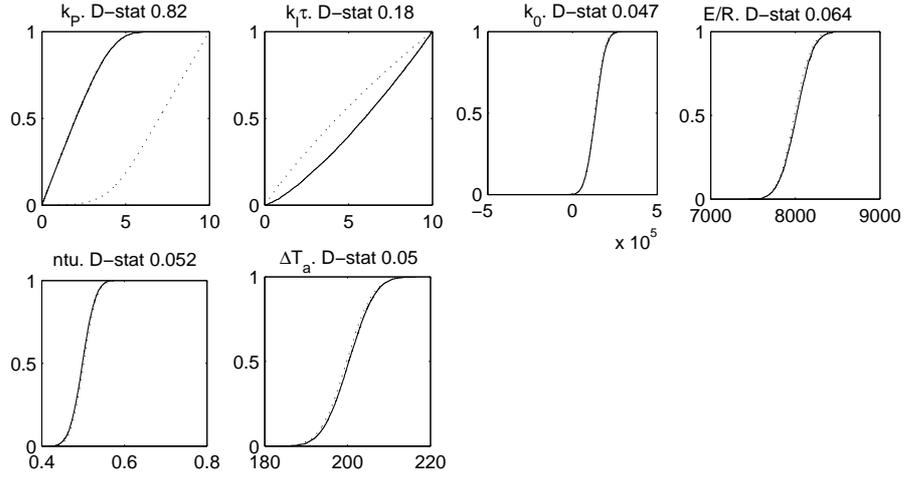
$(\mathbf{k}_0, E/R)$: assuming an estimate of kinetic coefficients, they are likely to be strongly correlated (Bard, 1974) and with much greater uncertainty for k_0 than for E/R . So we take a standard deviation of 35% for k_0 , 2% for E/R and a correlation coefficient of 0.96 (Paladino and Ratto, 2000).

(N) : the heat transfer coefficient has a standard deviation of 5%, i.e. $\sim N(0.5, 0.025)$.

(ΔT_a) : the adiabatic temperature difference has a standard deviation of 5 K, i.e. $\sim N(200, 5)$

We then proceed with the MCF analysis, as described below.

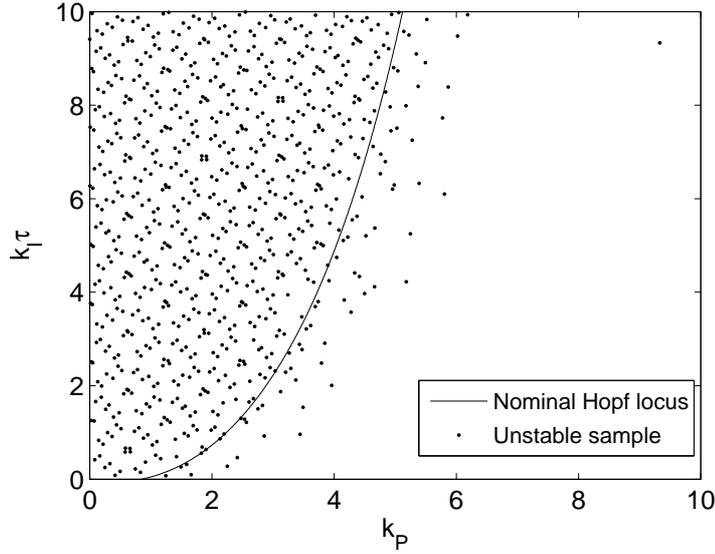
- First we perform the Smirnov test.



The sensitivity behaviour is still dominated by the control gains. In particular, sufficiently large values of k_P are still able to stabilize the reactor, whatever the values of $k_I \tau$ and all the uncertainties in the physico-chemical parameters. This is very important, since it allows us to design a stable control, i.e. robust against uncertainties. Looking very carefully at the physico-chemical parameters, we can also see that, as expected, the reactor has a very slight tendency towards instability if:

- the heat transfer coefficient decreases;
- the adiabatic temperature difference increases;
- the kinetic parameters increase.

- We then plot the bi-dimensional projection of the unstable sample \bar{B} onto the $(k_P, k_I \tau)$ plane.



This shows that unstable behaviour can be present beyond the Hopf locus computed at the nominal values of the physico-chemical parameters, implying that safe values for control gains have to be chosen according to uncertainty considerations. Constraining gains as $k_P > 6$ and $0 < k_I\tau < 4$ may be a good starting point.

2.3.2 Stability analysis of a small macro-economic model

Let us consider a simple macro-economic model: a Phillips curve. As usual in economics, this is a discrete-time dynamic model (typically with a quarterly sampling interval). Let c_t and π_t denote output gap and inflation, respectively. In macro-economic theory, the output gap denotes the cyclical component of GDP (gross domestic product) with respect to the long-term trend. In general terms, it is a quantity linked to the business cycle. We can write the hybrid Phillips curve as:

$$\pi_t = \omega_b \pi_{t-1} + \omega_f E_t \pi_{t+1} + \beta c_t + a_{\pi t} \tag{2.3.10}$$

$$c_t = 2A \cos(2\pi/\tau) c_{t-1} - A^2 c_{t-2} + a_{c,t} \tag{2.3.11}$$

where E_t denotes the expectation taken at period t , $0 < (\omega_b, \omega_f) < 1$, A and τ are the amplitude and period of c_t and $a_{\pi,t}, a_{c,t}$ are white noises.

The Phillips curve links the inflation dynamics to the output gap, in such a way that periods of economic expansion (i.e. with a positive output gap) are typically associated with an increase in inflation and vice versa. Moreover, the hybrid Phillips curve also says that inflation in the current period is linked with some persistence ω_b to the rate of inflation in the previous period and to the expected level of inflation in the following period, with a weight ω_f . This leads on to the theory of rational expectation behaviour of economic agents. In contrast to standard (physical) dynamic systems, the occurrence of a unique, stable solution in macro-economic rational expectations models requires that there be an equal number of explosive eigenvalues and forward-looking variables. In discrete-time dynamic models, stable roots have absolute values less than 1, while explosive ones are larger than 1. In this case, the Phillips curve has one lag π_{t-1} and one lead $E_t \pi_{t+1}$, so we need exactly one stable and one explosive eigenvalue. To help explain the stability conditions of economic rational expectations models to non-economists, we can say that the fact that the current level of inflation

depends on both past and future levels, makes the system like a two-point boundary system (similarly to certain types of differential equations in space describing advection-dispersion mechanisms). Hence, this implies the presence of initial and terminal conditions, corresponding to backward-looking and forward-looking components, respectively. Likewise in physical systems, backward-looking behaviour propagates the initial conditions into the future. This propagation is stable if it is associated to stable eigenvalues, thus assuring that the dynamic system will asymptotically converge to the steady state. The forward-looking components, on the other hand, propagate the terminal conditions into the past, i.e. in a symmetrical manner, reversing the orientation of the time axis. It intuitively makes sense that, if the orientation of time is reversed, explosive roots looking ‘towards the future’ become stable roots looking ‘towards the past’; therefore, in order to assure stability of the propagation into the past of forward-looking components, such components need to be associated to explosive roots.

The eigenvalues of this simple model can be computed analytically:

$$\begin{aligned}
 r_b &= \frac{1/\omega_f - \sqrt{1/\omega_f^2 - 4\omega_b/\omega_f}}{2} \\
 &= \frac{1 - \sqrt{1 - 4\omega_b\omega_f}}{2\omega_f} \\
 r_f &= \frac{1 + \sqrt{1 - 4\omega_b\omega_f}}{2\omega_f}.
 \end{aligned}
 \tag{2.3.12}$$

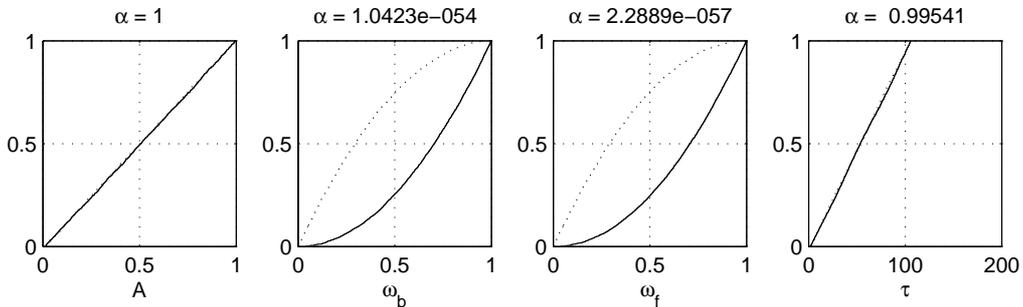
The restriction $1 - 4\omega_b\omega_f \geq 0$ ensures that the roots are real. The stability condition is verified if

$$\begin{aligned}
 \omega_f &< (1 - \omega_b) \\
 &\text{or} \\
 \omega_f &= (1 - \omega_b) \quad \text{and} \quad \omega_b > 0.5.
 \end{aligned}
 \tag{2.3.13}$$

We apply the Monte Carlo filtering technique to identify the stable behaviour. The support for the model coefficients is defined as:

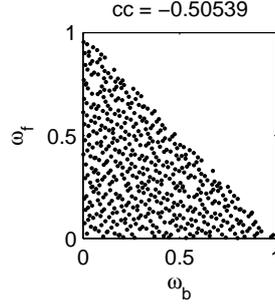
$$A \sim U[0, 1], \quad \omega_b \sim U[0, 1], \quad \omega_f \sim U[0, 1], \quad \tau \sim U[0, 100].$$

We first perform the Smirnov test for the separation of the B and \bar{B} subsets.



The occurrence of stable or unstable behaviour is clearly attributable to the coefficients ω_b and ω_f . Moreover, the shape of the dotted curves (corresponding to B) indicates that stable behaviour will more probably occur for smaller ω_b and ω_f values.

We subsequently perform a correlation analysis of the stable sample and plot the significant correlation selected: the bi-dimensional projection of the B subset onto the (ω_b, ω_f) plane.



The dots in this plot clearly indicate the first stability condition $\omega_f < (1 - \omega_b)$ in (2.3.13). The second condition in (2.3.13) is just one limit case, and tells us simply that only the half part $\omega_b > 0.5$ of the stability boundary line $\omega_f = (1 - \omega_b)$ provides stable behaviour. This is, of course, hardly visible in the plot.

2.3.3 Mapping propagation of the infection in the simple infection dynamics model

We consider an infective process at its early stage, where I is the number of infected individuals at time t and S is the number of individuals susceptible to infection at time t .

We assume that the infection is propagated through some kind of contact between individuals who, especially at the early stage, do not take any precautions to avoid the contagion.

It is reasonable to assume that the number of contacts per unit time is proportional to the number of individuals in each group (i.e. to $I \times S$) via a contact coefficient $k < 1$.

Also, the number of infections is proportional to the number of contacts through an 'infection coefficient' ($\gamma < 1$), which is the likelihood that the infection is passed on during a given contact.

Depending on the dangerousness of the infection, the infected individuals will end their condition in either of two ways: by recovering or by dying. It is presumed that recovery and death rates (r and d) are proportional to the number of infected individuals.

The number of susceptible individuals decreases with the number of infections, but can increase with new births b , proportional to S , or migration which happens at a constant rate m .

Two equations describe the dynamics of I and S , representing the model of the infection process:

$$\frac{dI}{dt} = \gamma k IS - rI - dI \tag{2.3.14}$$

$$\frac{dS}{dt} = -\gamma IS + bS + m \tag{2.3.15}$$

Let us investigate the evolution of the infection at its early stage $t \sim 0$, when we presuppose that the number of the susceptible individuals is much larger than that of the infected ($S(t) \gg I(t)$), and that S is changing slowly ($S(t) \sim S_0 = const$).

Equation (2.3.14) becomes linear and homogeneous:

$$\frac{dI}{dt} = (\gamma k S_0 - r - d)I \tag{2.3.16}$$

The solution is

$$I = I_0 \cdot \exp(Y),$$

where $Y = \gamma k S_0 - r - d$.

If $Y > 0$ the infection spreads, while if $Y < 0$ the infection dies out.

So, our MCF problem is to map $Y < 0$, i.e. the stable eigenvalue of the infection propagation dynamics.

Suppose that $S_0 = 1000$ (a small village), and that factors are distributed as follows:

- Infection coefficient $\gamma \sim U(0, 1)$. The infection is at an early stage, and no information is available about how it is acting.
- Contact coefficient. This distribution describes the probability of a person coming into contact with other individuals. In other words, the probability of meeting all the inhabitants of the village (and of meeting nobody) is low, while the probability of meeting an average number of persons is higher. We will analyze three different *beta* distribution

$$k \sim \text{beta}(2, 7);$$

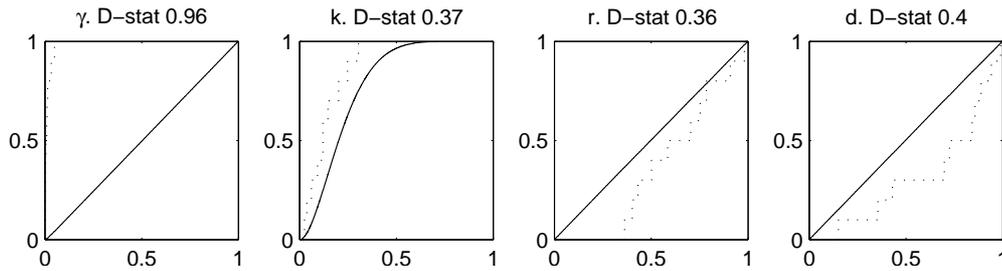
$$k \sim \text{beta}(0.5, 10);$$

$$k \sim \text{beta}(0.2, 15).$$

- Recovery rate $r \sim U(0, 1)$. We assume this to be uniform, as we do not know how it behaves at the beginning of the propagation.
- Death rate $d \sim U(0, 1)$, for the same reason as r .

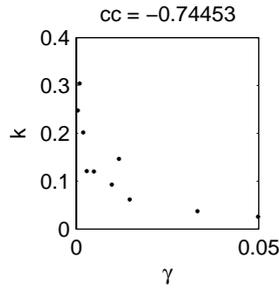
We obtain the following results:

$k \sim \text{beta}(2, 7)$. In this case, the probability that the infection will die out is very small (only 1% of the MC sample). Performing the Smirnov analysis we obtain the following picture:

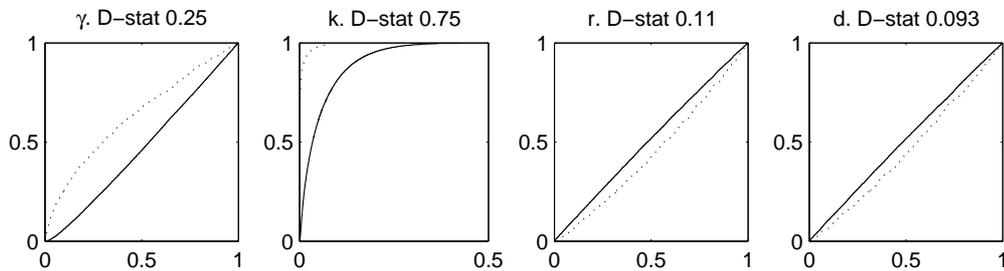


This shows that all input factors have a non-negligible effect in driving the propagation of the infection. However, γ is predominant, based on which it is clear that the infection can die out only for a very narrow range of γ values, in the lowest part of its range. Moreover, we can also see that smaller values of k and larger values of r and d tend to limit propagation of the infection.

The correlation analysis reveals a negative correlation between γ and k under the behavioural subset, which reflects the product interaction between these two factors. Note also that γ -values in this behavioural scatterplot are constrained in the range $[0, 0.05]$ out of a full sample in the range $[0, 1]$.

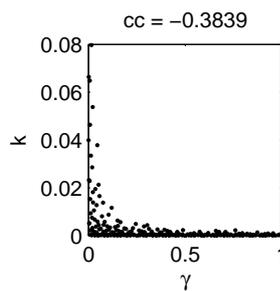


$k \sim \text{beta}(0.5, 10)$. Changing the prior distribution of k raises the probability that the infection will die out to about 20%. Performing the Smirnov analysis produces the following modified picture for the model parameters:

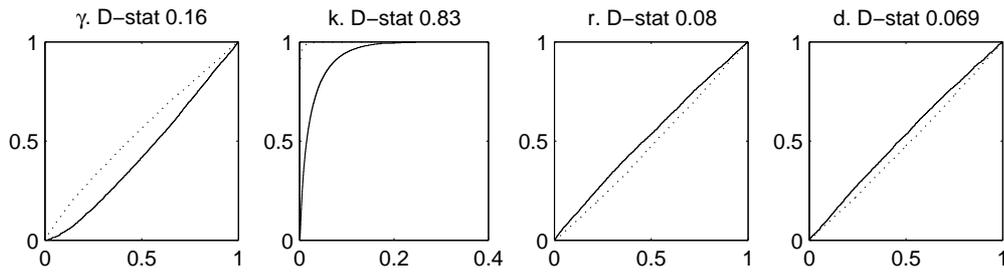


Now the importance of k has increased, and we can see that the behavioural runs are linked merely to k , falling in the lowest part of its range. Moreover, small γ values and large r and d values also tend to produce a declining infection dynamics.

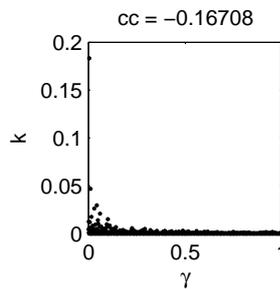
The correlation analysis confirms the negative correlation between γ and k , while the hyperbola shape linked to the product interaction of these two parameters in the model is now more sharply visible.



$k \sim \text{beta}(0.2, 15)$. In this case, the probability of a declining infection dynamics rises to 57%. The Smirnov analysis indicates that k is now the dominant factor in driving the behaviour of the infection dynamics model, leaving a minor role to the remaining parameters:

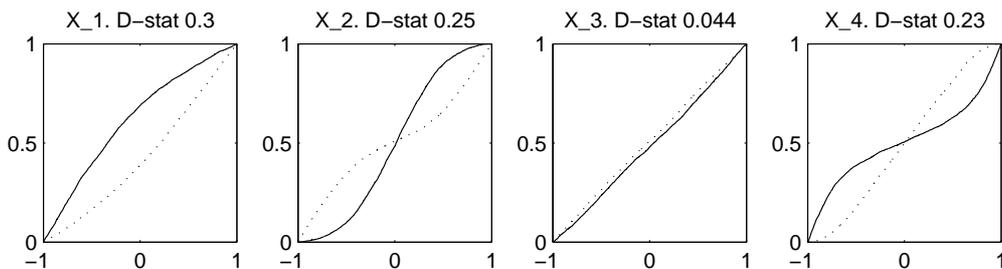


The correlation analysis still produces a negative correlation between γ and k with the hyperbola boundary between the B and \bar{B} subsets.

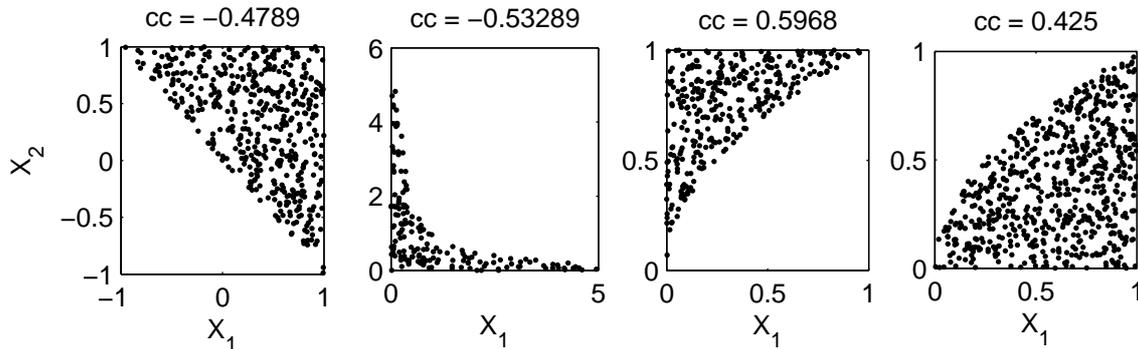


2.4 Exercises

Exercise 2.4.1. Interpret the plots below, which represent the Smirnov test for a set of input factors. The behavioural set is indicated by dotted lines, the non-behavioural by solid lines. Think of some functional forms that could produce such results. The D-stat above each plot indicates the value of the Smirnov statistic.



Exercise 2.4.2. Interpret the plots below, showing the bi-dimensional projections of behavioural subsets. Think of an analytic form of the types of interaction that produce the behavioural sets.



Exercise 2.4.3. Consider the model $Y = Z_1 + Z_2$, with

$$Z_1 \sim N(0, 1)$$

$$Z_2 \sim N(0, 3)$$

Discuss the mapping problem $Y > 0$ analytically and using the MCF techniques.

Exercise 2.4.4. For the same model as in Exercise 2.4.3, map the highest 5% quantile for Y onto the input space.

Exercise 2.4.5. For the same model as in Exercise 2.4.3, map the highest 5% quantile for Y when

$$Z_1 \sim N(0, 1)$$

$$Z_2 \sim N(0, 1).$$

Exercise 2.4.6. Consider the model $Y = X_1 \cdot X_2$, with $X_i \sim N(0, 2)$ and discuss the cases $Y > -1$, $Y > 0$, $Y > 1$.

Exercise 2.4.7. Consider the model $Y = X_1 \cdot X_2 \cdot X_3$, with $X_i \sim U(0, 2)$ and discuss the case $Y < 1$.

2.5 Solutions

Exercise 2.4.1 Three input factors out of four have a significant effect on the behavioural properties of the model. Only X_3 has a negligible effect.

The dotted cumulative distribution curve for X_1 is steepest on the right-hand side (it has shifted downwards), so high values of X_1 are more likely to produce behavioural model realizations. This kind of result suggests a monotonic mapping between X_1 and Y , e.g. a simple linear relationship.

Input factor X_2 , on the other hand, has two different behavioural regions: the dotted line has two separate sections of steepness, one for the smallest values and one for the largest. This implies that the extreme values of X_2 (either smallest or largest) are more likely to produce behavioural realizations of Y . This kind of result suggests a non-monotonic mapping between X_2 and Y , e.g. a quadratic form. Also, interaction effects can lead to the same type of Smirnov test (see Exercise 2.4.6 below).

For X_4 we have the opposite situation with respect to X_2 : the dotted line is steeper in the central part of the support, implying that the extreme values of X_4 have to be avoided in order to produce behavioural model realizations. This kind of result suggests a non-monotonic mapping between X_4 and Y , e.g. a quadratic form with *opposite* concavity with respect to X_2 ¹.

Exercise 2.4.2 The first shape is characterized by a negative correlation between X_1 and X_2 in producing behavioural model realizations. Negative correlations suggest that the two input factors act through sum or product relationships. In this case, a plausible description for the behavioural set is given by $X_1 + X_2 > 0$, i.e. an additive relationship.

The second shape is also characterized by a negative correlation, suggesting action through sum/product. In this case, the shape resembles an hyperbola, suggesting a plausible functional form as $X_1 \cdot X_2 < const.$

The third shape is characterized by a positive correlation, implying that the action is now through difference/quotient. Possible relationships are $a \cdot X_1 - b \cdot X_2^2 < 0$ or $X_1/X_2^2 < const.$

The fourth shape is again with positive correlation, but flipped with respect to the previous one, so plausible relationships are $a \cdot X_1 - b \cdot X_2^2 > 0$ or $X_1/X_2^2 > const.$

Exercise 2.4.3 The behavioural criterion $Y > 0$ is fulfilled by parameter combinations in the upper-right half plane delimited by the line $Z_1 = -Z_2$ (i.e. $Z_1 > -Z_2$). If we also consider the input factor distributions (Gaussian), we know that normal samples will fall into the range $\pm 1.96 \cdot \sigma$ with 95% probability, where σ is the standard deviation of the Gaussian distribution. So, the behavioural samples will be mainly concentrated (with 95% probability) in the upper part of an ellipse with vertical major axis of height 5.88 and horizontal minor axis of width 1.96, cut by the line $Z_1 = -Z_2$. Moreover, since Z_2 has a wider variance than Z_1 , it will also be clear that Z_2 drives the sign of Y more powerfully, i.e. extreme values of Z_2 will be able to drive the sign of Y regardless of the actual values of Z_1 .

We now perform the analysis applying the MCF approach (we give MATLAB commands as an example).

1. Generate a sample of 1000 elements from two normal distributions of standard deviation 1 and 3


```
x1 = randn(1000,1);
x2 = randn(1000,1).*3;
```
2. generate the output:


```
y = x1+x2;
```
3. look for behavioural elements of the sample:


```
ib = find(y>0);
and for the non-behavioural:
in = find(y<=0);
```

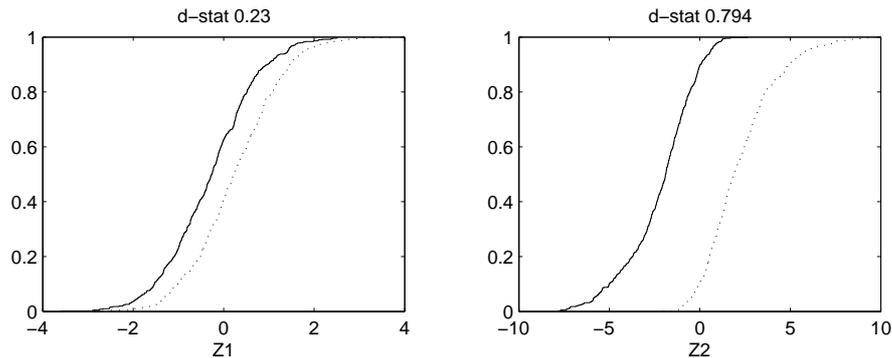
¹The true model used for this example was $Y = 4X_1 + (4X_2^2 - 1) - X_3^3 - (2X_4^4 - 1)$ with $X_i \sim U[-1, 1]$

4. compute the Smirnov statistics (e.g. the outputs `d1` and `d2` of the MATLAB Statistical Toolbox function `kstest2`):

```
[h1, p1, d1] = kstest2(x1(ib),x1(in));
[h2, p2, d2] = kstest2(x2(ib),x2(in));
```

5. plot the empirical cumulative density plots (e.g. using MATLAB Statistical Toolbox function `cdfplot`)

```
figure,
subplot(2,2,1)
h=cdfplot(x1(ib));
set(h,'linestyle',':'),
hold on, cdfplot(x1(in))
gca, title(['d-stat ',num2str(d1)]), xlabel('Z1'), ylabel('')
subplot(2,2,2)
h=cdfplot(x2(ib));
set(h,'linestyle',':'),
hold on, cdfplot(x2(in))
gca, title(['d-stat ',num2str(d2)]), xlabel('Z2'), ylabel('')
```



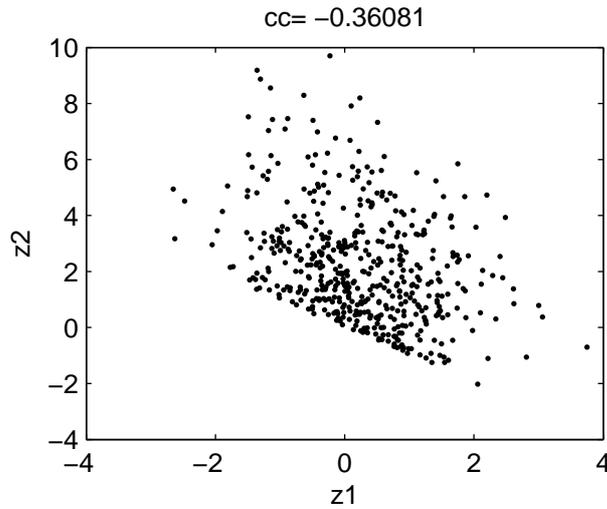
From the Smirnov analysis we can see that, while both input factors have a significant effect on the behavioural realizations of Y , Z_2 has the greater impact on the sign of Y (it has a far larger Smirnov statistic). As already mentioned, this is due to the larger variance of Z_2 with respect to Z_1 , which allows sufficiently large values of Z_2 to force a positive sign in Y , regardless of the values of Z_1 . The behavioural/non-behavioural subsets for Z_2 are therefore almost disjoint (they overlap only in the range $[-1.5, 1.5]$ within a full support of $[-9, 9]$), while for Z_1 the two subsets have a much larger degree of overlap.

6. compute the correlation coefficient under the behavioural subset:

```
cc = corrcoef(x1(ib),x2(ib))
```

and plot the bi-dimensional projection of the behavioural sample, which fills the half plane $Z_1 + Z_2 > 0$, as expected.

```
plot(x1(ib),x2(ib),'.')
xlabel('z1'), ylabel('z2'), title(['cc= ',num2str(cc(2,1))])
```



Exercise 2.4.4 The output Y is the sum of two normally distributed variables, so $Y \sim N(0, \sqrt{10})$. The 5% upper tail of a Gaussian distribution is located at a distance from the mean of $1.65 \cdot \sigma$. Hence the upper 5% tail of the output probability is given by the set $Y > Y^{95}$, where $Y^{95} = 1.65 \cdot \sqrt{10} = 5.22$. So the behavioural set of the input factors is given by the upper half plane delimited by the line $Z_1 + Z_2 = Y^{95}$.

Performing the same analysis applying MCF techniques requires the use of the same sample used in Exercise 3 and re-computing the new behavioural set.

- Sort the output values:

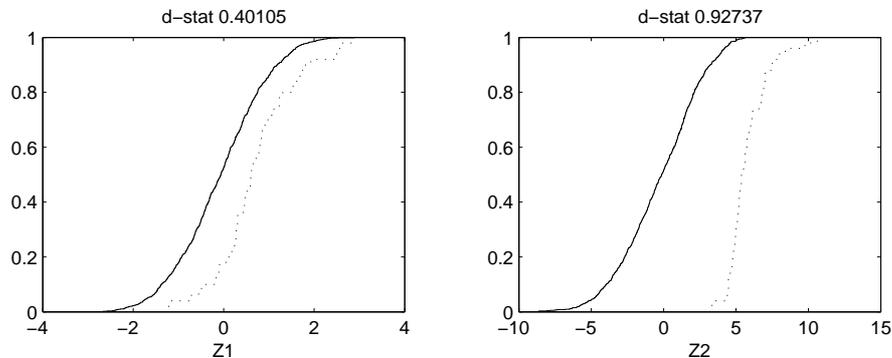
```
[ys, is]=sort(y);
```

- define the behavioural (upper 5 %) and non-behavioural sets (the rest of the sample)

```
ib = is(951:1000);
```

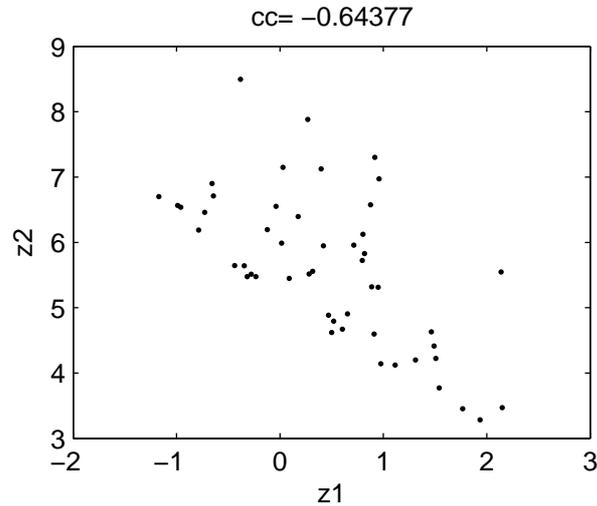
```
in = is(1:950);
```

- compute the Smirnov statistics and plot the cumulative distributions as in the previous exercise:



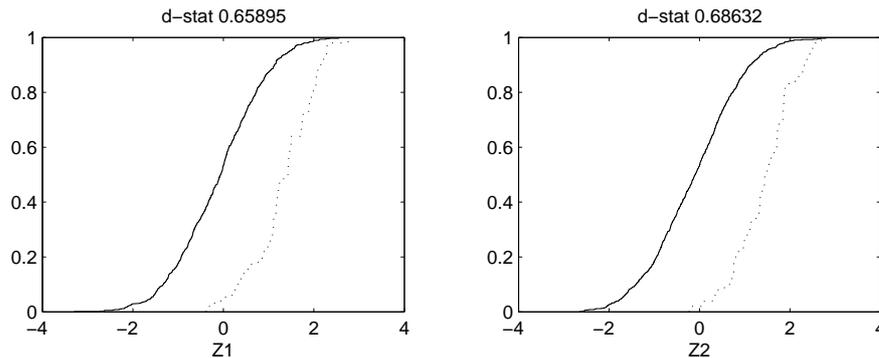
Again we can see that the behavioural and non-behavioural subsets of Z_2 are almost disjoint and sufficiently large values of Z_2 are capable of driving Y into the upper 5% quantile, whatever the value of Z_1 .

- compute the correlation coefficient of the behavioural set and plot the bi-dimensional projection of the input factor behavioural sample, which fills the half plane $Z_1 + Z_2 > Y^{95}$, as expected. This plot also confirms the Smirnov analysis by showing that Z_2 values have to remain significantly positive to drive Y to its upper values, while Z_1 values can range almost symmetrically around zero.

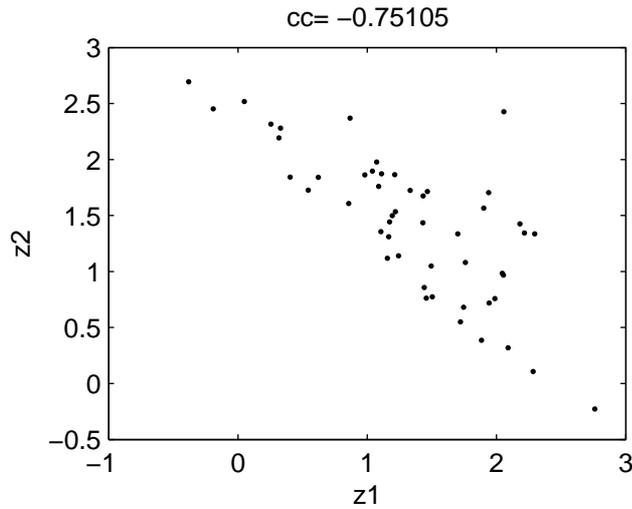


Exercise 2.4.5 In this case, the two input factors have the same variance, so we can expect an equal impact on the extreme values of Y . In analytic terms, Y now has a Gaussian distribution $N(0, \sqrt{2})$. Hence the upper 5% quantile is given by $Y > Y^{95} = 1.65 \cdot \sqrt{2} = 2.33$.

Performing the MCF analysis, we first obtain the Smirnov statistics and plot the cumulative distributions, which clearly display the similarity of the effects of the two input factors on Y .



We can subsequently perform the correlation analysis and plot the bi-dimensional projection of the behavioural sample, which also confirms the equivalent effect of Z_1 and Z_2 on Y .



Exercise 2.4.6 Analytically, the problems can be formalized as $Y > c$, with $c = -1, 0, 1$. This implies that $X_1 \cdot X_2 > c$.

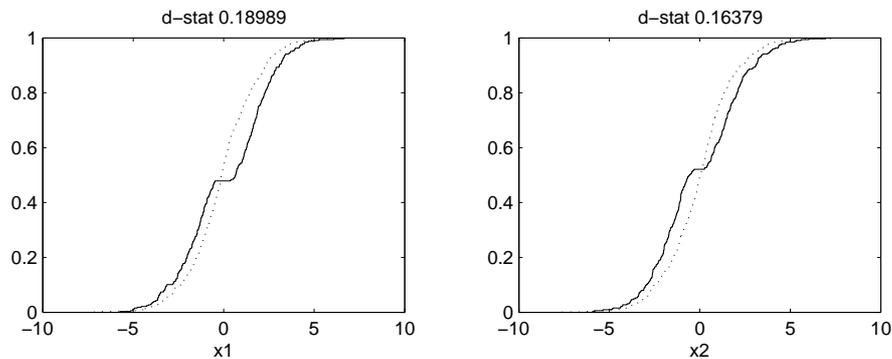
For $c = -1$, the behavioural condition is fulfilled for the portion of the (X_1, X_2) plane *between* the two branches of the hyperbola $X_1 = -1/X_2$. This also implies that the central part of the support of each input factor will be more likely to produce behavioural Y -realizations.

For $c = 0$, the behavioural condition is fulfilled if X_1 and X_2 have the same sign, i.e. all positive values of both X_1 and X_2 or all negative values of both X_1 and X_2 . This also implies that any value in the support of one input factor has an equal probability of producing a behavioural or non-behavioural run, conditional on the value of the other one. Therefore, the Smirnov test would not highlight any significant effect of X_1 and X_2 .

For $c = 1$, the behavioural condition is fulfilled for the two portions of the (X_1, X_2) plane *outside* the two branches of the hyperbola $X_1 = 1/X_2$. This also implies that the lower/upper part of the support of each input factor will be more likely to produce behavioural Y -realizations.

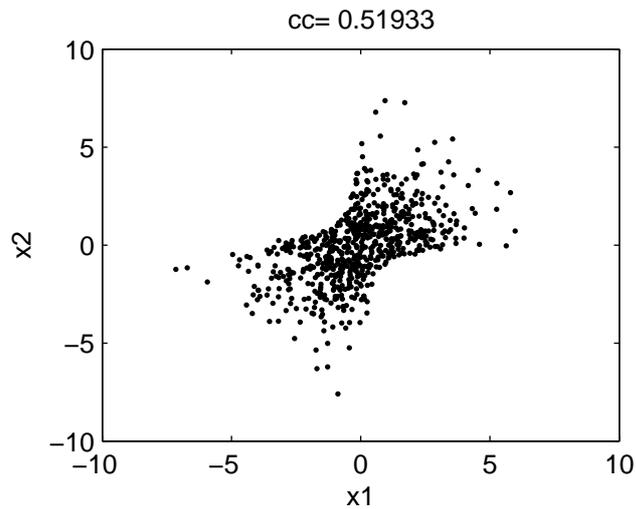
The MCF analysis can be implemented in a completely identical manner to the examples before.

$c = -1$. We first compute the Smirnov statistics and plot the cumulative distributions.



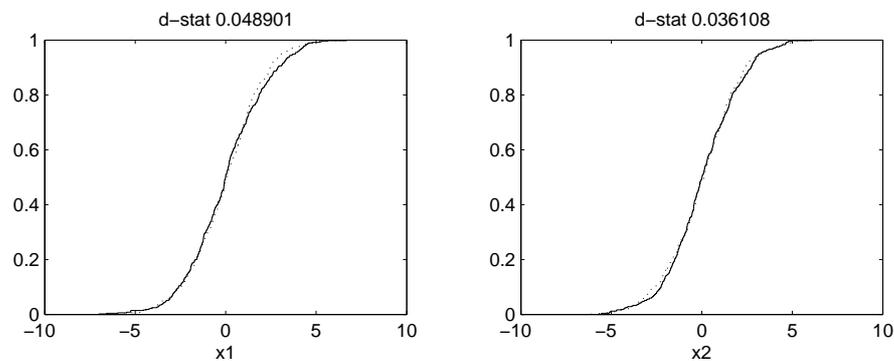
Analysing the steepness of the cumulative distributions under B and \bar{B} shows that the behavioural distribution is concentrated (steeper) in the central part of the initial support, while the non-behavioural is concentrated (steeper) in two disjoint subsets in the lower and upper part of the initial support.

We then compute the correlation coefficient and plot the bi-dimensional projection of the behavioural subset.



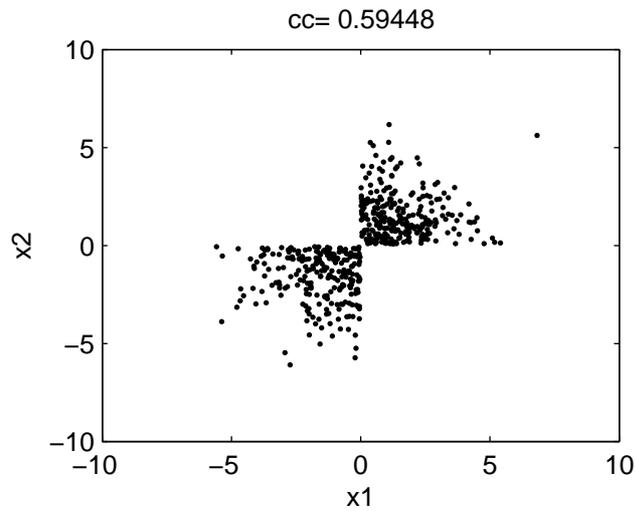
This shows neatly the portion of space between the two branches of hyperbola $X_1 = -1/X_2$ that produces behavioural Y -realizations.

$c = 0$. We first compute the Smirnov statistics and plot the cumulative distributions.



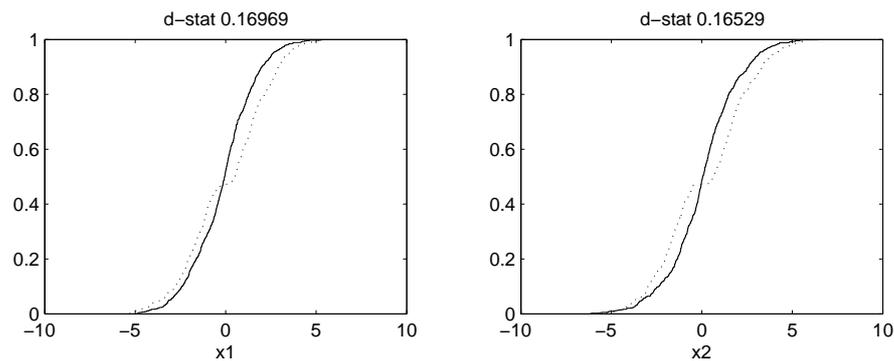
This shows that the two subsets have the same distribution, i.e. any value in the original support for each input factor is equally likely to produce behavioural or non-behavioural realizations.

We then compute the correlation coefficient and plot the bi-dimensional projection of the behavioural subset.



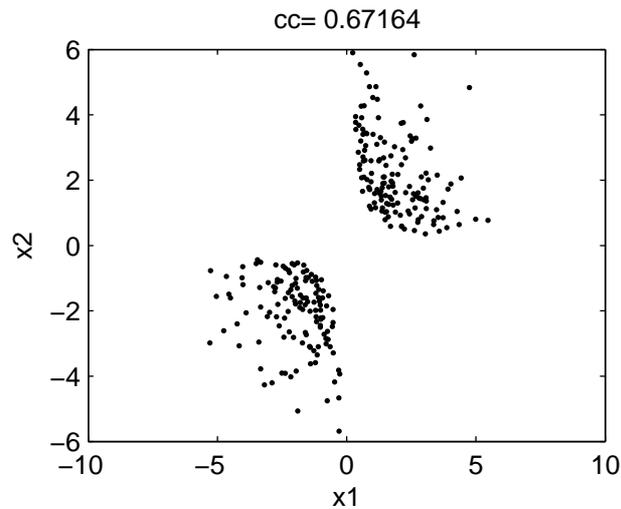
This allows us to understand the interaction mechanism between X_1 and X_2 that produces behavioural realizations.

$c = 1$. We first compute the Smirnov statistics and plot the cumulative distributions.



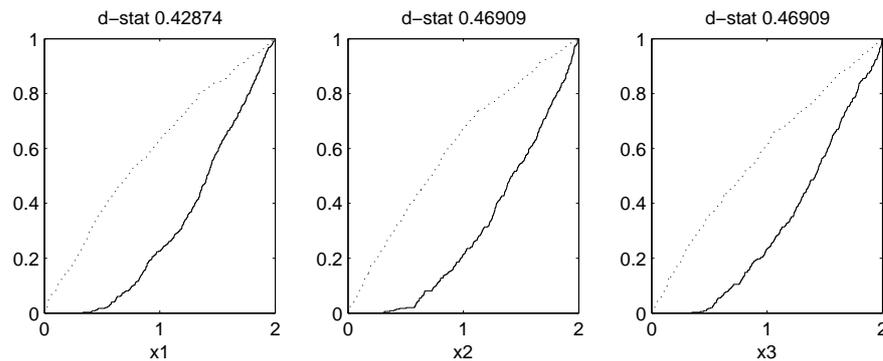
This shows that values of the input factors concentrated in the lower/upper part of each support are more likely to produce behavioural runs, while the central values of the original support are excluded from the behavioural set (the dotted cumulative lines are flat around the zero values of Z_i , implying a zero density of points of the behavioural subset in the central part of each support).

We then compute the correlation coefficient and plot the bi-dimensional projection of the behavioural subset.



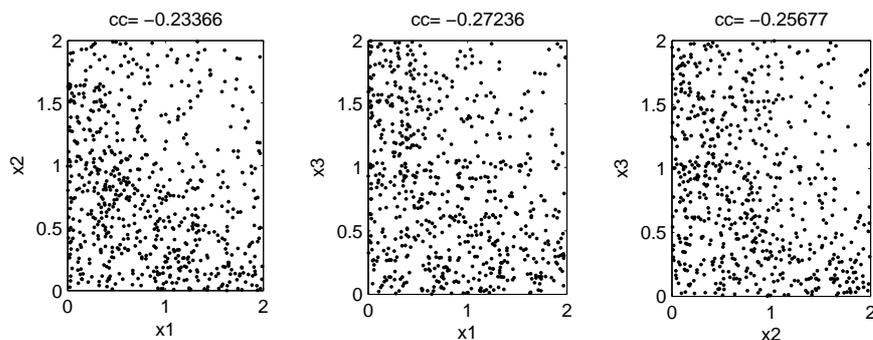
This shows the two disjoint portions of the (X_1, X_2) plane outside the two branches of the hyperbola $X_1 = 1/X_2$ that produce the behavioural Y realizations.

Exercise 2.4.7 We first compute the Smirnov statistics and plot the cumulative distributions.



This shows that, to maintain the output realization smaller than 1, all factors must be sufficiently small. The behavioural distribution is therefore steepest towards the smallest values of the original supports of X_i .

We then compute the correlation coefficient and plot the bi-dimensional projection of the behavioural subset.



The negative correlation coefficients indicate that the effect of input factors on Y is through sums/products. Moreover, these plots show nicely that the upper-right regions, combining large values of all the input factors, have to be avoided to assure that $Y < 1$.

Chapter 3

Metamodelling and the High-Dimensional Model Representation

Let us now consider the problem of representing in a clear and immediate way the relationship $Y = f(X_1, \dots, X_k)$, whose form is usually unknown to the analyst. This is essentially a problem of model approximation or metamodelling, whereby the analyst aims to identify a simple relationship between X_i 's and Y that fits the original model well and is less computationally demanding.

There is a vast literature on this subject. Local approximation methods take the value of f and its derivatives at a base point \mathbf{X}_0 and construct a function that matches the properties of f at \mathbf{X}_0 and in the nearby region (Taylor series). Interpolation methods look at 'nice' functions that go through a set of data points spanning the entire domain of the $Y = f(\cdot)$ mapping. The approximation is then identified by fixing p parameters (e.g. the coefficients of the polynomials) using p data points (Lagrange, Chebyshev interpolation). Tensor products of orthogonal polynomials or complete polynomials are usually applied to span the space of functions in R^k and to interpolate $Y = f(X_1, \dots, X_k)$. Regression/smoothing methods differ from interpolation in that a set of $N > p$ data points is used to identify the approximating function. For univariate functions f , the interpolation, regression and smoothing approaches can be extended by applying piecewise polynomials, constructing functions that are only piecewise smooth. Splines (cubic splines are the most popular) are a powerful and widely used approach to piecewise polynomial interpolation and regression/smoothing (in the latter case they are called smoothing splines). Splines are smooth where the polynomial pieces connect. In the multivariate case, radial basis function (RBF) networks can be seen as the equivalent of univariate piecewise interpolation, regression and smoothing approaches (RBF's are also classified under the heading of kernel regression smoothing methods).

Before proceeding with the description of the methodologies, we would like to draw attention to some additional properties of the approximating functions, which also uncover fundamental links between metamodelling and the theory of variance-based sensitivity analysis.

Let $g(\cdot)$ be the generic function approximating the true model $Y = f(X_1, \dots, X_k)$ and let us assume a quadratic loss function $E[(Y - g(\cdot))^2]$ as a measure of 'fit' for g .

If we were to approximate f with a function of *one single* parameter X_i , what function $g_i^*(X_i)$ would produce minimum loss?

It is well known, from any standard text on statistics, that the univariate function

$$g_i^* = E[Y|X_i], \quad (3.0.1)$$

i.e. the conditional expectation of Y given X_i , is the minimum loss approximation to f . The expression (3.0.1) tells us that at any p -location $X_i = x_{i,p}$, the value of g^* is obtained by integrating (averaging) Y over all the remaining $[X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_k]$ input factors.

Equation (3.0.1) can be generalized to any subset of the input factors \mathbf{X}_I , indexed by $I = (i_1, \dots, i_l)$, as

$$g_I^* = E[Y|\mathbf{X}_I]. \quad (3.0.2)$$

Equations (3.0.1,3.0.2) obviously link to the ANOVA-HDMR decomposition of f .

$$f(X_1, X_2, \dots, X_k) = f_0 + \sum_i f_i + \sum_i \sum_{j>i} f_{ij} + \dots + f_{12\dots k}, \quad (3.0.3)$$

where the connection between the f_i terms of the HDMR and the minimum loss approximating functions $g^*(\cdot)$ is made explicit by

$$\begin{aligned} f_0 &= E(Y) \\ f_i(X_i) &= E(Y|X_i) - f_0 = g_i^* - f_0 \\ f_{ij}(X_i, X_j) &= E(Y|X_i, X_j) - f_i(X_i) - f_j(X_j) - f_0 = g_{i,j}^* - f_i(X_i) - f_j(X_j) - f_0 \end{aligned} \quad (3.0.4)$$

Each term of the ANOVA-HDMR decomposition tells the analyst how much, on average, Y can be moved with respect to its mean level f_0 by acting on single input factors or groups of them. Moreover, the quantity $V(g_I^*(\mathbf{X}_I))/V(Y) = \text{corr}(g_I^*(\mathbf{X}_I), Y)$ is well known in statistics as ‘correlation ratio’ or ‘non-parametric R -squared’, and provides the fraction of the variability of Y that is explained with the *best* predictor based on \mathbf{X}_I . The equivalence between non-parametric R -squared and variance-based sensitivity indices is obvious and this closes the parallel between the problem of estimating and measuring the explanatory power of covariates in regression and variance-based sensitivity analysis.

Coming back to the metamodelling problem, kernel regression methods can be shown to provide, under certain regularity conditions, consistent estimators of $g^*(\cdot)$, which are asymptotically normal at the \sqrt{N} rate (see Doksum and Samarov, 1995, and references cited therein), i.e. as $N \rightarrow \infty$,

$$E \int (\hat{g}(\mathbf{X}_I) - g^*(\mathbf{X}_I))^2 d\mathbf{X}_I = o(N^{-1/2})$$

The metamodelling approach that we follow in this book can be classified under kernel regression smoothing methods. However, due to its recursive implementation, our approach differs with respect to other *en-bloc* methods. This has some advantages, such as the estimation of ‘smoothing parameters’ with maximum likelihood and greater flexibility in managing non-linearities in $f(\cdot)$.

3.1 Estimating HDMR's and meta-models

In the literature on sensitivity analysis there has been a growing interest in metamodelling and smoothing techniques. Storlie and Helton (2006) have reviewed smoothing methods for sensitivity analysis, from smoothing splines to various types of univariate and multivariate kernel regression approaches. Li et al. (2002, 2006) developed the so-called Random Sampling

HDMR, which involves approximating the truncated HDMR expansion up to order three, based on orthogonal polynomials. Pappenberger and Stauch (2007) uses spline smoothing to estimate sensitivity indices. Ratto et al. (2004, 2006, 2007) have developed a non-parametric approach which is very similar to smoothing splines and kernel regression approaches, but which is based on recursive filtering and smoothing estimation (the Kalman Filter, KF). Such a recursive least-squares implementation has some fundamental advantages: (a) it is coupled with optimal Maximum Likelihood estimation, thus allowing for an objective estimation of ‘smoothing’ parameters, and (b) it allows for greater flexibility in adapting to local discontinuities, heavy non-linearity and heteroscedastic error terms (see below). All such methods can be assigned to the regression/smoothing class of approximation approaches.

An example of interpolating metamodels, on the other hand, is given by Gaussian emulators (see Oakley and O’Hagan, 2004, and the references cited therein, for a detailed description) and kriging metamodels (Kleijnen, 2007b,a). Kriging metamodels are similar to Gaussian, except that they do not rely on Bayesian interpretation. While *theoretically* appealing, Gaussian emulators can be prone to the curse of dimensionality and to the smoothness assumptions of the function under analysis. This is because, instead of trying to identify the best predictors of Y based on a subset of input factors or on low-order ANOVA-HDMR terms, Gaussian emulators try to interpolate and predict the $f(\cdot)$ mapping by applying a Gaussian kernel of the same k -dimensionality as the input parameter space. Therefore, as k increases, the number of ‘hyper-parameters’ to be estimated (linked to the covariance structure of the k -dimensional Gaussian kernel), increases strongly, often implying problems with identification and over-parameterization.

Such problems are well known in standard interpolation and smoothing techniques based on k -dimensional kernel regressors. They imply that Gaussian emulators are only effective, *in practice*, for model structures having a small number of significant main effects and very mild interactions, for which such problems are made irrelevant by the very few highly identifiable elements of the $f(\cdot)$ mapping.

In other words, in regression/smoothing techniques, metamodels are based on subsets of input factors and/or truncated HDMR expansions of order smaller than k , and their identification and estimation *incorporate* sensitivity analysis criteria, in that non-significant contributions to Y are identified and eliminated *within* the process of construction of the approximation to Y . In contrast to this, Gaussian emulators aim first to estimate a full k -order mapping on the basis that a sensitivity analysis applied afterwards to the emulator will automatically reveal the significant contributions to Y .

Once identified, estimated and parameterized, metamodels provide a direct, albeit approximated, analytic expression of the $Y = f(X_1, \dots, X_k)$ mapping, which accounts for non-linearities and interaction terms of increasing order. As such, they can be used for various purposes:

- sensitivity analysis, by helping to highlight the most important input factors of the mapping;
- model simplification, by finding a surrogate model containing a subset of the input factors that account for most of the variability of Y ;
- model calibration, in which the metamodel is used to find directly the optimal parameterization for the fulfilment of the given calibration criteria.

A detailed description of all the available metamodelling techniques is beyond the scope of this book: readers can refer to the cited works for further information. We will concentrate here on non-parametric methods and, in particular, we will demonstrate simple implementations of univariate non-parametric smoothing methods that give the ‘flavour’ of the more sophisticated multivariate, recursive procedure of Ratto et al. (2004, 2006, 2007). In non-parametric methods, the function $E(Y|X_i)$ is not approximated by a basis of functions that span the entire domain of X_i ; rather many ‘local’ approximations are identified which move along the X_i -axis. Such ‘local’ functions are subsequently joined using *ad hoc* criteria, such as imposing some smoothness properties like continuity (piecewise linear interpolation is an example of this). This in practice gives a ‘look-up’ table of the function $f_i = g^*(X_i)$ which can subsequently be parameterized using functional bases, such as polynomials, Fourier expansions or linear wavelets.

Smoothing scatterplots using the Haar wavelet

Wavelets, and in particular the Haar wavelet, provide a very simple approach to smoothing signals.

The 2×2 Haar matrix is given by

$$H_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

Given an MC sample of $f(\cdot)$ whose length N is a power of two $(y_1, y_2, \dots, y_{2^{n-1}}, y_{2^n})$, i.e. $N = 2^n$, and where the sample is sorted with respect to the input factor X_i under analysis, we may group its elements as $((y_1, y_2), \dots, (y_{2^{n-1}}, y_{2^n}))$ and we may right-multiply each term by matrix H_2 ,

$$H_2 \begin{pmatrix} y_{2j-1} \\ y_{2j} \end{pmatrix} = \begin{pmatrix} s_j \\ d_j \end{pmatrix} \quad \text{for } j = 1, \dots, 2^{n-1}$$

obtaining two new sequences $(s_1, \dots, s_{2^{n-1}})$ and $(d_1, \dots, d_{2^{n-1}})$. The sequence s gives the *sum* between two consecutive points whilst the sequence d gives the *difference*. Since the original signal has a length equal to 2^n , one can recursively apply the same procedure to s -sequences up to n times. For each $\lambda = 1, \dots, n$ we call the associated sequence s^λ the λ th Haar approximation coefficients and d^λ the λ th Haar detail coefficients.

For each $\lambda = 1, \dots, n$ we may use s^λ to create an approximation \hat{f}^λ of the original signal f . This is obtained first by re-scaling each value of s^λ by $\sqrt{2^\lambda}$ and then by replicating each of them 2^λ times.

Example: Consider a sample of length $2^4 = 16$ from a function $Y = f(X) = (X - 0.5)^2 + \varepsilon$, where $X \sim U[0, 1]$ and ε is a white noise normally distributed $N(0, 0.03)$:

$$f = (0.21, 0.19, 0.14, 0.1, 0.05, 0.01, 0.03, -0.01, 0.01, -0.03, 0.01, 0.06, 0.04, 0.1, 0.13, 0.19).$$

- First stage ($\lambda = 1$). From the first pair of y -points we compute the first element of s^1 : $(0.21 + 0.19)/\sqrt{2} = 0.2828$. Repeating this with all pairs of y we get:

$$s^1 = (0.2828, 0.1697, 0.0424, 0.0141, -0.0141, 0.0495, 0.0990, 0.2263)$$

Scaling the first term of s^1 by $\sqrt{2}$ we get 0.2, which is replicated 2^1 times to construct the first two elements of \hat{f}^1 . This is repeated to get:

$$\hat{f}^1 = (0.2, 0.2, 0.12, 0.12, 0.03, 0.03, 0.01, 0.01, \\ -0.01, -0.01, 0.035, 0.035, 0.07, 0.07, 0.16, 0.16)$$

- Second stage ($\lambda = 2$): Replicating the steps of the first stage to the s^1 sequence, we get:

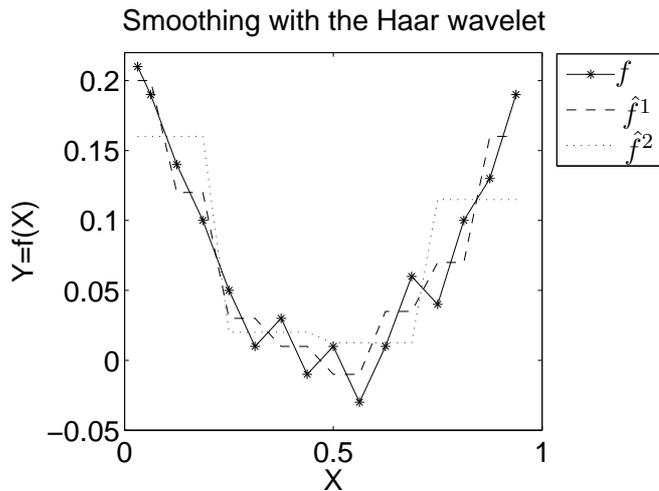
$$s^2 = (0.32, 0.04, 0.025, 0.23)$$

and

$$\hat{f}^2 = (0.16, 0.16, 0.16, 0.16, 0.02, 0.02, 0.02, 0.02, \\ 0.013, 0.013, 0.013, 0.013, 0.12, 0.12, 0.12, 0.12)$$

- ... and so on.

If we compare the graph of the original signal f and its first- and second-level approximations \hat{f}^1 , \hat{f}^2 we obtain the following picture:



The picture above shows that \hat{f}^λ may be used to infer the smoothed behaviour of the signal f , where λ plays the role of a smoothing parameter: the higher λ , the smoother the signal extraction. At the highest level of smoothing $\lambda = n$, the Haar wavelet will simply provide the unconditional mean $E(Y)$. Note also that this smoothing approach can be implemented recursively, and does not require any matrix inversion.

Exercise 3.1.1. Write a code for the Haar wavelet and reproduce the results shown in the previous example.

Spline smoothing

Smoothing splines are a non-parametric method which is very useful for understanding the more sophisticated methodology to be discussed later in this report. Given a Monte Carlo

sample of size N , a cubic smoothing spline is a function \hat{g} that minimizes the penalized residuals' sum of squares:

$$\sum_{j=1}^N [y_j - \hat{g}(x_j)]^2 + \lambda \int_a^b \left[\frac{d^2 \hat{g}(x)}{dx^2} \right]^2 dx \quad (3.1.1)$$

where $a \leq \min(x_j)$, $b \geq \max(x_j)$, $j = 1, \dots, N$ and λ is the Lagrange Multiplier, which plays the role of a smoothing parameter (i.e. the bigger λ , the smoother \hat{g}). In (3.1.1), the first term is the sum of squared residuals, measuring fit to the data, while the second term penalizes too high a curvature in \hat{g} . This approach is also known in numerical analysis as *regularization* (or deterministic regularization, DR).

The unique, explicit solution to (3.1.1) is given by a natural cubic spline with knots at each observed value x_j (see Hastie and Tibshirani, 1990). A cubic polynomial spline is a function that is a cubic polynomial on any interval defined by adjacent knots, has two continuous derivatives and a third which is a step function that jumps at the knots.

Spline smoothing can easily be implemented by taking the discrete formulation of (3.1.1), where the long-term signal \hat{t} is given by the solution of the minimization problem:

$$\min \left(\sum_{s=1}^N [y_s - \hat{t}_s]^2 + \lambda \sum_{s=2}^{N-1} [(\hat{t}_{s+1} - \hat{t}_s) - (\hat{t}_s - \hat{t}_{s-1})]^2 \right), \quad (3.1.2)$$

where the index s scans the data in a *sorted* order with respect to the input factor under analysis. In econometrics this is called the Hodrick-Prescott (HP) filter (Hodrick and Prescott, 1980), and it is used for trend extraction in economic time-series analysis (in such cases s scans the data in temporal order).

The solution to this problem can be written taking derivatives with respect to \hat{t}_s :

$$\begin{aligned} y_1 &= (1 + \lambda)\hat{t}_1 - 2\lambda\hat{t}_2 + \lambda\hat{t}_3 \\ y_2 &= -2\lambda\hat{t}_1 + (1 + 5\lambda)\hat{t}_2 - 4\lambda\hat{t}_3 + \lambda\hat{t}_4 \\ &\dots \\ y_s &= \lambda\hat{t}_{s-2} - 4\lambda\hat{t}_{s-1} + (1 + 6\lambda)\hat{t}_s - 4\lambda\hat{t}_{s+1} + \lambda\hat{t}_{s+2} \\ &\quad s = 3, \dots, n-2 \\ &\dots \\ y_{n-1} &= \lambda\hat{t}_{n-3} - 4\lambda\hat{t}_{n-2} + (1 + 5\lambda)\hat{t}_{n-1} - 2\lambda\hat{t}_n \\ y_n &= \lambda\hat{t}_{n-2} - 2\lambda\hat{t}_{n-1} + (1 + \lambda)\hat{t}_n. \end{aligned}$$

This can easily be put into matrix form,

$$\mathbf{y} = (\mathbf{I} + \lambda \mathbf{U}^T \mathbf{U}) \cdot \hat{\mathbf{t}}^\lambda, \quad (3.1.3)$$

where \mathbf{I} is the $N \times N$ identity matrix, \mathbf{y} is the $1 \times N$ vector of model output data and \mathbf{U} is the following $N \times N$ matrix:

$$\mathbf{U} = \begin{bmatrix} 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 1 & -2 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & 1 & -2 & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 & -2 & 1 \end{bmatrix}$$

The explicit solution is easily found to be:

$$\hat{\mathbf{t}}^\lambda = (\mathbf{I} + \lambda \mathbf{U}^T \mathbf{U})^{-1} \cdot \mathbf{y} = \mathbf{W}^{-1} \cdot \mathbf{y}. \quad (3.1.4)$$

The ‘trend’ $\hat{\mathbf{t}}^\lambda$ is a non-parametric estimate of the univariate ‘metamodel’ $f_i(X_i) = E(Y|X_i)$, which approximates the function $f(\cdot)$. The dependence of such an estimate on the smoothing parameter is made explicit in (3.1.3,3.1.4) by the λ exponent in $\hat{\mathbf{t}}^\lambda$. The appropriate value for λ is not immediately apparent: cross-validation, graphical methods and measures of degrees of freedom are used for this purpose (see Hastie and Tibshirani, 1990, Storlie and Helton, 2006, for more details) . Moreover, en-bloc smoothing methods like the one presented here require the inversion of a matrix of the same dimension N of the MC sample, so that they can be computationally very intensive and require a large memory size for large MC samples. This contrasts with recursive methods like the State Dependent Regression described next. Finally, note that in the en-bloc approaches it is also possible to obtain the standard errors of the estimated $\hat{\mathbf{t}}^\lambda$. Readers can refer to Hastie and Tibshirani (1990) for a detailed discussion of en-bloc methods and estimation of standard error bands and to Young and Pedregal (1999) for a combined discussion of recursive and en-bloc approaches.

Example: Consider the same example used above for the Haar wavelet. Using the same data, we get the following estimates for $\hat{\mathbf{t}}^\lambda$, at various smoothing levels:

- $\lambda = 1$:

$$\begin{aligned} \hat{\mathbf{t}}^1 &= (0.22, 0.18, 0.14, 0.096, 0.055, 0.026, 0.011, \\ &\quad -0.0012, -0.0064, -0.0069, 0.01, 0.036, 0.059, 0.095, 0.14, 0.18) \end{aligned}$$

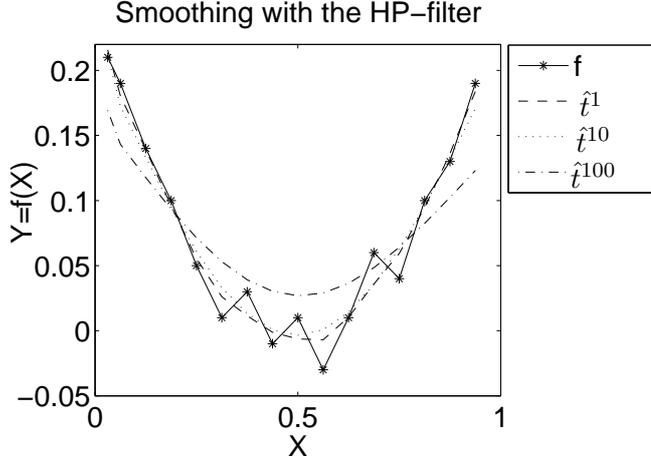
- $\lambda = 10$:

$$\begin{aligned} \hat{\mathbf{t}}^{10} &= (0.21, 0.17, 0.13, 0.095, 0.061, 0.033, 0.013, \\ &\quad 0.00069, -0.0035, 0.00068, 0.014, 0.036, 0.063, 0.096, 0.13, 0.17) \end{aligned}$$

- $\lambda = 100$:

$$\begin{aligned} \hat{\mathbf{t}}^{100} &= (0.17, 0.14, 0.12, 0.093, 0.071, 0.053, 0.039, \\ &\quad 0.03, 0.027, 0.029, 0.036, 0.048, 0.064, 0.082, 0.1, 0.12) \end{aligned}$$

Such smoothed estimates are shown in the following plot:



This plot shows the smoothing effect of increasing λ . If λ tends to infinity, the HP-filter simply provides the unconditional mean $E(Y)$. Judging by eye, the best smoothing seems to be obtained using λ values between 1 and 10.

Exercise 3.1.2. Write a code for the spline smoothing (HP-filter) and reproduce the results shown in the previous example.

State Dependent Regressions

Ratto et al. (2004, 2006, 2007) have recently presented a flexible and efficient approach to the estimation of $g_I^*(\mathbf{X}_I) = E(Y|\mathbf{X}_I)$ and of truncated ANOVA-HDMR decompositions. The estimation procedure is based on State Dependent Regression (SDR) models: a non-parametric approach, based on recursive filtering and smoothing estimation.

In brief, the fundamental concept underlying of SDR approach is that any term like $E(Y|\mathbf{X}_I)$ can be viewed as an SDR model of the form:

$$Y_s = p_{I,s}(\xi_I) + e_s \quad (3.1.5)$$

where the state dependent parameter (SDP) $p_{I,s}(\xi_I)$, $I = i_1, \dots, i_l$, depends on a state variable ξ_I that moves according to a generalized sorting strategy along the co-ordinates of the single factors or groups of factors indexed by I ; and e_s is the residual, i.e. the portion of variability of Y that cannot be explained by the group of factors indexed by I only.

Extending this definition, the truncated ANOVA-HDMR expansion can also be expressed as an SDR model as follows:

$$\begin{aligned} Y_s - f_0 &= \sum_i p_{i,s}(\xi_i) + \sum_{j>i} p_{ij,s}(\xi_{ij}) + \sum_{l>j>i} p_{ijl,s}(\xi_{ijl}) + e_s \\ &= \sum_i f_{i,s}(X_i) + \sum_{j>i} f_{ij,s}(X_i, X_j) + \sum_{l>j>i} f_{ijl,s}(X_i, X_j, X_l) + e_s \end{aligned} \quad (3.1.6)$$

where e_s now represents the higher-order terms of the ANOVA-HDMR. Note also that, in formulation (3.1.5), the SDP accounts for all cumulative effects within group I , be they first-order or interaction terms. As such, it can be applied to any type of dependency structure among input factors. On the other hand, in the HDMR formulation (3.1.6), each SDP accounts only for its associated first-order or interaction term of the HDMR.

According to the generalized sorting strategy adopted in (3.1.5) and (3.1.6), the group of input factors of interest I is characterized by a low-frequency spectrum (e.g. by some quasi-periodic pattern) while the remaining ones present a white spectrum. In this way, the estimation of the various HDMR terms is reduced to the extraction of the low-frequency component (i.e. of a ‘trend’) from the sorted output Y . To do so, the SDP’s are modelled by one member of the generalized random walk (GRW) class of non-stationary processes. For instance, the integrated random walk (IRW) process turns out to produce good results, since it ensures that the estimated SDP relationship has the smooth properties of a cubic spline.

Given the IRW characterization, the model (3.1.6) can be put into state-space form as:

$$\begin{aligned} \text{Observation Equation: } Y_s &= \mathbf{1}_s \mathbf{p}_s + e_s \\ \text{State Equations: } p_{I,s} &= p_{I,s-1} + d_{I,s-1} \\ d_{I,s} &= d_{I,s-1} + \eta_{I,s} \end{aligned} \quad (3.1.7)$$

where $\mathbf{1}_s$ is the unit constant regressor and e_s (observation noise) and $\eta_{I,s}$ (system disturbances) are zero-mean white-noise inputs with variance σ^2 and $\sigma_{\eta(I)}^2$ respectively. Given this formulation, SDP’s are estimated using the recursive Kalman Filter and associated recursive Fixed Interval Smoothing (FIS) algorithm (Kalman, 1960; Young, 1999b). The recursive state estimation requires that each SDP be estimated in turn, each with a different ordering of the data, within a backfitting procedure (Young, 2000, 2001). At each backfitting iteration, the hyper-parameters associated with (3.1.7), namely the white noise variances σ^2 and $\sigma_{\eta(I)}^2$, are optimized by maximum likelihood (ML), using prediction error decomposition (Schweppe, 1965). In fact, by a simple reformulation of the KF and FIS algorithms, each SDP and its stochastic IRW process model can be entirely characterized by one Noise Variance Ratio (NVR) hyper-parameter, where $\text{NVR}_I = \sigma_{\eta(I)}^2 / \sigma^2$. Hence, only NVR’s need to be optimized.

One very useful implication of the KF/FIS algorithms underlying SDR estimation is that the standard error $\sigma_{f,s}$ of the SDP’s is recursively produced in a natural manner by the covariance matrix of the state vector. This allows us to distinguish the significance of the estimated functions $E(Y|\mathbf{X}_I)$.

A similar backfitting procedure is also applied in the Generalized Additive Modelling (GAM) approach of Hastie and Tibshirani (1990). However, the SDP approach is couched with optimal ML estimation and seems more elegant and flexible than the scatterplot smoothing used by Hastie, Tibshirani and others. Moreover, the stochastic model for each SDP can be based on any member of the GRW family and does not have to be limited to the IRW model. For example, Random Walk (RW) or Smoothed Random Walk (SRW) might be identified as being preferable in certain circumstances because they yield less smooth estimates. The recursive formulation of the SDR modelling also implies much greater flexibility in the estimations, whenever required. For example, if any sharp changes or jumps seem possible following an initial identification phase in the analysis, then these can be handled using ‘variance intervention’ (see Ng and Young, 1990), i.e. the NVR at some specific locations is allowed to assume a large value in order to accommodate sharp changes in the underlying Y -signal. Moreover, in the case of heteroscedastic behaviour (i.e. change in variance) in the observation noise, this can be further modelled in some SDP manner, e.g. by making the NVR a function of the state ξ_I . This can then easily be plugged into the KF and FIS algorithms. In practice, these extensions enable the SDR models to adapt to such situations, *typical in non-linear systems*, where one single, constant smoothing parameter λ does not allow us to follow appropriately the observed patterns of the $f(\cdot)$ mapping.

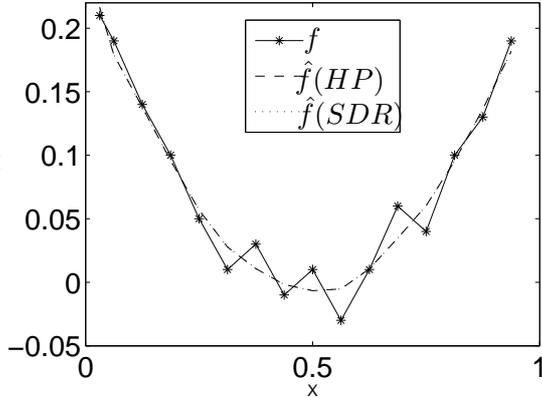
In its basic IRW formulation, the links between the SDR approach and the HP-filter

smoothing are clear: both of them have the properties of a smoothing cubic spline. Moreover, it is also easy to verify the equivalence between the NVR and λ , linked by the simple relationship: $\lambda = 1/\text{NVR}$. NVR therefore plays the role of the inverse of a smoothing parameter. SDR advantages, however, are in terms of the ML estimation of the NVR, which makes the choice of the smoothing parameter completely objective, and in the great flexibility features mentioned above. These properties provide optimal convergence properties of the SDR estimates to the best least-squares predictor of Y , given by $g^*(\mathbf{X}_I)$.

The SDR smoothing techniques based on GRW processes can also be seen as low-pass filters. In the case of IRW, the 50% cutoff frequency ω^* is linked to the NVR by the relationship $\text{NVR} = 4(1 - \cos(2\pi\omega^*))^2$ (see e.g. Young and Pedregal, 1999). The period $T^* = 2\pi/\omega^*$ which is obtained with maximum likelihood can be compared to N to get an idea of 'typical' T^*/N ratios and to identify some rule-of-thumb criterion for the smoothing parameter λ of the HP-filter. Our experience suggests that a reasonable rule of thumb can be to set $T^* = N/3 \rightarrow N/2$ and to derive λ accordingly.

In the last step of the SDR analysis, the smoothed non-parametric 'curves' obtained from the SDR model estimation are parameterized by, for example, a linear wavelet functional approximation, allowing us to build a full metamodel to replace the original one.

Example: Consider again the simple example used before. Using the same data, we get $\text{NVR}=0.59$ from ML estimation. This implies that the equivalent optimal λ for the HP-filter would be about 1.7, matching the values that 'looked' acceptable in the HP-filter analysis. The plots of the SDR estimate and the equivalent HP-filter estimate using $\lambda = 1/\text{NVR}$ are shown below:



This clearly shows substantial identity between the HP-filter and IRW smoothing (the curves overlap), as well as the importance of an estimation procedure linked to maximum likelihood, providing for optimal identification of the smoothing parameter.

Estimating sensitivity indices

Once the smoothed estimates \hat{g}_I of $E(Y|\mathbf{X}_I)$ have been obtained, the estimation of sensitivity indices $S_I = V[E(Y|\mathbf{X}_I)]/V(I)$ is straightforward. As discussed in Doksum and Samarov (1995), three estimators can be applied for this purpose. The one we use is:

$$\hat{S}_I = \frac{N^{-1} \sum_{s=1}^N (\hat{g}_I(\mathbf{x}_{I,s}) - \bar{g})^2}{\sigma_Y^2} \quad (3.1.8)$$

where $\bar{g} = N^{-1} \sum \hat{g}_I(\mathbf{x}_{I,s})$ and $\sigma_Y^2 = N^{-1} \sum (y_s - \bar{Y})^2$. Doksum and Samarov (1995) also provide the error estimate for (3.1.8), by showing that $N^{1/2}(\hat{S}_I - S_I)$ is asymptotically normal with mean zero and variance $(1 - S_I)^2 V[y^{*2} - u^2]$, where y^* and u are the standardized output and residual respectively, i.e. $y_s^* = (y_s - \mu_Y)/\sigma_Y$ and $u_s = [y_s - \hat{g}_I(\mathbf{x}_{I,s})]/(\sigma_Y(1 - S_I)^{1/2})$. Hence, the standard error of the estimate of S_I is given by:

$$SE(\hat{S}_I) = (1 - \hat{S}_I) \text{std}[y_s^{*2} - u_s^2]/N^{1/2} \tag{3.1.9}$$

3.2 Examples

3.2.1 A simple example

We give here an example of the smoothing estimation procedures. Let us consider the simple model

$$Y = X_1 + X_2^2 + X_1 \cdot X_2 \tag{3.2.1}$$

with input distributions $X_i \sim N(0, 1)$. This model has the simple ANOVA-HDMR representation:

$$\begin{aligned} f_0 &= 1 \\ f_1(X_1) &= X_1 \\ f_2(X_2) &= X_2^2 - 1 \\ f_{1,2}(X_1, X_2) &= X_1 \cdot X_2 \end{aligned} \tag{3.2.2}$$

We apply the smoothing procedures described above, by using an LP_τ sample of $N = 256$ model evaluations.

Haar wavelet smoothing

Let us first perform the smoothing with the Haar wavelet. As a rule-of-thumb criterion for the smoothing parameter of the Haar wavelet, we dictate that the extracted signal be made of $8 = 2^3$ values, i.e. we construct the smoothing by taking eight local averages from the sample of Y divided into eight bins. Given the sample of $N = 256 = 2^8 = 2^n$, this implies that $\lambda=8-3=5$. The sample of the output Y has to be sorted according to each input factor and the Haar smoothing procedure described above has to be applied for each sorted sample.

The results of this procedure are shown in Figure 3.1. This shows that the rule of thumb of eight local means is able to provide an illustrative idea of the f_i patterns.

Spline smoothing (HP-filter)

In this case we have to choose the value for the smoothing parameter of the HP-filter. Using $\lambda=50000$, which is in the range of rule-of-thumb values for the cut-off frequency ω^* described above, we get the results shown in Figure 3.2. We can observe the much nicer pattern provided by the cubic spline properties of the HP-filter with respect to the Haar wavelet. The problem, however, still lies in the approximate choice of λ .

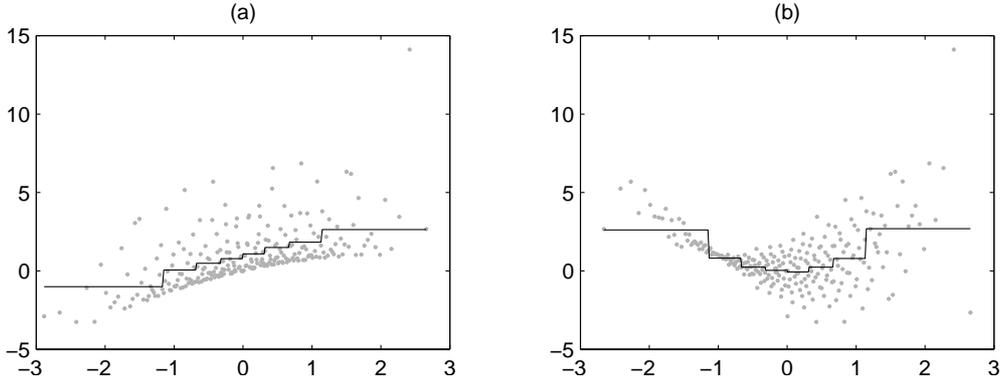


Figure 3.1: Haar estimation of the first-order HDMR of the simple model (3.2.1). (a,b) scatterplots of Y versus X_1 and X_2 (grey dots), with the smoothed estimates of the $f_i + f_0$ functions (solid lines).

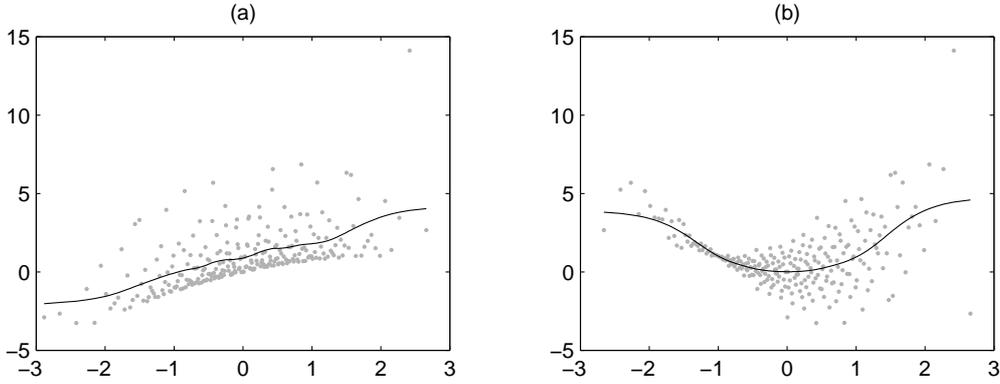


Figure 3.2: HP-filter estimation of the first-order HDMR of the simple model (3.2.1). (a,b) scatterplots of Y versus X_1 and X_2 (grey dots), with the smoothed estimates of the $f_i + f_0$ functions (solid lines).

SDR estimation

The maximum likelihood optimization of the NVR's gives the following results for the two input factors: $NVR_1=4.46e-6$, $NVR_2= 4.675e-5$.

In Figure 3.3(a,b) we show the scatterplots of Y versus the two model parameters together with the result of the SDR recursive filtering and smoothing estimation of the f_i terms ($f_i + f_0$ are plotted to match the same level of Y values). In Figure 3.3(c,d) we show the detail of the SDR estimates, compared to the analytical values (3.2.2). Apart from expectable border phenomena, the SDR estimates are excellent. In Figure 3.3(c,d) the dashed lines show the 95% error bands ($= \pm 2 \cdot \sigma_{f,s}$) of the estimated patterns. This permits us to assess the significance of the estimated patterns by simply checking whether the zero-line is always included in the error band (implying insignificance) or not.

Using $1/NVR_i$ as smoothing parameters for the HP-filter analysis produces the same results as for the SDR analysis, provided that the backfitting procedure is also applied – otherwise the results will be slightly different.

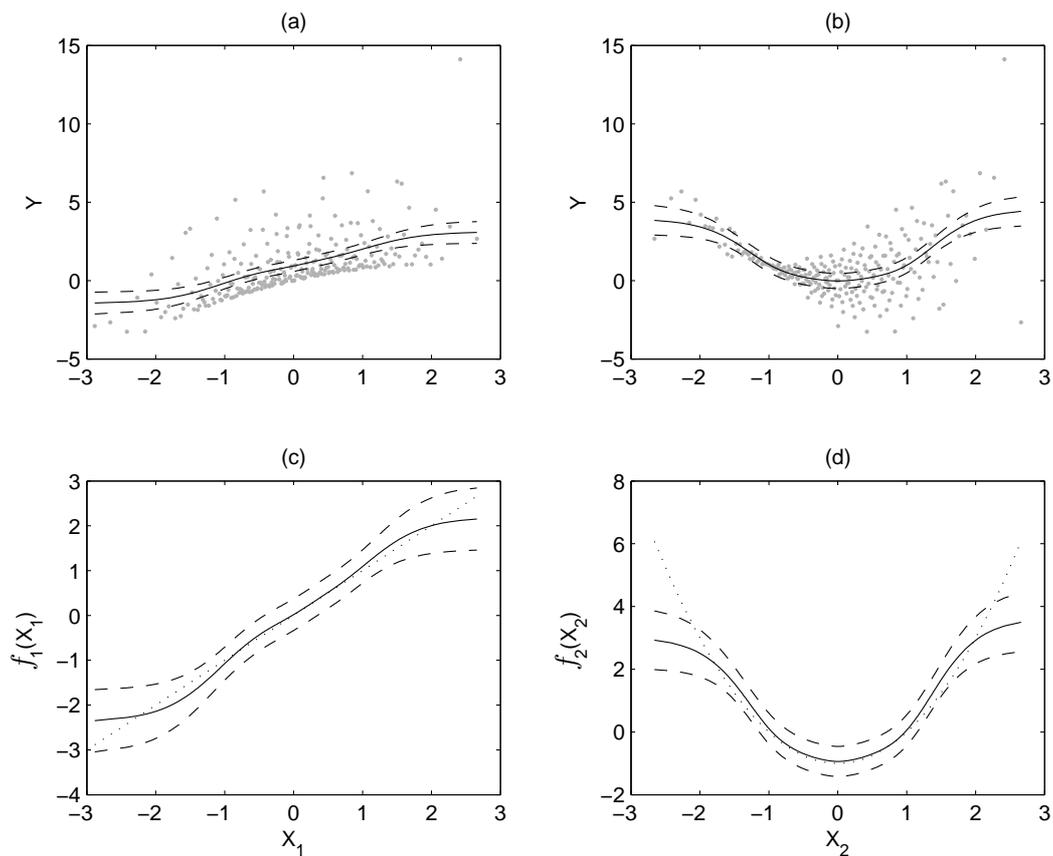


Figure 3.3: SDR estimation of the first-order HDMR of the simple model (3.2.1).

(a,b) Scatterplots of Y versus X_1 and X_2 (grey dots), with the smoothed estimates of the $f_i + f_0$ functions (solid lines) and their 95% error bands (dashed lines).

(c,d) Detail of the estimated f_i functions (solid lines) with estimated 95% error bands (dashed lines) and analytic f_i functions (dotted lines).

3.2.2 Another simple example

Let us consider the function

$$Y = X_1 \cdot X_2 + X_3 \tag{3.2.3}$$

with

$$X_i \sim U(-1, 1).$$

We perform an analysis using a Sobol' LP_τ sample of dimension 256. This model has only one non-zero main effect for X_3 and one second-order interaction term for (X_1, X_2) . In Figure 3.4 we show the SDP estimation of the first-order HDMR terms f_1 , f_2 and f_3 . We also report the standard error band, showing that only f_3 has a significant main effect.

In order to give the flavour of the generalized sorting strategy used in the SDR methodology for interaction terms, let us perform the analysis of second-order interaction term (X_1, X_2) . Since we want to compute the f_{12} interaction effect, the $2D$ sorting requires exploration of the (X_1, X_2) plane along a closed trajectory, like the one shown in Figure 3.5, with the sorting of the sample points carried out as they fall within the band delimited by two adjacent lines. This allows for the identification of an ordering in which (X_1, X_2) has low-frequency characteristics while X_3 maintains the white spectrum (Figure 3.6). The corresponding sorted output signal Y can then be analysed to identify the second-order interaction term. This is shown in Figure 3.7, where we compare the analytic values and the SDP estimates of the sorted $f_{1,2}$ interaction term.

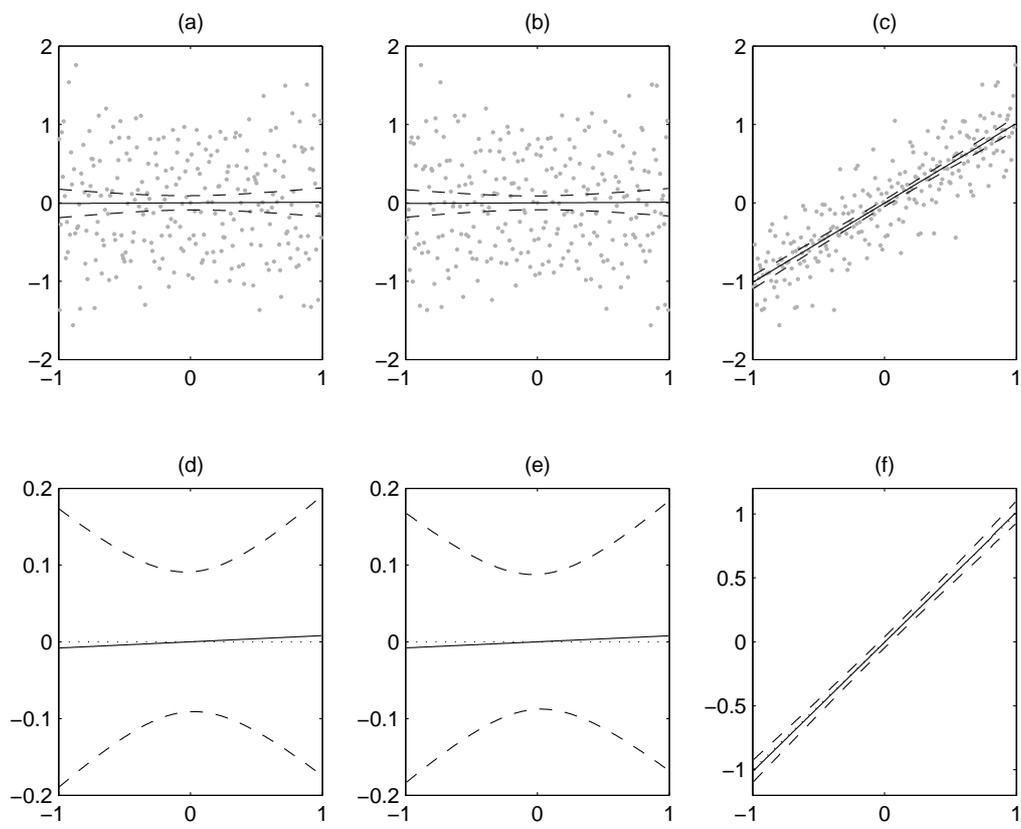


Figure 3.4: SDR estimation of the first-order HDMR of the simple model (3.2.3).
 (a,b,c) Scatter plots of Y versus X_i (grey dots), with the smoothed estimates of the $f_i + f_0$ functions (solid lines) and their 95% error bands (dashed lines).
 (d,e,f) Detail of the estimated f_i functions (solid lines) with estimated 95% error bands (dashed lines) and analytic f_i functions (dotted lines).

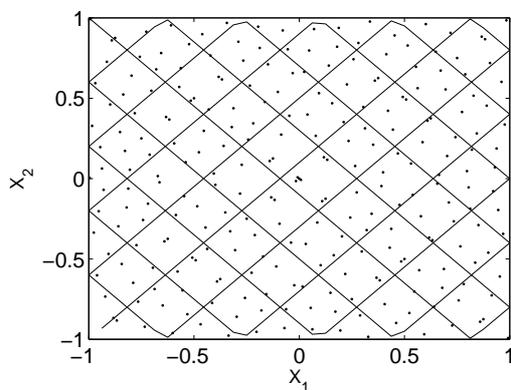


Figure 3.5: Sorting trajectory in the (X_1, X_2) plane.

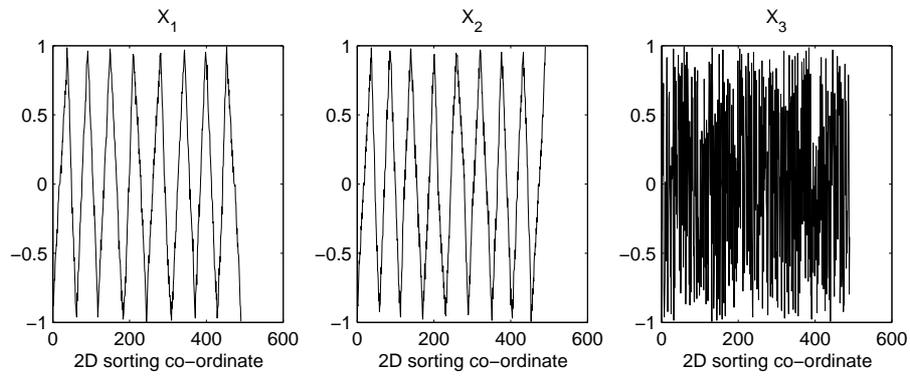


Figure 3.6: Sorted sample used to compute the (X_1, X_2) interaction.

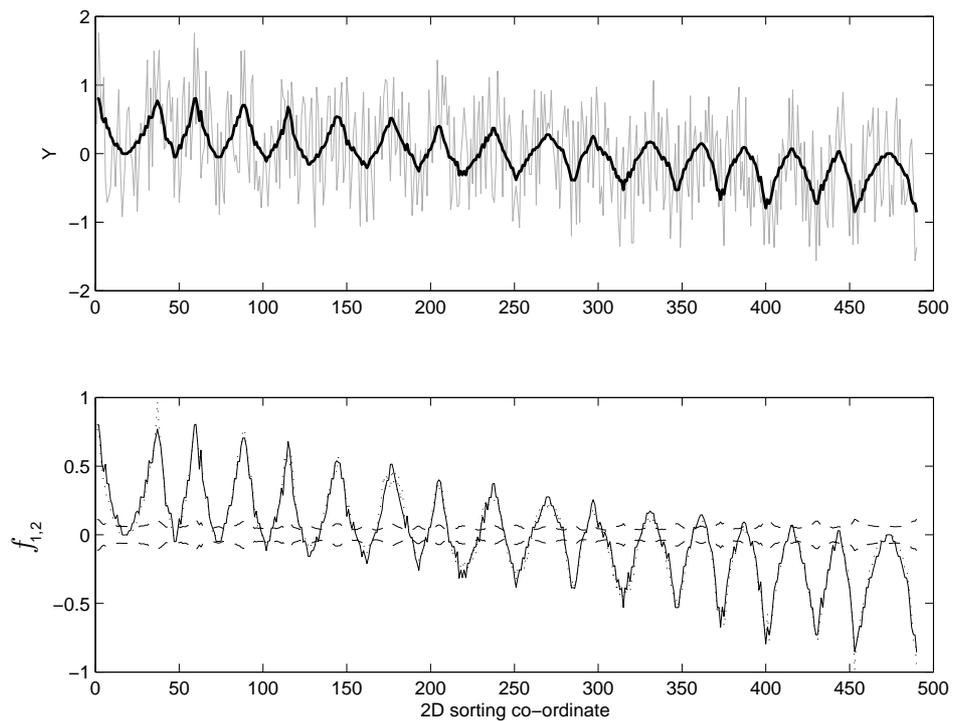


Figure 3.7: Upper panel: Sorted output signal (grey line) and smoothed low-frequency component (bold line) attributable to the (X_1, X_2) interaction. Lower panel: Smoothed low-frequency component (solid line) with standard error (dashed lines) and sorted analytic (X_1, X_2) interaction (dotted line).

3.3 Exercises

Exercise 3.3.1. Compute the analytic HDMR expansion of the function $Y = X_1 + X_2 + X_3$, with $X_i \sim N(\mu_i, \sigma_i)$. Assign values to (μ_i, σ_i) and perform the regression/smoothing analysis with the Haar wavelet and spline smoothing.

Exercise 3.3.2. Compute the analytic HDMR expansion of the Ishigami function (Ishigami and Homma, 1990):

$$Y = \sin X_1 + A \sin^2 X_2 + BX_3^4 \sin X_1$$

where $X_i \sim U(-\pi, \pi)$. Perform the regression/smoothing analysis with the Haar wavelet and spline smoothing, when $A = 7$ and $B = 0.1$.

Exercise 3.3.3. Compute the analytic HDMR expansion of the Sobol' g -function. Perform the regression/smoothing analysis with the Haar wavelet and spline smoothing, when:

- $a = [0, 1, 4.5, 9, 99, 99, 99, 99]$;
- $a = [0, 0, 1, 1, 4.5, 4.5, 9, 9, 99, 99, 99, 99, 99, 99, 99]$;
- $a = [0, 0.01, 0, 0.2, 0.3, 0, 0.5, 1, 1.5, 1.8, 3, 4.5, 8, 9, 99]$;

Exercise 3.3.4. Compute the HDMR expansion of the model $Y = \prod_{i=1}^k X_i$, with $X_i \sim U(0, MAX_i)$. Is there a clever way to map this function synthetically?

Perform the regression/smoothing analysis with the Haar wavelet and spline smoothing for the output Y and for the 'clever' transformation of Y , for $k = 3, 4$, assigning values to MAX_i .

Exercise 3.3.5. Compute the HDMR expansion of the model $Y = \prod_{i=1}^k X_i$, with $X_i \sim N(0, \sigma_i)$. Is there a clever mapping in this case as well?

Perform the regression/smoothing analysis with the Haar wavelet and spline smoothing for the output Y and for the 'clever' transformation of Y , for $k = 3, 4$, assigning values to σ_i .

3.4 Solutions

Exercise 3.3.1 First we compute the unconditional mean

$$f_0 = E(Y) = \mu_1 + \mu_2 + \mu_3.$$

Then the first-order terms $f_i = E(Y|X_i) - f_0$:

$$\begin{aligned} f_1 &= X_1 + \mu_2 + \mu_3 - f_0 = X_1 - \mu_1 \\ f_2 &= X_2 + \mu_1 + \mu_3 - f_0 = X_2 - \mu_2 \\ f_3 &= X_3 + \mu_1 + \mu_2 - f_0 = X_3 - \mu_3 \end{aligned}$$

It is easy to verify that any other term of order higher than one is null.

Exercise 3.3.2) Looking at the analytic form of the Ishigami function one can say that the decomposition will have the following terms

$$Y = f_0 + f_1(X_1) + f_2(X_2) + f_3(X_3) + f_{13}(X_1, X_3)$$

Given the uniform probability density of X_i , we have that $p(X_i) = 1/2\pi$ if $-\pi < X_i < \pi$ and $p(X_i) = 0$ elsewhere. The unconditional mean is therefore

$$\begin{aligned} f_0 &= E(Y) = \iiint_{-\pi}^{\pi} Y dX_1 dX_2 dX_3 / 8\pi^3 \\ &= \int_{-\pi}^{\pi} \sin X_1 dX_1 / 2\pi + \int_{-\pi}^{\pi} A \sin^2 X_2 dX_2 / 2\pi + \iint_{-\pi}^{\pi} B X_3^4 \sin X_1 dX_1 dX_3 / 4\pi^2 \\ &= 0 + A/2 + 0 \end{aligned}$$

The first-order terms are:

$$\begin{aligned} f_1 &= E(Y|X_1) - f_0 = \iint_{-\pi}^{\pi} Y dX_2 dX_3 / 4\pi^2 - A/2 \\ &= \sin X_1 + \int_{-\pi}^{\pi} A \sin^2 X_2 dX_2 / 2\pi + \sin X_1 \int_{-\pi}^{\pi} B \cdot X_3^4 dx_3 / 2\pi - A/2 \\ &= \sin X_1 (1 + B \cdot 2\pi^5 / 5 / (2\pi)) + A/2 - A/2 \\ &= \sin X_1 (1 + B\pi^4 / 5) \\ f_2 &= E(Y|X_2) - f_0 = \iint_{-\pi}^{\pi} Y dX_1 dX_3 / 4\pi^2 - A/2 \\ &= \int_{-\pi}^{\pi} \sin X_1 dX_1 / 2\pi + A \sin^2 X_2 + \iint_{-\pi}^{\pi} B X_3^4 \sin X_1 dX_1 dX_3 / 4\pi^2 - A/2 \\ &= 0 + A \sin^2 X_2 + 0 - A/2 \\ &= A(\sin^2 X_2 - 1/2) \\ f_3 &= E(Y|X_3) - f_0 = \iint_{-\pi}^{\pi} Y dX_1 dX_2 / 4\pi^2 - A/2 \\ &= (1 + B X_3^4) \cdot \int_{-\pi}^{\pi} \sin X_1 dX_1 / 2\pi + \int_{-\pi}^{\pi} A \sin^2 X_2 dX_2 / 2\pi - A/2 \\ &= 0 + A/2 - A/2 = 0 \end{aligned}$$

It is easy to verify that all second-order terms are null except:

$$\begin{aligned} f_{13} &= E(Y|X_1, X_3) - f_1 - f_3 - f_0 \\ &= \sin X_1 + B X_3^4 \sin X_1 + \int_{-\pi}^{\pi} A \sin^2 X_2 dX_2 / 2\pi - \sin X_1 (1 + B\pi^4 / 5) - 0 - A/2 \\ &= \sin X_1 - \sin X_1 + B X_3^4 \sin X_1 - B\pi^4 / 5 \sin X_1 + A/2 - A/2 \\ &= B(X_3^4 - \pi^4 / 5) \sin X_1 \end{aligned}$$

Exercise 3.3.3 In order to compute the HDMR expansion of the Sobol' g-function,

$$f = \prod_{i=1}^k g_i(X_i) = \prod_{i=1}^k \frac{|4X_i - 2| + a_i}{1 + a_i}$$

with $X_i \sim U(0, 1)$, it is useful to recall that:

$$\int_0^1 g_i(X_i) dX_i = \int_0^1 \frac{|4X_i - 2| + a_i}{1 + a_i} dX_i = 1$$

valid for any X_i and any a_i .

So the unconditional mean is

$$f_0 = \prod_i \int_0^1 g_i(X_i) dX_i = 1$$

The first-order terms are:

$$\begin{aligned} f_i &= E(Y|X_i) - f_0 \\ &= g_i(X_i) \prod_{j \neq i} \int_0^1 g_j(X_j) dX_j - f_0 \\ &= g_i(X_i) - 1 \end{aligned}$$

The second-order terms are:

$$\begin{aligned} f_{ij} &= E(Y|X_i, X_j) - f_i - f_j - f_0 \\ &= g_i(X_i) g_j(X_j) \prod_{l \neq (i,j)} \int_0^1 g_l(X_l) dX_l - f_i - f_j - f_0 \\ &= g_i(X_i) g_j(X_j) - (g_i(X_i) - 1) - (g_j(X_j) - 1) - 1 \\ &= g_i(X_i) g_j(X_j) - g_i(X_i) - g_j(X_j) + 1 \end{aligned}$$

All higher-order HDMR terms are obtained recursively, considering that, for any group of input factors indexed by $I = (i_1, \dots, i_l)$:

$$E(Y|X_I) = \prod_{i \in I} g_i(X_i)$$

The Sobol' g -function has non-zero interaction effects of any order.

Exercise 3.3.4 The unconditional mean is

$$\begin{aligned} f_0 &= E(Y) = \prod_{i=1}^k \int_0^{MAX_i} X_i dX_i / MAX_i \\ &= \prod_{i=1}^k MAX_i / 2 \end{aligned}$$

The first-order terms are:

$$\begin{aligned} f_i &= E(Y|X_i) - f_0 \\ &= X_i \prod_{j \neq i} \int_0^{MAX_j} X_j dX_j / MAX_j - f_0 \\ &= X_i \prod_{j \neq i} MAX_j / 2 - f_0 \\ &= (X_i - MAX_i / 2) \prod_{j \neq i} MAX_j / 2 \end{aligned}$$

The second-order terms are:

$$\begin{aligned}
 f_{ij} &= E(Y|X_i, X_j) - f_i - f_j - f_0 \\
 &= X_i \cdot X_j \prod_{l \neq (i,j)} \int_0^{MAX_l} X_l dX_l / MAX_l - f_i - f_j - f_0 \\
 &= X_i \cdot X_j \prod_{l \neq (i,j)} MAX_l / 2 - f_i - f_j - f_0 \\
 &= (X_i X_j - X_i MAX_j / 2 - X_j MAX_i / 2 - MAX_i \cdot MAX_j / 4) \prod_{l \neq (i,j)} MAX_l / 2
 \end{aligned}$$

All higher-order terms are obtained recursively, remembering that

$$E(Y|X_I) = \prod_{i \in I} X_i \prod_{j \notin I} MAX_j / 2$$

As with the g -function, this example also has non-null interaction effects of any order.

A clever way to map the function Y is to consider its log-transformed version. This is allowed, since $X_i > 0$ for all X_i . This gives:

$$\log(Y) = \sum_{i=1}^k \log(X_i)$$

which has an easy HDMR expansion of only first-order terms that describes 100% of Y !

In general, trying to analyse the log-transformed version of a function can produce interesting results. In particular, if the the log-transformation of Y is described up to a large extent (e.g. 90% of the variance of $\log(Y)$) by first-order HDMR terms, this suggests that the mapping between X_i and Y can be factorized.

Exercise 3.3.5 All the HDMR terms of this model are null except for the highest, k th order term:

$$f_{1,\dots,k} = \prod_{i=1}^k X_i = f$$

This is a very unattractive model. Even if the input factors had different σ_i values (i.e. different uncertainties), neither variance-based sensitivity indices nor HDMR analysis would be able to produce a ranking of the importance of input factors.

One possible way of analysing sensitivity for this kind of model is to apply MCF techniques and, similarly to Example 6 of Section 2.4, to try to map quantiles of Y , e.g. $Y > 1$ or $Y < -1$ (we leave this as an additional exercise for the reader) .

In metamodeling terms, however, one can still identify a clever transformation of Y , extracting useful hints about the effect of each input factor on the outcome Y .

The conditional expectation $E(Y|X_I)$ is null whatever the single input factor or group of input factors. So, Y is a pure ‘noise’ process, without any shift in the mean, and the effect of each input factor is simply to modulate the noise amplitude, similarly to heteroscedastic processes.

One could therefore conceive of trying to build a metamodel for the variance of Y . In order to do this, one could consider the transformation $\log(Y^2)$. By doing so, the function to be analysed becomes

$$\log(Y^2) = \sum_{i=1}^k \log(X_i^2)$$

which is a simple, additive mapping between the (log of the) squared input factors and the (log of the) squared output. Remembering that any expectation of Y (unconditional or not) is null, taking the square of Y is equivalent to taking its variance. Then, the main effect $f_i = E(\log(Y^2)|\log(X_i^2))$ describes effect of X_i in modulating the variance of Y , and the further X_i is from zero, the greater the variance of Y . So, the greatest impact in modulating the variance of Y will come from the input factor with the highest variance σ_i^2 . This allows for a nice ranking of the importance of input factors, which will be proportional to the magnitude of the σ_i 's.

This 'extreme' example shows that, in nasty cases where 'standard' GSA tools seem to fail in providing clear answers, functional transformations of Y can facilitate a better explanation of the model properties.

Chapter 4

Conclusions

We have discussed different methods capable of greatly accelerating the computation of the indices (in terms of reduced number of model executions), based on metamodelling. All smoothing and metamodelling techniques are based on fitting a model approximation based on a single Monte Carlo (or, better, quasi-Monte Carlo) sample. This is a major advantage with respect to ‘classical’ variance-based techniques, which require some ad hoc sampling. Moreover, metamodelling techniques usually converge much more quickly, i.e. they are more efficient. This is possible since such techniques rely on assumptions of regularity and smoothness in Y that allow us to infer the value of Y at untried points, based on the information available from MC samples at nearby points. This also implies that, in contrast to ‘classic’ variance-based estimators which rely only on square-integrability of Y , smoothing methods are not robust in the face of heavily discontinuous mappings, e.g. piecewise continuous functions, in which the values of Y jump continuously in an apparently random fashion. In this regard, the recursive approach of Ratto et al. (2006) provides elements of meaningful flexibility to adjust the estimated model approximation to local jumps and spikes.

When to use what? The decision is evidently based on a trade off between model execution cost and the analyst’s time.

- The method of Saltelli (2002) is straightforward to encode and comes at the cost of $N(k + 2)$ model executions, with $N = 1000$ or higher, to compute *both* the S_i ’s and the S_{T_i} ’s.
- The method of Tarantola et al. (2006) is also easy to encode and requires only a *single* set of N runs to compute the whole of set of S_i ’s. However, any information about interaction terms or total effects is missed.
- The approach of Ratto et al. (2004, 2006, 2007) herein discussed is, according to our experience, the most advisable meta-modelling practice in terms of model executions, robustness of estimation and flexibility of implementation. It can give a fairly precise estimate of all indices of the first order at the overall cost of 200–500 simulations, although to obtain *reliable* estimates of second and third-order interaction terms, somewhat longer samples are needed: e.g. 1000–2000. In its full formulation, it is considerably less simple to encode, even if the HP-filter formulation allows a very easy and quick analysis, at least at a preliminary stage of identification of the most relevant terms for explaining the variability of Y . Other methods mentioned here applying polynomial bases to estimate HDMR terms would also be very helpful at such a preliminary stage.

- The total effect estimation is a weak element of metamodelling. Such an estimate requires adding all the first-order and interaction terms associated with each input factor. However, the precision of the mapping of the function $f(\cdot)$ is inversely proportional to the order of interactions, while the maximum order of interactions that can be reliably estimated with metamodelling techniques can hardly exceed the order of three, if one wants to reduce to the minimum the number of model evaluations. This implies that some relevant high-order interaction might always be missed in the metamodelling exercise. In this case, performing an Elementary Effects (EE) analysis (see Saltelli et al. (2004) Campolongo et al. (2007)) would allow for a cheap and quite comprehensive assessment in terms of low-order ANOVA-HDMR terms and EE, which replaces variance-based total effects.

Ideally a model-building environment should have software for both approaches at its disposal, unless the model is truly inexpensive to run. Software developed for SDR modelling is available at

<http://sensitivity-analysis.jrc.cec.eu.int/>.

Bibliography

- Bard, Y. (1974). *Nonlinear parameter estimation*. New York: Academic Press.
- Beck, B. (1987). Water quality modelling: a review of the analysis of uncertainty. *Water Resources Research* 23, 1393–1442.
- Campolongo, F., J. Cariboni, and A. Saltelli (2007). An effective screening design for sensitivity analysis of large models. *Environmental Modelling and Software* 22, 1509–1518.
- Doksum, K. and A. Samarov (1995). Nonparametric estimation of global functionals and a measure of the explanatory power of covariates in regression. *The Annals of Statistics* 23(5), 1443–1473.
- Giona, M. and O. Paladino (1994). Bifurcation analysis and stability of controlled CSTR. *Comput. Chem. Engng* 18, 877–887.
- Grieb, T., N. Shang, R. C. Spear, S. A. Gherini, and R. A. Goldstein (1999). Examination of model uncertainty and parameter interaction in the global carbon cycling model. *Environ Intl* 25, 787–803.
- Hastie, T. and R. Tibshirani (1990). *Generalized Additive Models*. Chapman and Hall.
- Helton, J., J. Johnson, C. Sallaberry, and C. Storlie (2006). Survey of sampling-based methods for uncertainty and sensitivity analysis. *Reliability Engineering and System Safety* 91(10–11), 1175–1209.
- Hodrick, T. and E. Prescott (1980). Post-war US business cycles: an empirical investigation. Manuscript, Carnegie Mellon University.
- Hornberger, G. and R. Spear (1981). An approach to the preliminary analysis of environmental systems. *Journal of Environmental Management* 7, 7–18.
- Kalman, R. (1960). A new approach to linear filtering and prediction problems. *ASME Trans., Journal Basic Eng.* 82D, 35–45.
- Kleijnen, J. (2007a). *DASE: Design and analysis of simulation experiments*. Springer Science + Business Media. Forthcoming.
- Kleijnen, J. (2007b). Kriging metamodeling in simulation: a review. CentER discussion paper.
- Li, G., J. Hu, S.-W. Wang, P. Georgopoulos, J. Schoendorf, and H. Rabitz (2006). Random Sampling-High Dimensional Model Representation (RS-HDMR) and orthogonality of its different order component functions. *Journal of Physical Chemistry A* 110, 2474–2485.

- Li, G., W. S. W., and R. H. (2002). Practical approaches to construct RS-HDMMR component functions. *Journal of Physical Chemistry* 106, 8721–8733.
- Ng, C. and P. Young (1990). Recursive estimation and forecasting of non-stationary time series. *Journal of Forecasting* 9, 173–204.
- Oakley, J. and A. O’Hagan (2004). Probabilistic sensitivity analysis of complex models: a Bayesian approach. *J. Royal Stat. Soc. B* 66, 751–769.
- Paladino, O. and M. Ratto (2000). Robust stability and sensitivity of real controlled CSTRs. *Chemical Engineering Science* 55, 321–330.
- Paladino, O., M. Ratto, and P. Costa (1995). Chaos and chemical reactor models: sensitivity of dynamics on parameters uncertainty. *Chemical Engineering Science* 50, 3829–3833.
- Pappenberger, F., I. Iorgulescu, and K. J. Beven (2006). Sensitivity analysis based on regional splits and regression trees (SARS-RT). *Environmental Modelling and Software* 21 (7), 976–990.
- Pappenberger, F. and V. Stauch (2007). Estimating Sobol’ sensitivity with cubic splines. *Journal of Hydrology*. submitted.
- Pellegrini, L. and G. Biardi (1990). Chaotic behaviour of a controlled CSTR. *Computers Chem. Engng* 14, 1237–1247.
- Ratto, M., A. Pagano, and P. Young (2007). State dependent parameter meta-modelling and sensitivity analysis. under peer review.
- Ratto, M., S. Tarantola, A. Saltelli, and P. C. Young (2004). Accelerated estimation of sensitivity indices using State Dependent Parameter models. In K. M. Hanson and F. M. Hemez (Eds.), *Sensitivity Analysis of Model Output, Proceedings of the 4th International Conference on Sensitivity Analysis of Model Output (SAMO 2004) Santa Fe, New Mexico, March 8-11, 2004.*, pp. 61–70. <http://library.lanl.gov/ccw/samo2004/>.
- Ratto, M., S. Tarantola, A. Saltelli, and P. C. Young (2006). Improved and accelerated sensitivity analysis using State Dependent Parameter models. Technical Report EUR 22251 EN, ISBN 92-79-02036-6, Joint Research Centre, European Commission. <http://webfarm.jrc.cec.eu.int/uasa/doc/mrat/eur/EUR22251EN.pdf>.
- Saltelli, A. (2002). Making best use of model valuations to compute sensitivity indices. *Computer Physics Communications* 145, 280–297.
- Saltelli, A., S. Tarantola, F. Campolongo, and M. Ratto (2004). *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models*. John Wiley and Sons.
- Schweppe, F. (1965). Evaluation of likelihood functions for Gaussian signals. *IEEE Trans. on Information Theory* 11, 61–70.
- Spear, R. (1997). Large simulation models: calibration, uniqueness and goodness of fit. *Environmental Modelling and Software* 12, 219–228.
- Spear, R., T. Grieb, and N. Shang (1994). Factor uncertainty and interaction in complex environmental models. *Water Resources Research* 30, 3159–3169.

- Storlie, C. and J. Helton (2006). Multiple predictor smoothing methods for sensitivity analysis. SANDIA Report SAND2006-4693, Sandia Laboratories.
- Tarantola, S., D. Gatelli, and T. Mara (2006). Random balance designs for the estimation of first order global sensitivity indices. *Reliability Engineering and System Safety* 91(6), 717–727.
- Young, P. (1999a). Data-based mechanistic modelling, generalised sensitivity and dominant mode analysis. *Computer Physics Communication* 117, 113–129.
- Young, P. (1999b). Nonstationary time series analysis and forecasting. *Progress in Environmental Science* 1, 3–48.
- Young, P. (2000). Stochastic, dynamic modelling and signal processing: Time variable and state dependent parameter estimation. In W. J. F. et al. (Ed.), *Nonlinear and Nonstationary Signal Processing*, pp. 74–114. Cambridge: Cambridge University Press.
- Young, P. (2001). The identification and estimation of nonlinear stochastic systems. In A. I. M. et al. (Ed.), *Nonlinear Dynamics and Statistics*. Boston: Birkhauser.
- Young, P., S. Parkinson, and M. Lees (1996). Simplicity out of complexity: Occam’s razor revisited. *Journal of Applied Statistics* 23, 165–210.
- Young, P. and D. Pedregal (1999). Recursive and en-bloc approaches to signal extraction. *Journal of Applied Statistics* 26, 103–128.
- Young, P. C., R. C. Spear, and G. M. Hornberger (1978). Modeling badly defined systems: Some further thoughts. In *Proceedings SIMSIG Conference, Canberra*, pp. 24–32.

European Commission

EUR 21878 EN – Joint Research Centre

Title: Factor Mapping and Metamodelling

Author(s): Marco Ratto, Andrea Pagano and Peter Young

Luxembourg: Office for Official Publications of the European Communities

2007 pp.54 – 21 x 29,7 cm

EUR – Scientific and Technical Research series – ISSN 1018-5593

Abstract

In this work we present some techniques, within the realm of Global Sensitivity Analysis, which permit to address fundamental questions in term of model's understanding.

In particular we are interested in developing tools which allow to determine which factor (or group of factors) are most responsible for producing model outputs Y within or outside specified bounds ranking the importance of the various input factors in terms of their influence on the variation of Y . On the other hand, we look for representing in a direct way (graphically, analytically, etc.) the relationship between input factors $\{X_1, \dots, X_k\}$ and output Y in order to get a better understanding of the model itself.

The mission of the JRC is to provide customer-driven scientific and technical support for the conception, development, implementation and monitoring of EU policies. As a service of the European Commission, the JRC functions as a reference centre of science and technology for the Union. Close to the policy-making process, it serves the common interest of the Member States, while being independent of special interests, whether private or national.

