A Novel Approach to Quantification of Model Risk for Practitioners

Z. Krajčovičová*, P. P. Pérez Velasco, C. Vázquez

Abstract

Models continue to increase their already broad use across industry as well as their sophistication. Worldwide regulation oblige financial institutions to manage and address model risk with the same severity as any other type of risk, e.g. Federal Reserve (SR 11–7) (2011), which besides defines model risk as the potential for adverse consequences from decisions based on incorrect and misused model outputs and reports. Model risk quantification is essential not only in meeting these requirements but for institution's basic internal operative. It is however a complex task as any comprehensive quantification methodology should at least consider the data used for building the model, its mathematical foundations, the IT infrastructure, overall performance and (most importantly) usage. Besides, the current amount of models and different mathematical modelling techniques is overwhelming.

Our proposal is to define quantification of model risk as a calculation of the norm of some appropriate function that belongs to a Banach space, defined over a weighted Riemannian manifold endowed with the Fisher–Rao metric. The aim of the present contribution is twofold: Introduce a sufficiently general and sound mathematical framework to cover the aforementioned points and illustrate how a practitioner may identify the relevant abstract concepts and put them to work.

Keywords: model risk, uncertainty, Riemannian manifold, geodesics, exponential map, Fisher–Rao information metric

1. Introduction

Models are simplifying mappings of reality to serve a specific purpose aimed at applying mathematical, financial and economic theories to the available data. They focus on specific aspects of the reality and degrade or ignore the rest. There are always some areas that are deliberately left out of the model, as they are viewed as being irrelevant for the specific model purpose. As such, understanding the capabilities and limitations of the underlying assumptions is a key when dealing with a model and its outputs. Besides, not only the model limitations, but the materiality of their consequences are of interest in the management of model risk. Wrong design or implementation, misuse or inadequate knowledge regarding the model development and usage, may lead to exposing a bank to additional risks. Thus, the implied objective is to determine where models are useful 'enough' while understanding and managing their integral limitations implied by and inherent to their boundary conditions as wilful abstractions.

According to the Federal Reserve (SR 11–7) (2011) model risk is defined as

> *"[. . .] the potential for adverse consequences from decisions based on incorrect or misused model outputs and reports. Model risk can lead to financial loss, poor business and strategic decision making, or damage to bank's reputation. Model risk occurs primarily for two reasons:*
>
> - *the model may have fundamental errors and may produce inaccurate outputs when viewed against the design objective and intended business uses*
> - *the model may be used incorrectly or inappropriately"*

They state that model risk should be managed and addressed with the same severity as any other type of risk and that banks should identify the sources of model risk and assess their magnitude. Further, they emphasize that expert modelling, robust model validation and a properly justified approach are necessary elements in model risk moderation, though they are not sufficient and should not be used as an excuse for not improving models.

*Corresponding Author
Email address: zkrajcov@udc.es (Z. Krajčovičová*)

In spite of the rise of awareness of model risk and understanding its significant impact, there are no globally defined industry and market standards on its exact definition, management and quantification, even though a proper model risk management is required by regulators. Furthermore, there is a widespread lack of clarity as to what model risk management should achieve, and about which tools should be used for the purpose.

Within the finance literature, some authors have defined model risk as the uncertainty about the risk factor distribution (Gibson (2000)), the misspecified underlying model (Cont (2006)), the deviation of a model from a 'true' dynamic process (Branger and Schlag (2004)), the discrepancy relative to a benchmark model (Hull and Suo (2002)), and the inaccuracy in risk forecasting that arises from the estimation error and the use of an incorrect model (Boucher et al. (2014)). Model risk has been classified previously in all asset classes, see Morini (2011) for interest rate products and credit products, Christodoulakis and Satchell (2008) for portfolio applications, Saltelli et al. (2013) for asset backed securities, and Boucher et al. (2014) for relation to measuring marker risk.

The quantification, as an essential part of model risk management, is required for a consistent management and effective communication of model weaknesses and limitations to decision makers and users and to assess model risk in the context of the overall position of the organization. The main complication in assessing model risk arises in the context of quantifying uncertainty which stems from the selection of the mathematical techniques (e.g. focusing on fitting a normal distribution), the calibration methodology (e.g. different optimization algorithms may derive different parameter values), and from the limitations on the sample data (e.g. sparse or incomplete database). The current methods for quantifying model risk available in the literature fail to deploy a common scale to simultaneously assess the alternative sources of model risk and their impacts on usage and decision–making. Further, they are limited to very specific types of models, and at times provide a risk rating instead of risk quantification.

Model risk quantification poses many challenges stemming from the high diversity of models, the wide range of techniques, the different use of models, and a large number of model owners. Some model outputs drive decisions; other model outputs provide one source of management information among many, some outputs are further used as an inputs in other models; additionally, the model outputs may be completely overridden by expert judgement. Furthermore, model risk arises in every stage of model development lifecycle. Besides, to quantify model risk you need another model, which is again prone to model risk and thus a vicious cycle of risks gets created.

The most relevant areas of analysis for the quantification of model risk are: data and calibration, model foundation, model performance, IT infrastructure, model use, controls and governance, model sensitivity and scenario analysis. The model may be fundamentally wrong due to the errors in theoretical foundation and conceptual design that emerge from incorrect logic or assumptions, model misspecification or omission of variables. Data quality issues, inadequate sample sizes and outdated data contribute to model performance issues such as instability, inaccuracy or bias in model forecasts. Model risk also arises from inadequate controls over the model use. For example, unreported models resulting in unmeasured model risk, unapproved model use, improperly used models, and inadequate understanding of model limits or uncertainty. Wrong development and implementation in software systems, incorrect application, and unapproved on-top adjustments can result to model control issues and errors. Flawed test procedures or failure to perform consistent and comprehensive user acceptance tests can lead to material model risk. Thus, one has to realize that even appropriately applied and correctly used models can produce model errors, and that perfectly accurate model results can be misused or misapplied.

The focus of this paper is on developing a novel approach for quantifying model risk within the framework of differential geometry (Murray and Rice (1993)) and information theory (Amari et al. (1987)). In this work we introduce a measure of model risk on a statistical manifold where models are represented by a probability distribution function. Differences between models are determined by the geodesic distance (under the Fisher–Rao metric), that is a symmetric, inner metric that allows us to utilize the intrinsic structure of the manifold of densities and to respect the geometry of the space we are working on, i.e. it accounts for the non–linearities of the underlying space.

The rest of this paper is structured as follows. In section 2, we summarize basic facts about Riemannian geometry and introduce the terminology used throughout the paper. Modeling process steps and a general description of our proposed method for quantification of model risk are presented in Section 3, which is followed by a detailed discussion on the main quantification steps. Section 4 to 6 describe the construction of the neighbourhood containing material variations of the model, and the definition and construction of the weight function. The model risk measure is then defined and explained in Section 7. Finally, Section 8 provides some final conclusions and directions for future work.

2. Background on Riemannian Geometry

In this section, we summarize some basic facts about Riemannian geometry and introduce the necessary notation for the rest of the paper. The details can be found among other standard references in Amari et al. (1987), Murray and Rice (1993), Chavel (2006) and Do Carmo (1992).

In the following, we denote $\mathcal{M}$ a compact and connected manifold without boundary equipped with a Riemannian metric $< \cdot, \cdot >$ and a Riemannian connection $\bigtriangledown$, with $T_p\mathcal{M}$ the tangent space at a point $p \in \mathcal{M}$. The distance metric $d(p, q)$ between two point p and q on $\mathcal{M}$ is given by

$$d(p, q) := \inf_{\gamma} \int_a^b ||\gamma'(t)||dt$$

where γ ranges over all differentiable paths $\omega : [a, b] \to \mathcal{M}$ satisfying $\gamma(a) = p$ and $\gamma(b) = q$, and $||\gamma'||^2 = \langle \gamma', \gamma' \rangle$. The distance metric $d(\cdot, \cdot)$ turns $\mathcal{M}$ into a metric space.

A Riemannian manifold $\mathcal{M}$ is endowed with a Riemannian metric that is a correspondence which associates to each point $p \in \mathcal{M}$ an inner product $< \cdot, \cdot >_p$ on $T_p\mathcal{M}$ that varies from point to point. One natural metric on the Riemannian manifold $\mathcal{M}$ is the Fisher–Rao information metric (Rao (1945)) with entries given by

$$I_{ij}(p) = g_{ij}(p) = \mathbb{E}\Big[\frac{\partial \log(p)}{\partial x^i}\frac{\partial \log(p)}{\partial x^j}\Big] = \int p\frac{\partial \log(p)}{\partial x^i}\frac{\partial \log(p)}{\partial x^j}dx \tag{1}$$

The Fisher–Rao information metric is uniquely determined the invariance under sufficient re–parameterizations (Amari et al. (1987)). The $det\, I(p)$ represents the amount of information a sample point conveys with respect to the problem of estimating the parameter $\mathbf{x}$, and so $I(p)$ can be used to determine the dissimilarities between distributions.

Under a square–root representation, the Fisher–Rao metric becomes the standard $\mathbb{L}^2$ metric and the space of probability density functions becomes the positive orthant of the unit hypersphere in $\mathbb{L}^2$ (Lang (2012)). The square–root mapping is defined as a continuous mapping $\phi : \mathcal{M} \to \Psi$ where Ψ is the space containing the positive square–root of all possible density functions. Using this mapping, we define the square–root transform of probability density functions as $\phi(p) = \psi = \sqrt{p}$, where ψ is assumed to be non–negative to ensure uniqueness. We can compute quantities of interest in this orthant and then map them back to $\mathcal{M}$ using the inverse mapping, $\phi^{-1}(\psi) = p = \psi^2$. The space of all such functions is defined as

$$\Psi = \Big\{\psi(x)\Big| \int_{\mathcal{X}} |\psi(s)|^2 ds = 1, \forall s\ \psi(s) \geq 0, x \in \mathcal{X}\Big\}$$

where $\mathcal{X}$ is the sample space. In this case, the associated natural Hilbert space, $\mathcal{H}$, equipped with a symmetric inner product, g_{ij}, induces a spherical geometry, i.e. the sum $\sum(\sqrt{p})^2$ is equal to unity (Lang (2012)). If the density function is parametrized by the set of parameters $\theta = (\theta^1, \ldots, \theta^n)$ then for each value of θ^i we have a corresponding point on the unit sphere S in $\mathcal{H}$. In this setting the geodesics are available in closed form and can hence be computed quickly and exactly. For any two tangent vectors $v_1, v_2 \in T_\psi\Psi$, the Fisher–Rao metric is given by

$$\big\langle v_1, v_2 \big\rangle \quad = \quad \int_{\mathbb{R}} v_1(s)v_2(s)ds = \left\langle \frac{\partial\{p(\cdot, \theta)\}^{1/2}}{\partial\theta_i}, \frac{\partial\{p(\cdot, \theta)\}^{1/2}}{\partial\theta_j} \right\rangle = \frac{1}{4}g_{ij} \tag{2}$$

The geodesic on the sphere and the distance given two points ψ_1, ψ_2 are given by

$$\begin{aligned} \gamma(t) &= \cos(t||v||)\psi + \sin(t||v||)\frac{v}{||v||} \\ d(\psi_1, \psi_2) &= \cos^{-1}(\langle \psi_1, \psi_2 \rangle) \end{aligned}$$

Since the compactness of $\mathcal{M}$ implies geodesic completeness (Chavel (2006)), there exists for every $p \in \mathcal{M}$ and $v \in T_p\mathcal{M}$ an unique geodesic $\gamma : \mathbb{R} \to \mathcal{M}$ satisfying $\gamma(0) = p$ and $\gamma'(0) = v$. Moreover, the Hopf–Rinow Theorem (Hopf and Rinow (1931)) ensures that any two points $p, q \in \mathcal{M}$ can be joined by a minimal geodesic of length equal to the distance between the points, $d(p, q)$. Through the geodesic γ, one can define the exponential map $\exp_p : T_p\mathcal{M} \to \mathcal{M}$ by

$$\exp_p tv := \gamma(t), \quad \forall t \in \mathbb{R}, \forall v \in \mathcal{M}$$

The exponential map for square–root transformation (Joshi et al. (2007)) has the form

$$\exp_{\psi_i} tv := \cos\Big(||tv||_{\psi_i}\Big)\psi_i + \sin\Big(||tv||_{\psi_i}\Big)\frac{tv}{||tv||_{\psi_i}}$$

where $v \in T_{\psi_i}(\Psi)$ is a tangent vector at ψ_i and $||v||^2_{\psi_i} = \big\langle v, v\big\rangle_{\psi_i}$. The inverse exponential map from ψ_i to ψ_j is given by

$$\exp^{-1}_{\psi_i}(\psi_j) := \frac{u}{\sin(u)}\Big(\psi_j - \cos(u)\psi_i\Big). \tag{3}$$

Next, we turn to the notions of cut point and cut locus. For every vector $v \in T_p\mathcal{M}$ we define

$$F(v) := \sup_{t>0}\{tv \in T_p\mathcal{M} : d(p, \gamma_v(t)) = t\}$$

as the maximal distance in direction v for which $\exp_p$ is isometric. The point $\gamma(F(v)) = \exp_p(F(v)v)$ is referred to as the cut point of p along the geodesic $\gamma(t)$ ([20]). Let $\mathcal{S}(p, 1)$ be a unit sphere on tangent space, the largest open subset of $\mathcal{M}$ in which the exponential map of a point p is diffeomorphic is $D_p = \exp_p(V(p))$, where $V(p) = \{tv : 0 \leq t < F(v), v \in \mathcal{S}(p, 1)\}$. The cut locus of p is then defined as $C_p = \exp_p\{F(v)v : v \in \mathcal{S}(p, 1), F(v)\text{is finite}\}$. Note that $\mathcal{M} = D_p \cup C_p$.

An open set $U \subset \mathcal{M}$ is said to be a normal neighbourhood of $p_0 \in U$, if $\exp_{p_0}$ is a diffeomorphism on a neighbourhood V of the origin of $T_{p_0}\mathcal{M}$ onto U, with V such that $tv \in V$ for $0 \leq t \leq 1$, if $v \in V$.

3. Modeling Process Steps and Quantification of Model Risk

Model risk cannot easily be spoken of as a single entity. There are different types and aspects of model risk that tend to easily overlap, co–occur, or co–vary. However, by viewing model risk in terms of a process of steps from design to implementation we can gain an understanding of particular areas where more attention is needed in the process of designing, validating, and implementing models. In this context, we propose four rough model creation steps: Data, Calibration, Model Selection and Testing, and Implementation and Usage. This may occur in an iterative fashion, but they result in a general linear flow that ends with institutional use (implementation and maintenance) to direct decision making (often encoded into an IT system). Limitations in any of these areas can impair reliance on model results.

1. **Data** refers to the definition of the purpose for modeling, the specification of the modeling scope, human and financial resources, the specification of data and other prior knowledge, their interpretation and preparation. The data may be obtained from both internal and external sources, and they are further prepared by cleaning and reshaping into a readily usable form for performing next steps of model development. Model risk may arise from data deficiencies in terms of both quality and availability, including, among others, error in data definition, lack of historical depth, lack of critical variables, insufficient sample, data migration, inaccurate proxies, high sensitivity of expert judgments, or wrong interpretation.
2. **Calibration** includes the selection of the types of variables and the nature of their treatment, the adjustments and tuning of free parameters to fit the model outputs, as well as links between system components and processes. Model risk or estimation uncertainty may occur due to errors in simplifications, approximations, flawed and inadequate assumptions, inappropriate calibration due to lack of data, wrong selection of some relevant subset, errors in statistical estimation or in market benchmarks which may lead to the decay in predictive power between re–estimations, limitations of the optimization algorithm, complexity of the problem, computational limitations, or use of unobservable parameters, among many others.
3. **Model Selection and Testing** involves the choice of estimation performance criteria and techniques, the identification of model structure and parameters, which is generally an iterative process with the underlying aim to balance sensitivity to system variables against complexity of representation. Further, it is related to the conditional verification which includes testing the sensitivity to changes in the data and to possible deviations of the data and system from the initial assumptions. In this step, model risk stems from, e.g., inadequate and incorrect modeling assumptions, outdated model due to parameter decalibration, model instability, use of new methodology not supported by academic research or not well–known methodologies, model misspecification, insufficient analytical capabilities, deviance from modeled process, or lack of market consensus on the model's functional form.
4. **Implementation and Usage** refers to the deployment of the model into production which is followed by a regular maintenance and monitoring. Sources of model risk in this step include using the model for unintended purposes, lack of re–estimation and re–calibration, IT failures, lack of communication between modelers and users, lack of understanding on model limitations, and differences between regulatory and management practices.

Certain risk factors may be mitigated at source by proper validation process, internal audit, monitoring, controls, skilled employees, underwriting, operational practices or by application of an expert judgment, they can never be eliminated.

Quantification of model risk, from a best practice perspective, should be quick and reliable, without refitting or building models, without reference to particular structure and methodologies, and with prioritizing analysis (getting immediate assurance on shifts that are immaterial). The base for such an approach offers the rich framework of differential geometry and information theory. In this framework, a model is represented by a particular probability distribution, $p : \mathcal{X} \to \mathbb{R}_+$ that belongs to the set of probability measures $\mathcal{M}$, so called statistical manifold, available for modelling. The manifold $\mathcal{M}$ can be further equipped with the information–theoretic geometric structure that, among other things, allows us to quantify variations and dissimilarities between probability distribution functions (models).

The set of possible probability measures may be further parametrized in a canonical way by a parameter space Θ, i.e. $\mathcal{M} = \{p(x;\theta) : \theta \in \Theta\}$. This set forms a smooth Riemannian manifold $\mathcal{M}$, so called model manifold, embedded in the space of all distributions. Every distribution is a point in this space, and the collection of points created by varying the parameters of the model, $p \in \mathcal{M}$, gives rise to a hypersurface (a parametric family of distributions) in which similar distributions are mapped to nearby points. The model manifold can be either a low dimensional parametric space as in parametric statistics, or the space of all possible models as in non–parametric statistics. The natural Riemannian metric is shown to be the Fisher–Rao metric (Rao (1945)) which is the unique intrinsic metric on the statistical manifold. As aforementioned, it is the only metric that is invariant under re–parametrization (i.e., coordinate transforms) (Amari et al. (1987)).

Let us consider a given model p_0. We assume that the examined model can be uniquely parametrized[1] using the n–dimensional vector parameter $\theta_0 = (\theta_0^1, \ldots, \theta_0^n)$ over the sample space $\mathcal{X}$ and can be described by the probability distribution $p_0 = p(x;\theta_0)$ with respect to the Lebesgue measure, i.e. p_0 belongs to the n–dimensional statistical manifold[2]. This probability distribution belongs to a set (family) of distributions $\mathcal{M} = \{p(x;\theta) : \theta \in \Theta \subset \mathbb{R}^n\}$ that forms a model manifold. We assume that for each $x \in \mathcal{X}$ the function $\theta \mapsto p(x;\theta)$ is C^∞. Thus, $\mathcal{M}$ forms a differentiable manifold and we can identify models in the family with points on this manifold. Since the set of all density functions on $\mathcal{X}$ is a subset of the L^1 space of functions in $\mathcal{X}$, $\mathcal{M}$ is considered to be a subset of the L^1 space. Choosing a particular model is the same as fixing a parameter setting $\theta \in \Theta$.

Example

To help fix ideas, we introduce an illustrative simple example and develop it further throughout the paper. Let $\mathbb{X}$ denote a vector of profit and loss, $P\&L$, over a two year time horizon (520 days) that is used to calculate the Value at Risk (VaR). VaR is derived from a distribution of $P\&L$ as the quantile loss at the portfolio level and is defined by

$$\mathbb{P}(\mathbb{X} \leq VaR) = 1 - \beta$$

where β is the confidence level, set to 99.9%. The output of the VaR is then further used either for risk measurement or decision making in risk management. Assume that the given model considers $\mathbb{X}$ to be normally distributed $p_0 = \mathcal{N}(\mu_0, \sigma_0)$ with parameters $\mu_0 = 2, \sigma_0 = 10$ once calibrated. This model belongs to a family of normal distributions that forms a differentiable manifold $\mathcal{M} = \{p(x;\mu,\sigma) : \mu \in \mathbb{R}, \sigma > 0\}$ where μ is the mean and σ is the standard deviation. Every point $p \in \mathcal{M}$ corresponds to a normal distribution $p(x,\theta)$ with $\theta = (\mu, \sigma)$.

In our univariate normally distributed case parametrized by a 2–dimensional space, $\theta = (\mu, \sigma)$, the Riemannian

[1] In case of non–parametric probability distributions we can construct a unique vector of canonical parameters $\theta = (\theta^1, \ldots, \theta^n)$ via invertible mappings $\theta^i = \log(p_i/p_0), \forall i$, where the chosen model p_0 stands for the reference probability and p_i are other models in the manifold. These canonical parameters may stand as a global coordinate system for $\mathcal{P}$.

[2] Note that this does not impose any restrictions. For a general non–statistical manifold, e.g. a model family $\mathcal{N}$ parametrized by $\psi = (\psi^1, \ldots, \psi^d)$ but not represented by a probability distribution, one needs to construct mappings from $\mathcal{N}$ to $\mathcal{M}$. Following such a mapping, any tangent vector $\frac{\partial}{\partial \psi^i}$ of $\mathcal{N}$ can be pushed forward to $T_p\mathcal{M}$ as $\frac{\partial}{\partial \psi^i} = \sum_{k=1}^n \left(\frac{\partial \theta^k}{\partial \psi_i} \frac{\partial}{\partial \theta^k} \right)$. In this way, the inner product $\left\langle \frac{\partial}{\partial \psi^i}, \frac{\partial}{\partial \psi^j} \right\rangle$ can be used to define a Riemannian metric.

matrix defined by 2 is given by

$$I = [I_{ij}(\mu, \sigma)] = \begin{bmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{bmatrix} = \begin{bmatrix} 0.01 & 0 \\ 0 & 0.02 \end{bmatrix}.$$

□

We define the model risk for a given model p_0 at the scale of an open neighbourhood around p_0 that contains alternative models that are not too far in a sense quantified by the relevance to (missing) properties and limitations of the model (i.e., the uncertainty connected with the model selection). The model risk is then measured with respect to all models inside this neighbourhood as a norm of an appropriate function of the output differences over a weighted Riemannian manifold endowed with the Fisher–Rao metric and the Levi–Civita connection[3]. The analysis consists of five steps:

1. Embedding of the model manifold into one that considers missing properties[4] in the given model p_0.
2. Choosing a proper neighbourhood around the given model.
3. Choosing an appropriate weight function, that assigns relative relevance to the different models inside the neighbourhood.
4. Calculating the measure of model risk with respect to all models inside the neighbourhood, through the corresponding norm.
5. Interpretation of the measure with respect to the specific use of the model risk quantification.

Each step addresses and aligns different limitations of the model and the uncertainty in various areas related to the model[5]. In the following sections we further develop these steps and describe the intuition behind.

4. Neighbourhood Around the Model

Recall that the given model p_0 belongs to a n–dimensional manifold $\mathcal{M}$ where each dimension represents different piece of information inherited in p_0. To consider missing properties, the uncertainty surrounding the data and the calibration, the additional information about the limitations of the model, or wrong underlying assumptions, we may need to adjoin a new dimensions to $\mathcal{M}$, and thus, consider a higher–dimensional space within which $\mathcal{M}$ is embedded for further analysis.

The proper neighbourhood around p_0 we define by the help of the tangent space $T_{p_0}\mathcal{M}$ at a point p_0. $T_{p_0}\mathcal{M}$ is a vector space that describes a first order approximation, infinitesimal displacements or deformations on the manifold in the position of the point p_0. Intuitively, it contains all possible directions of the qualitative and quantitative variations of the given model p_0 within the selected manifold as a response to new information or changes in the underlying assumptions. From a practical point of view, not all perturbations are relevant, e.g., variations that are not supported by the market, business or past experience. Thus, taking into account the materiality with respect to the intended purpose of the model, its usage, business and market, we consider only a small subset of the tangent space.

We define the open set, $\mathcal{U}$, around p_0 as a subset of some normal neighbourhood V such that

$$\mathcal{U} := \{tv \in V \subset T_{p_0}\mathcal{M} : 0 < t \leq \alpha(v) \leq inj_{p_0},\, v \in S_{p_0}\mathcal{M}\}$$

This set provides a class of directions of all relevant perturbations of the model p_0 up to a certain level $\alpha(v)$. The level $\alpha(v)$ depends on the tangent vectors, since the degree of our uncertainty on p_0 might not be constant across the canonical parameter space; for instance we could assume more uncertainty in the tails of the distribution p_0 than in its body. We can interpret $\alpha(v)$ as a means to control uncertainty regarding the choice of the model p_0, and it is appropriately chosen based on the intended purpose of the model, usage and business. Besides, the optimal level of $\alpha(v)$ depends on the

[3]The Levi–Civita connection parallely transports tangent vectors defined at one point to another and is compatible with the geometry induced by the Riemannian metric (Amari et al. (1987)). Additionally, for this choice of connection, the shortest paths are geodesics.

[4]Or properties not appropriately modelled, for which there is no consensus, cannot be adequately calibrated, among many others.

[5]Such as data, calibration, model selection, model performance, model sensitivity and scenario analysis, and most importantly the usage of the model

uncertainty surrounding the nature of the data and calibration.

Since $\mathcal{U}$ is a subset of the normal neighbourhood around p_0, the exponential map is well defined and we can construct a corresponding set of models close enough to p_0. The subset of $\mathcal{M}$ containing these models is defined as

$$U := \exp_{p_0}(\mathcal{U}) = \{p \in \mathcal{M} : d(p_0, p) \leq \alpha(v)\}$$

From now on, we shall require the boundary $\partial U = \{\alpha(v)v | v \in S_{p_0}\mathcal{M}\}$ to be continuous and piecewise regular. Moreover, U shall be a *star–shaped set with respect to* p_0 that is defined as follows:

Definition 1. *A compact subset U of a Riemannian manifold $\mathcal{M}$ is called star–shaped with respect to $p_0 \in U$ if $\forall p \in U, p \neq p_0$ there exists a minimizing geodesic γ with $\gamma(0) = p_0$ and $\gamma(T_p) = p$ such that $\gamma(t) \in U$ for all $t \in [0, T_p]$, where $T_p > 0$.*

One advantage of the exponential map in this setting is that we can avoid calibration of different alternative models inside U. For each unit vector $v \in \mathcal{U}$ there exists a unique geodesic connecting points on the boundary of U with the point p_0. This geodesic is given by $\gamma(t) = \exp_{p_0}(t|v|)$ for $t \in [0, \alpha(v)]$.

Example

Demonstrating that the model is suitable for intended purpose is a critical part of the analysis of model risk. For illustrative purposes, we focus in our example only on the uncertainty about the selected statistical distribution. We want to evaluate the impact of relaxing the assumption of symmetry for the underlying $P\&L$ distribution, i.e. the impact of not including the skew in the model. Hence, we embed the model manifold $\mathcal{M}$ into a larger manifold of skew–normal distributions, $\bar{\mathcal{M}} = \{p(x; \mu, \sigma, s) : \mu \in \mathbb{R}, \sigma > 0, s \in \mathbb{R}\}$, where μ is the location, σ the scale parameter, and s the shape parameter (Azzalini (1985)). Note, that for $s = 0$, we re–obtain the initial normal distribution, $\mathcal{N}(\mu, \sigma)$. The skew normal distribution family emerges to take into account the skewness property.

After considering various time windows, data sequences, fittings and estimates, we determine the neighbourhood of our model to be the geodesic connecting the base model, $p_0 = \mathcal{N}(2, 10) = \mathcal{SN}(2, 10, 0)$, and the skew–normal distribution, $p_1 = \mathcal{SN}(\mu, \sigma, s)$, with parameters $\mu = 1.95, \sigma = 9.98$, and $s = 2$. The geodesic distance between these two distributions is $d(\sqrt{p_0}, \sqrt{p_1}) = 0.6809$. To form the neighbourhood, we first construct the related perturbation tangent vector associated with the directions to the boundary point, $\sqrt{p_1}$, using the inverse exponential map defined by 3

$$v_{p_1} \quad = \quad \exp^{-1}_{\sqrt{p_0}}\left(\sqrt{p_1}\right) = \Big[\frac{0.6809}{\sin(0.6809))}\Big(\sqrt{p_1} - \cos\left(0.6809\right)\Big)\sqrt{p_0}\Big)\Big]$$

This provides a class of variations of the initial model by moving away from it in the direction v_{p_1} which determines the whole neighbourhood U given by

$$U = \{\gamma(t) = (\exp_{\sqrt{p_0}}(tv_{p_1}))^2; t \in [0, 1]\}$$

with boundaries $\partial_1 U = \{p_0\}$ and $\partial_2 U = \{p_1\}$ Thus, by varying t from 0 to 1, one traces the geodesic path from p_0 to p_1, and we obtain a set of all distributions in the directions p_1. The neighbourhood U around p_0 includes all distributions on the geodesic γ for $t \in [0, 1]$.

□

5. Weight Function Definition

Variations of the chosen model are not equally material and they all might take place with different probabilities. By placing a non–linear weight function (kernel), K, over the set U we can easily place relative relevance to each alternative model, and assign the credibility of the underlying assumptions that would make alternative models partially or relatively preferable to the nominal one p_0. The particular choice of the structure of the kernel depends on various factors, such as usage of the model, distance from p_0, or sensitivity to different changes. Additionally, it may as well be influenced by the importance of the outcomes with connection to decisions making, by the business, or by its intended purpose.

In what follows, we define a general weight function K, and show that under certain conditions it is well defined and unique. In general, we consider K to be a non–negative and continuous function that depends on the local geometry of $\mathcal{M}$ by incorporating a Riemannian volume[6] associated to the Fisher–Rao information metric given by $dv(p) = \sqrt{det(I(\theta))}d\theta$. The volume measure is the unique Borel measure on $\mathcal{M}$ (Federer (2014)). With respect to a coordinate system, the information density of p represents the amount of information the single model possesses with respect to the parameters. For example, a small $dv(p)$ means that the model contains much uncertainty and requires many observations to learn.

As the underlying factors[7] that influence the perturbations of the given model happen with some likelihood, we treat all models inside $\mathcal{M}$ as random objects. As a consequence, we require K to be a probability distribution with respect to the Riemannian volume, i.e. $\int_{\mathcal{M}} K dv(p) = 1$. With this assumption we prescribe the corresponding weights but also our belief about the likelihood of the models. Additionally, we state that the right model does not exist and that the choice of p_0 was to some extent a subjective preference.

Definition 2. *An admissible weight function K defined on $\mathcal{M}$ satisfies the following properties:*

($K1'$) *K is continuous on $\mathcal{M}$*

($K2'$) *$K \geq 0$ for all $p \in \mathcal{M}$*

($K3$) $\int_{\mathcal{M}} K dv(p) = 1$

Recall that to compute the n–dimensional volume of the objects in $\mathcal{M}$, one considers a metric tensor on the tangent space $T_p\mathcal{M}$ at the point $p \in \mathcal{M}$. In particular, the Fisher–Rao information metric I on $\mathcal{M}$ maps each point $p \in \mathcal{M}$ to a volume $dv(p)$ which is a symmetric and bilinear form that further defines an n–dimensional volume measure on any measurable subset $U \subset \mathcal{M}$ by $Vol(U) := \int_U dv(p)$. A smooth probability distribution K over $\mathcal{M}$ with respect to the Riemannian measure induces a new absolutely continuous probability measure ζ with respect to Vol, that is,

$$\zeta(U) = \int_U d\zeta = \int_U K dv(p) \tag{4}$$

for all measurable $U \subset \mathcal{M}$, and $\zeta(\mathcal{M}) = 1$. The pair $(\mathcal{M}, \zeta)$ is then called a weighted manifold, or a Riemannian metric–measure space and is proved to be a nontrivial generalization of Riemannian manifolds (Morgan (2005)).

The weight function K of the Definition 2 represents a general characterization of a probability distribution over the Riemannian manifold $\mathcal{M}$. To tune K for proper analysis of model risk, we need to impose additional properties which are connected with the specific uncertainties surrounding the given model.

From a practitioner point of view, the materiality of perturbations of the given model is bounded. Models that do not belong to the chosen neighbourhood U are not relevant from the perspective of model risk, and so do not add any uncertainty. Therefore, we assume the weight function to be non–negative only over the neighbourhood U and zero elsewhere. Moreover, translation of the changes in various underlying assumptions, data or calibration into the changes in output and further usage of the model are going to vary with respect to the direction of the change. Hence, we require K to be continuous along the geodesic curves γ uniquely determined by $v \in \mathcal{S}(p_0, 1) \subset T_{p_0}U$ starting at p_0 and ending at the points on ∂U. These additional properties are a modified property $(K1')$ and $(K2')$ defined as follows:

($K1$) *K is continuous along all geodesics γ starting at p_0 for all unit vectors on $\mathcal{S}(p_0, 1)$*

($K2$) *$K > 0$ $\forall p \in U \backslash \{\partial U\}$ and $K \geq 0$ $\forall p \in \partial U$, and $K = 0$ $\forall p \in \mathcal{M} \backslash \{U\}$*

The weight function satisfying properties $(K1) - (K3)$ takes into consideration and is adjusted according to the different directions of the changes, i.e. prescribes different sensitivities to different underlying factors, just as to their size, i.e. the relevance of the change in comparison to the other possible changes, based on the adequate knowledge of the practitioner about the limitations and weaknesses of the model, market and business.

[6]The Riemannian volume plays the same role as the Lebesgue measure in $\mathbb{R}^n$. The difference is that in the case of a Riemannian manifold, the measure is now different at each point since the local expression of the metric is changing. Intuitively, the induced manifold measure represents the notion of uniformity according to the chosen Riemannian metric.

[7]For example the uncertainty surrounding data, calibration or model selection.

6. Weight Function Construction

The construction of a weight function on a given Riemannian manifold is technically very difficult since it requires precise knowledge of the intrinsic geometry and the structure of the manifold. The challenge here is to gain the spatial intuition of the geometrical properties of the abstract manifold in order to employ the particular properties of the weight function for the different directions and sizes of the material perturbations.

To determine a weight function K that satisfies all of the required properties and in order to overcome this difficulty we introduce a continuous mapping from a manifold endowed with an Euclidean geometry to the model manifold endowed with a Riemannian geometry that preserves the local properties. Euclidean geometry is well understood and intuitive, and thus a construction of a function on this space is considerably easier and more intuitive. In total, we construct three mappings: the exponential map $\exp_{p_0}$, the polar transform P and a further coordinate transform Λ_ρ.

Every Riemannian manifold $\mathcal{M}$ is locally diffeomorphic to the Euclidean space $\mathbb{R}^n$, and so in a small neighbourhood of any point the geometry of $\mathcal{M}$ is approximately Euclidean. All inner product spaces of the same dimension are isometric, therefore, all the tangent spaces $T_p\mathcal{M}$ on a Riemannian manifold $\mathcal{M}$ are isometric to the n–dimensional Euclidean space $\mathbb{R}^n$ endowed with its canonical inner product. Hence, all Riemannian manifolds have the same infinitesimal structure not only as manifolds but also as Riemannian manifolds.

The weight function is defined with respect to the neighbourhood U and is continuous on the geodesic curves γ connecting p_0 to the points on the boundary ∂U. All material perturbations, i.e. alternative models inside U, are uniquely described by the distances from the base model p_0 and by the vectors tangent to the unique geodesics γ that pass through them. To maintain these properties, we consider an n–dimensional cylinder $C^n = [0,1]\times\mathbb{S}^{n-1} = \{(t,\nu) : t \in [0,1], \nu \in \mathbb{S}^{n-1}\} \subset \mathbb{R}^{n+1}$, where the parameter t stands for the normalized distance of geodesics, and where $\mathbb{S}^{n-1}$ denotes the $(n-1)$–dimensional unit sphere on $\mathbb{R}^n$ containing all the unit tangent vectors of $T_{p_0}\mathcal{M}$. Then the boundaries of C^n, defined by

$$\partial_1 C^n = \{(0,\nu) : \nu \in \mathbb{S}^{n-1}\}, \quad \partial_2 C^n = \{(1,\nu) : \nu \in \mathbb{S}^{n-1}\},$$

represent the end points of the geodesics, i.e. $\partial_1 C^n$ will be constructed to the point p_0 and $\partial_2 C^n$ will be transformed to become ∂U.

The Riemannian structure on C^n is given by the restriction of the Euclidean metric in $\mathbb{R}^{n+1}$ to C^n. Hence, C^n is a compact smooth Riemannian manifold with a canonical measure given by the product measure $dt \times d\nu$. This manifold allows us to construct an appropriate function on C^n, and then using a proper transformation to obtain a weight function satisfying all required properties $(K1) - (K3)$. Besides, it gives us the freedom to choose a suitable function based on the uncertainty surrounding data and other input information, the calibration, the model selection and the model intended purpose.

As a first step to obtain a mapping from C^n to $\mathcal{M}$, we consider the exponential map from the tangent space at the point p_0 onto the neighbourhood U (see Section 2 for definition). Since U is compact and, hence, topologically complete, the geodesic γ can be defined on the whole real line $\mathbb{R}$ (Hopf and Rinow (1931)). Thus, the exponential map is well–defined on the whole tangent space $T_{p_0}\mathcal{M}$. Further, since U is a subset of the normal neighbourhood of p_0, the exponential map defines a local diffeomorphism from $T_{p_0}U$ to U. Then the geodesics γ are given in these coordinates by rays emanating from the origin.

Example

The weight function is constructed with respect to the neighbourhood U that in our example represents the geodesic γ with boundary points p_0 and p_1. We parametrize γ by $t \in [0,1]$, and define the one-dimensional cylinder as, i.e. $C^1 = [0,1]\times\mathbb{S}$ with boundaries $\partial_1 C = \{(0,\nu)\}$ and $\partial_2 C = \{(1,\nu)\}$ where $\mathbb{S} = \{\nu \in \mathbb{R}^n : ||\nu||^2 = \nu^2 = 1\}$. The left boundary $\partial_1 C^1$ will be contracted to the given model p_0 and $\partial_2 C^1$ to p_1.

The construction of the weight function reduces to the construction of the weight function on the real line on the interval $[0,1]$.

□

Next, we introduce a polar coordinate transformation on $T_{p_0}U$. For the sake of distinction, we denote by $\mathcal{S}(p_0,1)$ the $(n-1)$–dimensional unit sphere in $T_{p_0}U$, and by $\mathbb{S}^{n-1}$ the unit sphere in $\mathbb{R}^n$. We define the transformation and its

inverse as

$$P:\ [0,\infty)\times\mathcal{S}(p_0,1)\to T_{p_0}\mathcal{M}:\ \ P(t,v)=tv;$$
$$P^{-1}:\ T_{p_0}\mathcal{M}\backslash\{0\}\to(0,\infty)\times\mathcal{S}(p_0,1):\ \ P^{-1}(v)=\Big(||v||,\frac{v}{||v||}\Big)$$

In order to precisely describe the neighbourhood U, we define $\rho(v)$ as the length $d(p_0,p)\geq 0$ of the geodesic γ connecting the point p_0 with a boundary point $p\in\partial U$ in the direction $v\in\mathcal{S}(p_0,1)$. The distance ρ, considered as a real valued function on $\mathcal{S}(p_0,1)$, is strictly positive and Lipschitz continuous on $\mathcal{S}(p_0,1)$. Using the distance function, we define the coordinate transformation

$$\Lambda_\rho:(t,v)\to\Big(\rho(v)t,v\Big).$$

The mapping $\mathbb{S}^{n-1}\mapsto\mathcal{S}(p_0,1)$, $\nu\mapsto v$ is well defined in the sense that there exists a canonical identification between a unit vector $\nu\in\mathbb{S}^{n-1}\subset\mathbb{R}^n$ and the element $v\in\mathcal{S}(p_0,1)\subset T_{p_0}U$. Since the distance function $v\mapsto\rho(v)$ is strictly positive and Lipschitz continuous on $\mathcal{S}(p_0,1)$, so is the inverse $v\mapsto\dfrac{1}{\rho(v)}$. Therefore, the mapping Λ_ρ defines a bi–Lipschitz mapping from $[0,1]\times\mathbb{S}^{n-1}$ onto the subset $[0,\rho(v)]\times\mathcal{S}(p_0,1)$.

Then the composition $\exp_{p_0}P\Lambda_\rho$ defines a mapping from C^n onto U that maps ∂_1C^n onto the point $\{p_0\}$ and the right hand side boundary onto ∂U. Moreover, it preserves continuity for any continuous function h defined on C^n that satisfies the following consistency condition:

Definition 3. *A continuous function h defined on a cylinder C^n of $\mathbb{R}^n$ is called consistent with a continuous function f on U under the mapping $\exp_{p_0}P\Lambda_\rho$ if $h(t,\nu)=\Lambda_\rho^{-1}f^{-1}(t,\nu)$ for all $(t,\nu)\in C^n$. In this case, h satisfies the following conditions:*

$$\begin{aligned}(i)\ h(0,\nu_1)&=h(0,\nu_2)\ \ \forall\nu_1,\nu_2\in\mathbb{S}^{n-1}\\ (ii)\ h(1,\nu_1)&=h(1,\nu_2)\ \ \textit{if}\ \ \exp_{p_0}P\Lambda_\rho(1,\nu_1)=\exp_{p_0}P\Lambda_\rho(1,\nu_2)\ \textit{on}\ \mathcal{M}\end{aligned}$$

The first condition (i) implies that h is constant on the boundary ∂_1C^n. Then when the function h on C^n is consistent with f, the constant value at ∂_1C^n corresponds exactly with the value $f(p_0)$. The second condition ensures compatibility of function h with function f at the points of the boundary ∂U, i.e. if $\exp_p P\Lambda_\rho$ maps two different points $(1,\nu_1)$ and $(1,\nu_2)$ in C^n onto the same point $p\in\partial U$, then $h(1,\nu_1)=h(1,\nu_2)=f(p)$.

Lemma 1. *The weight function K satisfying assumptions $(K1)-(K3)$ is equivalent to assuming that a consistent function $h(t,\nu)$ defined on C^n with codomain $\mathbb{R}^n$ satisfies the following properties:*

(H1) *$h(t,\nu)$ is a continuous function on the compact manifold C^n*

(H2) *$h(t,\nu)\geq 0,\ \ (t,\nu)\in[0,1)\times\mathbb{S}^{n-1}$*

(H3) *$h(1,\nu)=\kappa(\nu)$ for all $\nu\in\mathbb{S}^{n-1}$, where κ is some non–negative function of ν*

(H4) *$h(0,\nu_1)=h(0,\nu_2)=const.$ for all $\nu_1,\nu_2\in\mathbb{S}^{n-1}$*

(H5) *$\int_{C^n}h(t,\nu)d\nu=1$, where $d\nu=dt\times d\mu$*

Proof. See Appendix. Using this result, the construction of the weight function becomes easier and more intuitive. One chooses the appropriate function h defined on C^n with respect to the particular model and the uncertainty surrounding it. Then, applying the above transformation one obtains an appropriate weight function K defined on U satisfying properties $(K1)-(K3)$ relevant for model risk analysis. Besides, for a chosen function h the weight function K is unique and well defined.

Theorem 4. *A continuous function h defined on C^n satisfying conditions $(H1)-(H5)$, determines a unique and well defined weight function K on U given by*

$$K(p,t)=\frac{1}{\eta_{p_0}(p)}t^{1-n}\rho(v)^{-n}h\left(\frac{t}{\rho(v)},v\right)\tag{5}$$

where $\eta_{p_0}(p)$ is the volume density with respect to p_0, $t\in[0,\rho(v)]$ is a scaling parameter, and $\rho(v)$ is the distance function defined above.

Proof. See Appendix.

Example

In line with our example, we construct a suitable weight function adjusted to the uncertainty surrounding the VaR model. We have seen in the previous section that the underlying process suggests small deviations from the normal distribution and indicate a negative skew. Thus, to determine the weight function we construct a continuous function h that has the maximum value at the point representing our given model p_0, and is monotonically–decreasing with the distance from p_0. This choice means that we are interested more on how sensitive is the model to small variations around p_0. We define h on $[0,1] \times \mathbb{S}$ as follows

$$h(\tau, v) = c\Big(1-t\Big), \quad \forall \nu \in \mathbb{S},\ t \in [0,1]$$

where the normalizing constant c ensures the assumptions $(H5)$ and equals[8] to $\Gamma\left(\frac{n}{2}\right)\pi^{-n/2}$. By applying the continuous mapping $\exp_{p_0} P\Lambda_\rho$ we obtain the weight function K along the geodesics γ:

$$\begin{aligned} K(p,t) &= \frac{1}{\eta_{p_0}(p)} d(p_0,p_1)^{-1}\Gamma\left(\frac{1}{2}\right)\pi^{-1/2}\left(1-\frac{t}{d(p_0,p_1))}\right) \\ &= 1.4686\Big(1-1.47t\Big) \end{aligned}$$

□

7. Measure of Model Risk

In this section we shall introduce a mathematical definition of the quantification of model risk, relate to the concepts introduced so far and study some actual applications.

Recall that we have so far focused on a weighted Riemannian manifold $(\mathcal{M}, I, \zeta)$ with I the Fisher–Rao metric and ζ as in eq. 4. The model in previous sections was assumed to be some distribution $p \in \mathcal{M}$. More likely, a practitioner would define the model as some mapping $f : \mathcal{M} \to \mathbb{R}$ with $p \mapsto f(p)$, i.e. a model outputs some quantity.[9]

We shall formally introduce the normed space $(\mathcal{F}, \|\cdot\|)$ such that $f \in \mathcal{F}$. Though not strictly necessary at this informal stage we shall assume completeness so $(\mathcal{F}, \|\cdot\|)$ is a Banach space.

Definition 5. *With notation as above, let $(\mathcal{F}, \|\cdot\|)$ be a Banach space of measurable functions with respect to ζ. The* model risk *Z of $f \in \mathcal{F}$ and p_0 is given by*

$$Z(f, p_0) = \left\| \frac{f}{f(p_0)} \right\|. \tag{6}$$

In general, f will have units (e.g. monetary for economic capital calculation) so the functional form of eq. 6 allows us to obtain a dimensionless number which is in general a desirable property. It is always possible to add the appropriate units afterwards if so desired. The deviance of the result of eq. 6 from 1.0 can be interpreted as a percentage over the value associated to p_0, e.g. 1.185 would be an expected increase of 18.5% over $f(p_0)$.[10] Another easy to interpret possibility, but with units, is $\|f - f(p_0)\|$.

The quantification of model risk itself can be thought of as a model with a purpose such as provisions calculation or comparison of modelling approaches. Possibilities are endless so we might have started with some $T : \mathcal{F} \to \mathcal{F}$ and set $Z(f, p_0) = \|T \circ f\|$; however, we think eq. 6 is general enough for our present purposes.

---

[8] The volume of the $(n-1)$–dimensional ball $S(0,1)$ is $2\pi^{1/2}\backslash\Gamma\left(\frac{1}{2}\right)$. Thus we have

$$1 = c\int_{[0,1]\times S^{n-1}} \Big(1-t\Big)dt \times d\mu \quad \Rightarrow \quad c = \Gamma\Big(\frac{n}{2}\Big)\pi^{-\pi/2}$$

[9] This is not always the case but we can proceed along these lines depending on the usage to be given to the quantification itself. For example, an inter(extra)polation methodology on a volatility surface is a model whose output is another volatility surface, not a number. If we want to quantify the model risk of that particular approach for Bermudans we might consider its impact on their pricing.

[10] Almost equivalent to eq. 6 would be $\left\| \frac{f - f(p_0)}{f(p_0)} \right\|$.

In what follows we address four examples of Def. 5. Their suitability very much depends among other factors on the purpose of the quantification, as we shall see below.

1. $Z^1(f, p_0)$
$$Z^1(f, p_0) = \left\| \frac{f}{f(p_0)} \right\|_1 = \int_{\mathcal{M}} \left| \frac{f}{f(p_0)} \right| d\zeta$$
This measure represents the total relative change in the outputs across all relevant models.
2. $Z^2(f, p_0)$
$$Z^2(f, p_0) = \left\| \frac{f}{f(p_0)} \right\|_2 = \Big(\int_{\mathcal{M}} \Big(\frac{f}{f(p_0)} \Big)^2 d\zeta \Big)^{1/2}$$
With this measure we put more importance on a big changes in the outputs (big gets bigger and small smaller). So it might be better suited when this emphasis is appropriate, or when it is relevant to keep certain consistency with the norm used in the calibration process (e.g., with maximum likelihood or least square methodologies).
3. $Z^\infty(f, p_0)$
$$Z^\infty(f, p_0) = \left\| \frac{f}{f(p_0)} \right\|_\infty = \operatorname*{ess\,sup}_{\mathcal{M}} \left| \frac{f}{f(p_0)} \right|$$
Calculating Z^∞ one finds the relative worst–case error in the examined model that would be incurred through a deviation from the preferred model, given a precise constraint on the plausibility of the deviation. Further, it can point to the sources of model sensitivity that lead to the largest errors in measuring the risk: using the inverse exponential map, $\exp_{p_0}^{-1}$, we can detect the corresponding direction and size of the change in underlying assumptions that leads to the largest variation in the output and so, ascertain the biggest weakness or limitation of the model.
4. $Z^{s,p}(f, p_0)$
$$Z^{s,p}(f, p_0) = \left\| \frac{f}{f(p_0)} \right\|_{s,p} = \left(\sum_{|k| \le s} \int_{\mathcal{M}} \Big| \partial^k \Big(\frac{f}{f(p_0)} \Big) \Big|^p d\zeta \right)^{1/p}$$
A Sobolev-like norm takes into account the derivatives of f and it can be of interest in those cases when not only f is relevant but its rate of change. An example can be a derivatives model used not only for pricing but also for hedging.

The interpretation of the measures of model risk depends on the way and where the models are subsequently used. For example, in case of the VaR model, the output is further used for decision making in calculating capital requirements, or for setting limits. Capital models dealing with extreme percentiles are vulnerable to statistical error, and so the model risk in capital management may involve the consequences of poor decisions relying on those models.

Example

For illustration in our example we provide the total measure of model risk $Z^1(VaR, p_0)$, and the worst case model risk $Z^\infty(VaR, p_0)$ that equal

$$Z^1(VaR, p_0) = 0.5448 \quad \text{and} \quad Z^\infty(VaR, p_0) = 5.2348$$

$Z^1(VaR, p_0)$ means that there is expected total increase of 54.48% over $VaR(p_0)$. The supremum norm indicates that the biggest expected difference in VaR is 523% over the output from the given model p_0. These results are not surprising since VaR heavily depends on the underlying distribution. However, as sated above more concrete interpretation of these measures depend on the specific use of the model risk quantification. □

Sound methodology for model risk quantification should at least consider the data used for building the model, the model foundation, the IT infrastructure, overall performance, model sensitivity, scenario analysis and most importantly usage. Within our framework we address and measure the uncertainty associated with the aforementioned areas and the information contained in the models. The choice of the embedding and proper neighbourhood of the given model take into account the knowledge and the uncertainty of the underlying assumptions, the data and the model foundation. The weight function that assigns relative relevance to the different models inside the neighbourhood considers the model sensitivity, scenario analysis, the importance of the outcomes with connection to decision making, the business, the intended purpose, and it addresses the uncertainty surrounding the model foundation. Besides, every particular choice of the norm provides a different information and picture of the model with respect to model risk it induces. Last and most important, the model risk measure considers the usage of the model represented by the mapping f [11].

[11] Or equivalently by any possible transformation $T : \mathcal{F} \to \mathcal{F}$.

In order to explore the further financial implications of model risk, it is necessary to know the explicit link between models and decisions. The potential impact of model risk will depend on the very specific features of the business use that the model is put into, for example pricing, financial planning, hedging, or capital management. The plausible size and direction of such financial impact provides us with a measure of the materiality of model risk, in the context of a specific application.

8. Conclusions and Further Research

In this paper we introduce a general framework for the quantification of model risk using differential geometry and information theory. We also present a sound mathematical definition of model risk using Banach spaces over weighted Riemannian manifolds, applicable to most modelling techniques using statistics as a starting point.

Our proposed mathematical definition is to some extent comprehensive in two complementing ways: First, it is capable of coping with relevant aspects of model risk management such as model usage, performance, mathematical foundations, model calibration or data. Second, it has the potential to asses many of the mathematical approaches currently used in any financial institution: Credit risk, market risk, derivatives pricing and hedging, operational risk or XVA (valuation adjustments).

It is worth noticing that the approaches in the literature, to our very best knowledge, are specific in these same two ways: They consider very particular mathematical techniques and are usually very focused on selected aspects of model risk management.

There are many directions for further research, all of which we find to be both of theoretical and of practical interest. We shall finish naming a few of them.

Banach spaces are very well known and have been deeply studied in the realms of for example functional analysis. On the other hand, weighted Riemannian manifolds are non–trivial extensions of Riemannian manifolds, one of the building blocks of differential geometry. The study of Banach spaces over weighted Riemannian manifolds shall broaden our understanding of the properties of these spaces as well as their application to the quantification of model risk.

As introduced in the present contribution, our framework can include data uncertainties by studying perturbations and metrics defined on the sample, which are then transmitted to the weighted Riemannian manifold through the calibration process.

The general methodology can be tailored and made more efficient for specific risks and methodologies. For example, one may interpret the local volatility model for derivatives pricing as an implicit definition of certain family of distributions, extending the Black–Scholes stochastic differential equation (which would define the lognormal family).

Related to the previous paragraph, and despite the fact that there is literature on the topic, the calculation of the Fisher–Rao metric itself deserves further numerical research in order to derive more efficient algorithms.

9. Appendix

Lemma 2. *The weight function K satisfying assumptions $(K1) - (K3)$ is equivalent to assuming that a consistent function $h(t, \nu)$ defined on the C^n with codomain $\mathbb{R}^n$ satisfies the following properties:*

(H1) *$h(t, \nu)$ is a continuous function on the compact manifold C^n*

(H2) *$h(t, \nu) \geq 0,\ (t, \nu) \in [0, 1) \times \mathbb{S}^{n-1}$*

(H3) *$h(1, \nu) = \kappa(\nu)$ for all $\nu \in \mathbb{S}^{n-1}$, with $\kappa(\nu) \geq 0$*

(H4) *$h(0, \nu_1) = h(0, \nu_2) = const.$ for all $\nu_1, \nu_2 \in \mathbb{S}^{n-1}$*

(H5) *$\int_{C^n} h(t, \nu) d\nu = 1$, where $d\nu = dt \times d\mu$*

Proof. To prove the equivalence we need to show that the function h defined in Lemma 2 preserves the required properties of K under the continuous mapping $\exp_{p_0} P\Lambda_\rho$. First we show that the composition that consists of three different mappings is well defined.

As a first step, we define an n–dimensional cylinder

$$C^n := [0,1] \times \mathbb{S}^{n-1} = \{(t,\nu) : t \in [0,1], \nu \in \mathbb{S}^{n-1}\} \subset \mathbb{R}^{n+1}$$

where $\mathbb{S}^{n-1} := \{\nu \in \mathbb{R}^n : ||\nu||^2 = \nu_1^2 + \cdots + \nu_n^2 = 1\}$ denotes the $(n-1)-$dimensional unit sphere in $\mathbb{R}^n$. The cylinder C^n is a differentiable submanifold of $\mathbb{R}^{n+1}$ with boundaries

$$\partial_1 C^n := \{(0,\nu) : \nu \in \mathbb{S}^{n-1}\}, \;\; \partial_2 C^n := \{(1,\nu) : \nu \in \mathbb{S}^{n-1}\}$$

A Riemannian structure on C^n is given by the restriction of the Euclidean metric in $\mathbb{R}^{n+1}$ to C^n. Thus, C^n is a compact Riemannian manifold. A canonical measure on C^n is given by the product measure $dt \times d\mu(\nu)$, where μ denotes the standard surface measure on $\mathbb{S}^{n-1}$.

We define $\rho(v)$ as the length $d(p_0,p) \geq 0$ of the geodesic γ connecting the point p_0 with a boundary point $p \in \partial U$ in the direction $v \in \mathcal{S}(p_0,1)$, where $\mathcal{S}(p_0,1)$ denotes the $(n-1)-$dimensional unit sphere in the tangent space $T_{p_0}\mathcal{M}$.

Note that since U is the subset of the normal neighbourhood with respect to p_0, the exponential map is isometric. From now on we will assume that U is a compact star–shaped subset of a Riemannian manifold $\mathcal{M}$ and the distance function ρ is Lipschitz continuous on $S(p_0,1) \subset T_{p_0}\mathcal{M}$. Lipschitz continuity of $\rho(v)$ is equivalent to the assumption of continuity and piecewise regularity of ∂U.

Now we define an $n-$dimensional subset of C^n by

$$C^n_\rho := \{(t,v) : t \in [0,\rho(v)], v \in \mathbb{S}^{n-1}\} \subset [0,1] \times \mathbb{S}^{n-1}$$

with boundary

$$\partial_1 C^n_\rho := \{(0,v) : v \in \mathbb{S}^{n-1}\}, \;\; \partial_2 C^n_\rho := \{(\rho(v),v) : v \in \mathbb{S}^{n-1}\}$$

The new set C^n_ρ is a compact subset of C^n. In order to map C^n onto C^n_ρ we define the following coordinate transform:

$$\Lambda_\rho : C^n \to C^n_\rho, \quad (t,v) \to \Big(\rho(v)t, v\Big).$$

Since the distance function $v \mapsto \rho(v)$ is strictly positive and Lipschitz continuous on $\mathcal{S}(p_0,1)$, so it is the inverse function $v \mapsto \dfrac{1}{\rho(v)}$. Therefore, the mapping Λ_ρ defines a bi–Lipschitz mapping from C^n onto C^n_ρ. The Jacobian determinant of Λ_ρ equals ρ almost everywhere on C^n.

Next, we are going to introduce polar coordinates on the tangent space $T_{p_0}\mathcal{M}$. We define the transformation P as

$$P: \;\; [0,\infty) \times \mathcal{S}(p_0,1) \to T_{p_0}\mathcal{M} : \quad P(t,v) = tv$$

and its inverse by

$$P^{-1} : \;\; T_{p_0}\mathcal{M} \backslash \{0\} \to (0,\infty) \times \mathcal{S}(p_0,1) : \quad P^{-1}(v) = \Big(||v||, \frac{v}{||v||}\Big)$$

P is well defined by continuity in $T_{p_0}\mathcal{M}$, and maps C^n_ρ onto $\mathcal{U} \subset T_{p_0}\mathcal{M}$. Moreover, the transformation P defines a diffeomorphism from $C^n_\rho \backslash \{\partial_1 C_\rho, \partial_2 C_\rho\}$ onto the open set $\mathcal{U} \backslash \{0, \partial\mathcal{U}\}$. Combining P with the exponential map $\exp_{p_0}$, we have

$$\exp_{p_0} P(C^n_\rho) = U$$

The composition $\exp_{p_0} \circ P$ defines a diffeomorphism from $C^n_\rho \backslash \{\partial_1 C^n_\rho, \partial_2 C^n_\rho\}$ onto $U \backslash \{p_0, \partial U\}$. Furthermore, the boundary $\partial_1 C^n_\rho$ is mapped onto $\{p_0\}$ and the boundary $\partial_2 C^n_\rho$ onto the boundary ∂U [12]. Then the points $(t,v) \in C^n_\rho$ induce geodesic polar coordinates on $\mathbb{R}^n$.

[12] When $p_0 \in \partial U$, the boundary $\partial_2 C^n_\rho$ is mapped onto $\partial U \backslash \{p_0\}$

We have introduced three mappings

$$C^n \xrightarrow{\Lambda_\rho} C^n_\rho \xrightarrow{P} \mathcal{U} \xrightarrow{\exp_{p_0}} U \subset \mathcal{M}$$

The composition $\exp_{p_0} P\Lambda_\rho$ is a continuous mapping from C^n onto U. Moreover, $\exp_{p_0} P\Lambda_\rho$ maps the boundary $\partial_1 C^n$ of the cylinder C^n onto the point p_0 and the boundary $\partial_2 C^n$ onto ∂U.

Now we prove that a consistent function satisfying properties $(H1)-(H5)$ uniquely determines the weight function satisfying $(K1)-(K3)$:

- It is straightforward to see, that properties $(K1)-(K2)$ are satisfied by construction. The composition $\exp_{p_0} P\Lambda_\rho$ preserves connectedness and compactness and it a continuous mapping from C^n onto $\mathcal{M}$. Moreover, $\exp_{p_0} P\Lambda_\rho$ maps the left hand boundary $\partial_1 C^n$ onto the point p_0 and the right hand boundary $\partial_2 C^n$ onto the boundary of U. Hence, the image $f(\exp_{p_0}(\rho(v)tv))$ of a continuous function f on $\mathcal{M}$ is also continuous on the cylinder C^n, and every function g defined on C^n satisfying consistency properties $(i)-(ii)$ of Definition 3 is the image of a continuous function on $\mathcal{M}$ under the pull–back operator $\Lambda_\rho^{-1} P^{-1} \exp_{p_0}^{-1}$.

  The composition $\exp_{p_0} P\Lambda_\rho$ applied to a function h that is continuous on C^n and satisfies the consistency conditions $(i)-(ii)$ of Definition 3 will give us a continuous function K on $\mathcal{M}$ that by construction is continuous along the geodesics starting at p_0 and ending at the points of the boundary. That means, property $(K1)$ is satisfied. The same argument applies to any non–negative function h on C^n. Thus, properties $(H1)-(H2)$ ensures $(K1)-(K2)$ under the composition $\exp_{p0} P\Lambda_{\rho(v)}$.

- Further, it remains to prove that the weight function K is indeed a probability density on $\mathcal{M}$ with respect to the measure $dv(p)$, i.e. to show that $\int_M d\zeta = 1$.

  $$\int_{\mathcal{M}} d\zeta = \int_{\mathcal{M}} K(p,t)dv(p) = \int_{\mathcal{M}} K(\exp_{p_0}(v),t)\eta(v)d\xi$$

  where $d\xi$ is the standard Lebesgue measure on the Euclidean space $T_{p_0}\mathcal{M}$ and $\eta_{p_0}(v) = det((d\exp_{p_0})_v)$ is the Jacobian determinant of the exponential map. Note that $\eta(v)$ represents the density function that is a positive and continuously differentiable function on $U \subset T_{p_0}\mathcal{M}$ and the zeros of η lie at the boundary of $\mathcal{M}$. Further, we have

  $$\int_{\mathcal{M}} K(\exp_{p_0}(v),t)\eta(v)d\xi = \int_{\mathcal{S}(p_0,1)} \int_0^{\rho(v)} t^{n-1} K(\exp_{p_0}(tv),t)\eta_{p_0}(tv)dtd\mu(v)$$

  where t^{n-1} is the Jacobian determinant of the polar coordinate transformation and $d\mu(v)$ is the standard Riemannian measure on the unit sphere $\mathcal{S}(p_0,1)$. The last step is the mapping from C_ρ to C^n:

  $$\int_{\mathcal{S}(p_0,1)} \int_0^{\rho(v)} t^{n-1} K(\exp_{p_0}(tv))\eta_{p_0}(tv)dtd\mu(v) = \int_{S^{n-1}} \int_0^1 \frac{1}{\rho(v)} K\Big(\exp_{p_0}(\rho(v)tv)\Big)\big(\rho(v)t\big)^{n-1}\eta_{p_0}(\rho(v)tv)dtd\mu(v)$$

  where the Jacobian determinant is $\dfrac{1}{\rho(v)}$. Then using the expression for K we have that the expression above is equal to:

  $$\begin{aligned} &= \int_{S^{n-1}} \int_0^1 \frac{1}{\rho(v)} \frac{1}{\eta_{p_0}(\rho(v)tv)} t^{1-n}\rho(v)^{-n} h\left(\frac{t}{\rho(v)}, v\right)\big(\rho(v)t\big)^{n-1}\eta_{p_0}(\rho(v)tv)dtd\mu(v) \\ &= \int_{S^{n-1}} \int_0^1 h\left(\frac{t}{\rho(v)}, v\right)dtd\mu(v) = 1 \end{aligned}$$

□

Theorem 6. *A continuous function h defined on the cylinder C^n satisfying conditions $(H1)-(H4)$ determines a unique and well defined weight function K on $\mathcal{M}$ given by*

$$K(p,t) = \frac{1}{\eta_{p_0}(p)} t^{1-n}\rho(v)^{-n} h\left(\frac{t}{\rho(v)}, v\right) \tag{7}$$

where $\eta_{p_0}(p)$ is the volume density with respect to p_0, $t \in [0,\rho(v)]$ is a scaling parameter, and $\rho(v)$ is the distance function defined in the Proof of Lemma 2.

Proof. Note that the composition $\exp_{p_0} P\Lambda_\rho$ induces a change of variables for integrable function f that yields to the following formula:

$$\begin{aligned}\int_{\mathcal{M}} f(p)d\zeta &= \int_U f(p)d\zeta = \int_{\mathcal{U}} f(\exp_{p_0}(v))\eta_{p_0}(v)dv \\ &= \int_{\mathcal{S}_{p_0}}\int_0^{\rho(v)} f(\exp_{p_0}(tv))t^{1-n}\eta_{p_0}(t,v)dtdv \\ &= \int_{\mathbb{S}_{p_0}}\int_0^1 f(\exp_{p_0}(\rho(\nu)\nu))\frac{1}{\rho(\nu)}\eta_{p_0}(t\rho(\nu),\nu)dtd\nu\end{aligned}$$

The volume density η_{p_0} is a well-defined, non-negative function with zeros at the cut locus of the point p_0. Besides, η_{p_0} is continuous and differentiable function on $\mathcal{M}$[13]. The distance function ρ is a well defined, strictly positive and Lipschitz continuous function on $\mathcal{S}(p_0,1)$, and thus is the inverse $1/\rho(v)$. Therefore, the mapping Λ_ρ defines a bi-Lipschitzian mapping from C^n to C^n_ρ. Moreover, the composition $\exp_{p_0} P$ defines a diffeomorphism from C^n_ρ $\{\partial_1 C^n_\rho, \partial_2 C^n_\rho\}$. The using the fact that the point set $\{p_0\}$ and the boundary of U are subsets of $t\nu-$measure zero, we can conclude that the mapping $\exp_{p_0} P\Lambda_\rho$ is isomorphism. Then for any h defined on C^n satisfying conditions $(i)-(ii)$ of Definition 3, the associated weight function K is well defined on U. The uniqueness of K follows after specifying a function h that satisfies properties $(H1)-(H5)$. □

[13]Note that when $\mathcal{M}$ is $\mathbb{R}^n$ with the canonical metric, then $\eta_{p_0}(p) = 1$ for all $p \in \mathbb{R}^n$.

References

[1] Amari, S-I, OE Barndorff-Nielsen, RE Kass, SL Lauritzen, and CR Rao (1987), "Differential geometry in statistical inference." *Lecture Notes-Monograph Series*, i–240.

[2] Azzalini, Adelchi (1985), "A class of distributions which includes the normal ones." *Scandinavian journal of statistics*, 171–178.

[3] Bhattacharyya, Anil (1946), "On a measure of divergence between two multinomial populations." *Sankhyā: the indian journal of statistics*, 401–406.

[4] Boucher, Christophe M, Jón Daníelsson, Patrick S Kouontchou, and Bertrand B Maillet (2014), "Risk models-at-risk." *Journal of Banking & Finance*, 44, 72–92.

[5] Branger, Nicole and Christian Schlag (2004), "Model risk: A conceptual framework for risk measurement and hedging." In *EFMA 2004 Basel Meetings Paper*.

[6] Chavel, Isaac (2006), *Riemannian geometry: a modern introduction*, volume 98. Cambridge university press.

[7] Christodoulakis, George and Stephen Satchell (2008), "The validity of credit risk model validation methods." *The Analytics of Risk Model Validation*, 27–44.

[8] Cont, Rama (2006), "Model uncertainty and its impact on the pricing of derivative instruments." *Mathematical finance*, 16, 519–547.

[9] Do Carmo, Manfredo Perdigao (1992), *Riemannian geometry*.

[10] European Banking Authority (2012), "Discussion paper on draft regulatory technical standards on prudent valuation, under article 100 of the draft capital requirements regulation (crr)." *Discussion Paper*.

[11] Federal Reserve (SR 11–7) (2011), "Supervisory guidance on model risk management." *Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, SR Letter*, 11–7.

[12] Federer, Herbert (2014), *Geometric measure theory*. Springer.

[13] Gibson, Rajna (2000), *Model risk: concepts, calibration and pricing*. risk books.

[14] Hopf, H. and W. Rinow (1931), "Uber den begriff der vollstandigen differentialgeometrischen flachen." Helv. 3, 209–225.

[15] Hull, John and Wulin Suo (2002), "A methodology for assessing model risk and its application to the implied volatility function model." *Journal of Financial and Quantitative Analysis*, 37, 297–318.

[16] Joshi, S, A Srivastava, and IH Jermyn (2007), "Riemannian analysis of probability density functions with applications in vision." IEEE.

[17] Lang, Serge (2012), *Fundamentals of differential geometry*, volume 191. Springer Science & Business Media.

[18] Morgan, Frank (2005), "Manifolds with density." *Notices of the AMS*, 853–858.

[19] Morini, Massimo (2011), *Understanding and Managing Model Risk: A practical guide for quants, traders and validators*. John Wiley & Sons.

[20] Murray, Michael K and John W Rice (1993), *Differential geometry and statistics*, volume 48. CRC Press.

[21] Rao, C Radhakrishna (1945), "Information and accuracy attainable in the estimation of statistical parameters." *Bull Calcutta. Math. Soc.*, 37, 81–91.

[22] Saltelli, A, S Tarantola, F Campolongo, and M Ratto (2013), "Sensitivity analysis in practice: A guide to assessing scientific models. 2004." *Chichester, England: John Wiley & Sons*, 46556, 48090–9055.