A New Approach to Quantification of Model Risk for Practitioners

Zuzana Krajčovičová
Dept. of Mathematics, University of A Coruña, Campus de Elviña, 15071 A Coruña, Spain,
e-mail: z.krajcovicova@udc.es

Pedro Pablo Pérez-Velasco
Model Risk Division, Banco Santander, Av. Cantabria s/n, 28660-Boadilla del Monte, Spain,
e-mail: pedpperez@gruposantander.com

Carlos Vázquez
Dept. of Mathematics, University of A Coruña, Campus de Elviña, 15071 A Coruña, Spain,
e-mail: carlosv@udc.es

Abstract

Worldwide regulation obliges financial institutions to manage model risk with the same severity as any other risk. Its quantification is essential not only in meeting these requirements but also for institution's basic internal operative. In this article we address the quantification of model risk by the calculation of the norm of an appropriate function defined on a Riemannian manifold endowed with Fisher–Rao metric. The aim is twofold: introduce a sufficiently general and sound mathematical framework to cover the main points in model risk and illustrate how a practitioner may identify the relevant abstract concepts and put them to work.

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 deliberately focus on specific aspects of the reality and degrade or ignore the rest. Understanding the capabilities and limitations of the underlying assumptions is key when dealing with a model and its outputs. According to US [10], 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"

Thus, US Federal Reserve identifies the two main reasons for model risk (inappropriate usage and fundamental errors). Furthermore, 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. The US Federal Reserve also emphasizes 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.

In spite of the rise of awareness of model risk and understanding its significant impact, there are no globally defined industry nor market standards on its exact definition, management and quantification, even though as aforementioned a proper model risk management is required by regulators.

Within the finance literature, some authors have defined model risk as the uncertainty about the risk factor distribution as in [12], the misspecified underlying model in [9], the deviation of a model from a 'true' dynamic process in [5], the discrepancy relative to a benchmark model in [13], and the inaccuracy in risk forecasting that arises from the estimation error and the use of an incorrect model in [4]. Model risk has been previously classified in all asset classes, see [18] for interest rate products and credit products, [8] for portfolio applications, [23] for asset backed securities, and [4] for relation to measuring market 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 quantification of model risk should consider the uncertainty stemming from the selection of the mathematical techniques (e.g. focusing on fitting a normal distribution hence leaving aside other distribution families), 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).

Model risk quantification poses many challenges that come from the high diversity of models, the wide range of techniques, the different use of models, among others. Some model outputs drive decisions; other model outputs provide one source of management information, some outputs are further used as an inputs in other models. Additionally, the model outputs may be completely overridden by expert judgment, not to mention that in order to quantify model risk you need another model, which is again prone to model risk.

It is common practice to consider the most relevant areas of analysis for the quantification of model risk to be data and calibration, model foundations, model performance, IT infrastructure, model use, controls and governance, and model sensitivity. The model may be fundamentally wrong due to errors in its theoretical foundation or conceptual design that emerge from incorrect 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. Flawed test procedures or failure to perform consistent and comprehensive user acceptance tests can lead to material model risk. To name just a few.

The focus of this paper is on developing a new approach for quantifying model risk within the framework of differential geometry, see [20], and information theory, see [1]. 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. This metric 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 the central concepts from 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 and Section 5 describe the construction of the neighbourhood containing material variations of the model, and the definition of the weight function, respectively. The model risk measure is then defined and explained in Section 6. Section 7 is dedicated to an empirical example of the model risk calculation of a credit risk model. Section 8 provides some final conclusions and directions for future work, and finally, the Appendix contains the construction of the weight function and the proofs of the main results.

2 Background on Riemannian Geometry

In this section we introduce some required concepts from differential geometry and information theory, and fix the notation and terminology used throughout the article. For more details we refer the reader to e.g. [1] or [20].

Let $\mathcal{M}$ be a statistical manifold consisting of probability density functions $p(x|\theta)$ of random variable $x \in \mathcal{X}$ with respect to a measure μ on $\mathcal{X}$, such that every distribution is uniquely parameterized by an n–dimensional vector $\theta = (\theta^i) = (\theta^1, \dots, \theta^n)$. [1] Specifically, let

$$\mathcal{M} = \{p(x|\theta) \mid \theta \in \Theta \subset \mathbb{R}^n\},$$

with a one–to–one mapping between θ and $p(x|\theta)$. Additionally, under the assumptions that the parametrization of $\mathcal{M}$ is differentiable and C^∞ diffeomorphism, the parametrization θ forms a coordinate system of $\mathcal{M}$, [1]. The local coordinate system $\theta = (\theta^1, \cdots, \theta^n)$ then induces a basis $\dfrac{\partial}{\partial\theta} = (\partial_1, \cdots, \partial_n)$ of the tangent spaces (∂_i is a shorter notation for $\partial/\partial\theta^i$).

The structure of $\mathcal{M}$ is specified by a Riemannian metric, $g = (g_{ij})$, that is defined by a local product on tangent vectors at each point $p \in \mathcal{M}$ denoted by $g : T_p\mathcal{M} \times T_p\mathcal{M} \to \mathbb{R}$, that is symmetric, bi-linear, positive definite and C^∞ differentiable in p. By the bi–linearity of the inner product of g, for any two tangent vectors $u, v \in T_p\mathcal{M}$

$$g(u, v) = \sum_{i=1}^{n}\sum_{j=1}^{n} v_i u_j g(\partial_i, \partial_j)$$

[1] We describe only the case for continuum on the set $\mathcal{X}$, however if $\mathcal{X}$ were discrete, the given framework will still apply by switching $\int(\cdot)$ with $\sum(\cdot)$.

where $\{\partial_i\}_{i=1}^n$ form the basis elements of $T_p\mathcal{M}$. For a statistical manifold, the Fisher–Rao information metric is given by

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

where $I(p) = I\big(p(x|\theta)\big)$ can be considered as a Riemannian metric. Note that $det\,(I(p(x|\theta)))$ represents the amount of information a sample point conveys with respect to the problem of estimating the parameter θ, and so $I(\cdot)$ can be used to determine the dissimilarities between distributions. It measures the ability of the random variable x to discriminate the values of the parameter θ' from θ for θ' close to θ.

Example (statistical manifold of normal distributions): The normal distribution with mean μ and variance σ^2 is given by

$$\mathcal{N}(x|\mu,\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}}\exp\Big\{-\frac{1}{2\sigma^2}(x-\mu)^2\Big\},\ x \in \mathbb{R}, \mu \in \mathbb{R}, \sigma > 0.$$

Set $\Theta = \{(\theta^1,\theta^2) \in \mathbb{R}^2 | \theta^2 > 0\}, \mathcal{X} = \mathbb{R}$ and $p(x|\theta) = \mathcal{N}(x|\sqrt{2}\theta^1, (\theta^2)^2)$. Then the statistical manifold with respect to $\{p(\cdot|\theta)|\theta \in \Theta\}$ has a Riemannian metric $g = 2(\theta^2)^{-2}\sum d\theta^i d\theta^i$ of constant curvature $-1/2$.

The Riemannian metric encodes how to measure distances, angles, areas and curvature on the manifold by specifying the scalar product between tangent vectors at a particular point. If we consider a curve $\gamma : [a,b] \subset \mathbb{R} \to \mathcal{M}$ on the manifold, its length $\ell(\gamma)$ can be defined as

$$\ell(\gamma) = \int_a^b \left\|\frac{d\gamma}{dt}\right\| dt = \int_a^b \sqrt{g_{ij}\dot{\gamma}^i\dot{\gamma}^j}dt = \sqrt{\langle\dot{\gamma},\dot{\gamma}\rangle}.$$

where $\dot{\gamma}^i$ is the derivative of $\dot{\gamma}^i = \dfrac{\partial\gamma}{\partial\theta^i}$. Then, the distance between two points $p, q \in \mathcal{M}$ is defined by the infimum of the length of all smooth curves between these two points and is give by

$$d(p,q) = \int_{\gamma\in\Gamma} \ell(\gamma)$$

where Γ is a set of all smooth curves between these two points. The locally length–minimizing smooth curve $\gamma(t) : [0,1] \to \mathcal{M}$ is called geodesic and is characterized by the fact that it is autoparallel, e.g. the field of tangent vectors $\dot{\gamma}(t)$ stays parallel along γ (velocity is constant along the geodesic $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$ on γ). In local coordinates, a curves is a geodesic *iff* it is the solution of the system of n second order Euler–Lagrange equations:

$$\frac{d^2\theta^k}{dt^2} + \sum_{i,j=1}^n \Gamma_{ij}^k \frac{\partial\theta^i}{dt}\frac{\partial\theta^j}{dt} = 0 \quad \forall k = 1,\ldots,n,$$

where Γ_{ij}^k are the Christoffel symbols of the second kind.

From the result of existence and uniqueness of solutions of differential equations, for each $p \in \mathcal{M}$ and for each tangent vector $v \in T_p\mathcal{M}$, there exists an open interval I_v with $0 \in I_v$ and a unique geodesic $\gamma_v : I_v \to \mathcal{M}$, such that $\gamma_v(0) = p$ and $\dot{\gamma}_v(0) = v$. Therefore, the exponential mapping $\exp_p : T_p\mathcal{M} \to \mathcal{M}$ is defined by $\exp_p(v) = \gamma_v(1) = \gamma_{v_1}(\|v\|)$, with $v_1 = v/\|v\|$. For each $p \in \mathcal{M}$, there exists a neighborhood $\mathcal{V}$ of the origin in $T_p\mathcal{M}$, such that $\exp_p$ is a diffeomorphism from $\mathcal{V}$ onto a neighborhood V of p. The neighborhood $\mathcal{V}$ is star–shaped, i.e. for any point belonging to $\mathcal{V}$, the line joining the point to the origin is contained in $\mathcal{V}$. The image of a star–shaped neighborhood of the origin under the exponential map is a neighborhood of p on the manifold (also called normal neighborhood).

A notion of connection ∇ defines a map between any neighboring tangent spaces. The canonical affine connection on a Riemannian manifold is the Levi–Civita connection and it is directly defined from the covariant derivative, i.e. orthogonal projection of the usual derivative on the vector fields onto tangent space. It parallel transports a tangent vector along a curve while preserving its inner product (it is compatible with the metric, i.e. the covariant derivative of the metric is zero). The Levi–Civita coefficients are defined, in each local chart by the Christoffel symbols of the second kind Γ_{ij}^k given by

$$\nabla_{ij}^k = \Gamma_{ij}^k = \frac{1}{2} g^{kl} \left(\frac{\partial g_{jl}}{\partial \theta^i} + \frac{\partial g_{il}}{\partial \theta^j} - \frac{\partial g_{ij}}{\partial \theta^l} \right)$$

$\forall i, j, k, l = 1, \ldots, n$ is used the Einstein summation convention, g^{kl} is the metric inverse. When the connection coefficients of ∇ with respect to a coordinate system of $\mathcal{M}$ are all identically 0, then ∇ is said to be flat, or alternatively, $\mathcal{M}$ is flat with respect to ∇.

3 Modeling Process Steps and Quantification of Model Risk

There are different types and aspects of model risk that tend to easily overlap, co–occur or co–vary. In this context, we propose four rough model creation steps: Data, Calibration, Model Selection and Testing, and Implementation and Usage. These steps may occur in an iterative fashion, although 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. Model risk may arise from data deficiencies in terms of both quality and availability, including, among others, error in data definition, insufficient sample, inaccurate proxies, sensitivity to expert judgments, or misinterpretation.

2. **Calibration** includes the selection of the types of variables and the nature of their treatment, the tuning of free parameters and links between system components and processes. Estimation uncertainty may occur due to simplifications, approximations, flawed assumptions, inappropriate

calibration, errors in statistical estimation or in market benchmarks, computational or algorithmic limitations, or use of unobservable parameters.

3. **Model Selection and Testing** involves the choice of the 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. Furthermore, it is related to the conditional verification which includes checking the sensitivity to changes in the data and to possible deviations 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 or model misspecification.

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 recalibration, IT failures, lack of communication between modelers and users, lack of understanding on model limitations.

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).

In this article we aim to show that differential geometry and information theory offer a base for such an approach. For this purpose, in our setting a model is represented by a particular probability distribution, $p : \mathcal{X} \to \mathbb{R}_+$, that belongs to a manifold 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 allows us to quantify variations and dissimilarities between probability distribution functions (models), among other things.

The set of probability measures may be further parameterized in a canonical way by a parameter space Θ, $\mathcal{M} = \{p(x;\theta) \mid \theta \in \Theta\}$. This set forms a smooth Riemannian manifold $\mathcal{M}$. 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 natural Riemannian metric is shown to be the Fisher–Rao metric (see [21]), which is the unique intrinsic metric on the statistical manifold. It is the only metric that is invariant under re–parametrization, see [1] for example.

Let us consider a given model p_0 which can be uniquely parametrized using the vector $\theta_0 = (\theta_0^1, \ldots, \theta_0^n)$ over the sample space $\mathcal{X}$ and which can be described by the probability distribution $p_0 = p(x;\theta_0)$. This probability distribution belongs to a set (family) of distributions $\mathcal{M} = \{p(x;\theta) \mid \theta \in \Theta \subset \mathbb{R}^n\}$ 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. Thus, choosing a particular model is equivalent to fixing a parameter vector $\theta \in \Theta$.

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. 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.[2] The analysis consists of five steps:

1. Embedding the model manifold into one that considers missing properties[3] 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.[4] 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 aspects of information inherited in p_0. In order to incorporte 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 add new dimensions to $\mathcal{M}$, and thus consider a higher–dimensional space into which $\mathcal{M}$ is embedded. [5]

We define a neighbourhood [6] around p_0 with the help of the tangent space $T_{p_0}\mathcal{M}$. Note that $T_{p_0}\mathcal{M}$ is a vector space that describes at first order approximation, infinitesimal displacements or deformations on the manifold at the point p_0. From a practical point of view, not all perturbations are relevant, thus taking into account the materiality with respect to the intended purpose of the model, its usage,

[2] 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 ([1]). Additionally, for this choice of connection, the shortest paths are geodesics.

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

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

[5] Consider for example a case when the underlying model space represents the family of normal distributions, i.e. 2–dimensional manifold $\mathcal{M}_0$, we may want to consider the family of skew normal distributions, i.e. 3–dimensional manifold $\mathcal{M}$, for which $\mathcal{M}_0 \subset \mathcal{M}$, in order to also examine skewness.

[6] Note that throughout the paper we do not refer to the neighbourhood as to a strictly topological neighborhood.

business and market, we consider only a small subset of the tangent bundle.

Let $\mathcal{U}$ be the open set around p_0 of normal neighbourhood V such that

$$\mathcal{U} = \{tv \in V \subset T_{p_0}\mathcal{M} \mid 0 < t \leq \alpha(v),\ v \in \mathcal{S}(p_0, 1) \text{ and normal coordinates are defined}\},$$

where $\mathcal{S}(p_0, 1) = \{v \in T_{p_0}\mathcal{M}, \|v\| = 1\}$ is the unit sphere on $T_{p_0}\mathcal{M}$.

The set $\mathcal{U}$ includes the directions of all relevant perturbations of the model p_0 up to a certain level $\alpha(v)$in direction 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, or higher sensitivity to one of the parameters than to others. The level $\alpha(v)$ can be interpreted as a means to control uncertainty regarding the choice of the model p_0, the uncertainty surrounding the data and calibration, and it is appropriately chosen based on the usage of the model.

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 to p_0:

$$U = \exp_{p_0}(\mathcal{U}) = \{p \in \mathcal{M} \mid d(p_0, p) \leq \alpha(v)\}, \tag{2}$$

From now on, we shall require the boundary $\partial U = \{\exp_{p_0}(\alpha(v)\, v) \mid v \in \mathcal{S}(p_0, 1)\}$ to be continuous and piecewise regular. Moreover, U shall be a compact *star–shaped set with respect to* p_0 that is defined as follows:

Definition 4.1 *A compact subset U of a Riemannian manifold $\mathcal{M}$ is called star–shaped with respect to $p_0 \in U$ if for all $p \in U, p \neq p_0$ there exists a geodesic segment γ with $\gamma(0) = p_0$ connecting p_0 and p such that $\gamma(t) \in U$ for all $t \in [0, \alpha(v)]$, where $\alpha(v) > 0$, with $v \in \mathcal{U}$.*

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 $\gamma(t) = \exp_{p_0}(tv)$ for $t \in [0, \alpha(v)]$.

5 Weight Function Definition

Variations of the chosen model are not equally material and they all might take place with different probabilities. By introducing 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.

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 associated to the Fisher–Rao information metric, which given by $dv(p) = \sqrt{det(I(\theta))}d\theta$. The volume measure is the unique Borel measure on $\mathcal{M}$ ([11]). With respect to a coordinates system, the information volume 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 density with respect to the Riemannian volume, i.e. $\int_{\mathcal{M}} K dv(p) = 1$. 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 5.1 *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 $T_p\mathcal{M}$. In particular, the Fisher–Rao information metric I on $\mathcal{M}$ maps each $p \in \mathcal{M}$ to a volume $dv(p)$ which is a symmetric and n–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 density K over $\mathcal{M}$ with respect to the Riemannian measure induces a new absolutely continuous probability measure ζ with respect to Vol

$$\zeta(U) = \int_U d\zeta = \int_U K dv(p) \tag{3}$$

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 ([17]).

The weight function K of the Definition 5.1 represents a general characterization of a probability density over the Riemannian manifold $\mathcal{M}$. In order 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.

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

From a practitioner point of view, models that do not belong to the chosen neighbourhood U are not relevant from the perspective of model risk, and so they do not add any uncertainty. Therefore, we assume the weight function K 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 $v \in \mathcal{S}(p_0, 1) \subset T_{p_0}\mathcal{M}$ starting at p_0 and ending at the points on ∂U. These additional properties motivate the following respective modifications of $(K1')$ and $(K2')$ in previous Definition 5.1:

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

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

A 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.

The construction of a weight function on a given Riemannian manifold may become technically difficult since it requires precise knowledge of the intrinsic geometry and the structure of the manifold. In order to overcome this difficulty and to determine a weight function K that satisfies all the required properties, 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. In summary, we construct three mappings: the exponential map $\exp_{p_0}$, the polar transform P and a further coordinate transform Λ_ρ. Euclidean geometry is well understood and intuitive, and thus a construction of a function on this space is considerably easier and more intuitive. The steps of the construction and the associated proofs are presented in the Appendix (see Sec. 9), as we believe this will improve the readability and flow of the article.

6 Measure of Model Risk

In this section we introduce a definition of the quantification of model risk, relate it to the previously introduced concepts and study some realistic 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. (3). The model was assumed to be identified with a 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.[8]

[8] 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 Bermudan derivatives we might consider its impact on their pricing.

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

Definition 6.1 *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*

$$Z(f, p_0) = \|f - f(p_0)\| . \tag{4}$$

Note that the measure represents the standard distance. All outcomes are constrained by the assumptions used in the model itself and so, the model risk is related to the changes in the output while relaxing and changing them.

The quantification of model risk itself can be thought of as a model with a purpose such as provisions, capital calculation or comparison of modelling approaches. Possibilities are endless, so we might have started with some $T : \mathcal{F} \to \mathcal{F}$ and set $Z(f, p_0) = \|T \circ f\|$;[9] however, we think Eq. (4) is general enough for our present purposes.

In what follows we consider four examples of Def. 6.1. Their suitability very much depends among other factors on the purpose of the quantification, as we shall see later.

1. $Z^1(f, p_0)$ for $f \in L^1(\mathcal{M})$ represents the total relative change in the outputs across all relevant models:
$$Z^1(f, p_0) = \|f - f(p_0)\|_1 = \int_{\mathcal{M}} \Big| f - f(p_0) \Big| d\zeta.$$

2. $Z^2(f, p_0)$ for $f \in L^2(\mathcal{M})$ puts more importance on big changes in the outputs (big gets bigger and small smaller). It would allow to keep consistency with some calibration processes such as the maximum likelihood or least square algorithms:
$$Z^2(f, p_0) = \|f - f(p_0)\|_2 = \left(\int_{\mathcal{M}} \Big(f - f(p_0) \Big)^2 d\zeta \right)^{1/2}.$$

3. $Z^\infty(f, p_0)$ for $f \in L^\infty(\mathcal{M})$ finds the relative worst–case error with respect to p_0:
$$Z^\infty(f, p_0) = \|f - f(p_0)\|_\infty = \operatorname*{ess\,sup}_{\mathcal{M}} \Big| f - f(p_0) \Big|.$$
Furthermore, it can point to the sources of the largest deviances: Using $\exp_{p_0}^{-1}$ we can detect the corresponding direction and size of the change in the underlying assumptions.

[9] For example, another possibility is to use $\left\| \frac{f}{f(p_0)} \right\|$ or $\left\| \frac{f - f(p_0)}{f(p_0)} \right\|$. These functional forms would allow us to obtain a dimensionless number which might be a desirable property.

4. $Z^{s,p}(f, p_0)$ for $f \in W^{s,p}(\mathcal{M})$ is a Sobolev norm that can be of interest in those cases when not only f is relevant but its rate of change:[10]

$$Z^{s,p}(f, p_0) = \|f - f(p_0)\|_{s,p} = \left(\sum_{|k| \leq s} \int_{\mathcal{M}} \left| \partial^k \Big(f - f(p_0) \Big) \right|^p d\zeta \right)^{1/p}.$$

A 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 takes 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 also addresses the uncertainty surrounding the model foundation. Besides, every particular choice of the norm provides different information of the model. Last and most important, the model risk measure considers the usage of the model represented by the mapping f.[11]

7 Application to Capital Calculation

In this section we contextualize the proposed framework by applying it to a credit risk model used by a commercial bank. More specifically, we employ the proposed methodology for the quantification of model risk to the probability of default (PD) model of high default portfolios used for capital calculation.[12] We analyze different scenarios and risk parameter assumptions in order to assess how these scenarios affect the economic capital based on methodologies commonly applied by IRB institutions.

In general, the purpose of the credit risk model is to estimate the PD of future credit losses on a bank portfolio. For a given time horizon, the model generates a distribution—a probability density function— of future losses that can be used to calculate the losses associated with any given percentile of the distribution. In practice, banks concentrate on two such loss components: expected loss (EL) and unexpected loss (UL). EL is the mean of the loss distribution and represents the amount that a bank expects to lose on average on its credit portfolio over a given time horizon. In contrast, UL refers to the risk of the portfolio that is computed as the losses associated with some high percentile of the loss distribution (e.g., the $99.99th$ percentile), thus covering all but most extreme events. For more details about credit risk, see for example [3].

[10] An example can be a derivatives model used not only for pricing but also for hedging.

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

[12] Note that the purpose of the analysis is only to illustrate the proposed framework, and so it does not fully comprehend all potential risks inherent in PD calculation.

Capital models are usually based on three risk parameters: the PD, the loss given default (LGD) and the exposure at default (EAD). Under the Basel II IRB framework, the PD per rating grade is the average percentage of obligors that will default over a one–year period, EAD gives an estimate of the amount outstanding if the borrower defaults, and LGD represents the proportion of the exposure that will not be recovered after default. These parameters are aggregated from obligor level (risk bucket) to portfolio level with the correlation set by the regulator or the financial entity.

Let N be the number of borrowers in a given loan portfolio. Assuming a uniform value of LGD, the aggregated expected loss amount, L, can be calculated as the sum of individual Ls in the portfolio, i.e.,

$$L = \sum_{i=1}^{N} EAD_i \cdot LGD_i \cdot PD_i. \tag{5}$$

Based on the loss distribution, the capital requirement for a bank under the IRB approach at confidence level α can be calculated simply as the difference between EL and the percentile for the level being considered:

$$C_\alpha(L) = q_L(\alpha) - EL \tag{6}$$

where $q_L(\alpha)$ is the α-quantile of L defined by $\mathbb{P}(L \leq q_L(\alpha)) = \alpha$. In practice, the portfolio is categorized into homogeneous risk buckets, $j = 1, \ldots, M$, and the capital calculation is done on the risk–bucket level. The same default probability PD_j and LGD_j is then assigned for all borrowers in each of these buckets.

When modeling credit risk losses, several sources of model risk may arise. Examples are the scarcity of default events, lack of data driving calibration and backtesting, correlations between failures, wrong way exposure (growing utilization of credit lines in case of an increase in PD), or independence of PD and LGD.

For the purpose of illustration, we only focus on the quantification of the model risk arising from the PD estimation on the level of capital, and so we assume both LGD and EAD be known and independent of PD.

7.1 PD model (p_0)

PD is the likelihood that an obligor will default within one–year given all currently available information. We consider a particular internal model for the long–run PD estimation that serves the regulatory purposes. The PD model is built on internal behavioral data and bureau information. Each customer account is scored and the portfolio is categorized, with respect to the scoring, into M homogeneous risk buckets. The number of defaults in each bucket is assumed to follow a binomial random variable, in which the defaults are independent across customers and over time, and defaults occur with common probability. The underlying Point–in–time (PIT) PD model is calibrated to the latest observed

PD, i.e. observed default frequency (ODF), that is just the approximation of the maximum likelihood estimator of the parameter of the binomial distribution. PIT PD is then adjusted by the Central Tendency (CT) to generate the long–run PD based on a combination of historical misalignment of the underlying model and expert judgment. The Through–the–Cycle (TTC) PD, i.e. pooled PD for each risk bucket, is given by the formula:

$$PD_i = ODF_i \cdot \frac{CT}{ODF}, \tag{7}$$

where:

- ODF_i is the observed default frequency obtained for each risk bucket i over the most recent quarter of the calibration sample.
- CT refers to the Central Tendency used for the TTC adjustment that is based on external macroeconomic data series in order to extend the internal ODF series.[13]
- ODF is the average ODF over the given segment portfolio.

Thus, the long–run observed default frequencies are calculated for each risk bucket, and are adjusted to the average PD observed for each portfolio over a complete economic cycle. The PD therefore gives the likelihood for obligors with a particular rating grade at the start of a given time period defaulting within this period. The distribution of defaults under our simplifying assumptions would equal the loss distribution of a portfolio for which all the borrowers had an EAD of 1, an LGD of 0.45 and a maturity of 1 year.

The empirical analysis is based on a hypothetical portfolio that was elaborated in order to represent the behavior of one particular segment consisting of $N = 9860$ clients. All customers are categorized by risk scoring into buckets that are heterogeneous at a defined confidence level and represent the grade of credit quality. We assume a portfolio consisting of ten buckets ($M = 10$) with a monotonous PD related to the model scores: higher scores imply lower ODFs.

All customers within the same risk bucket are assigned the same "pooled" PD, which can be thought of as the average of individual PDs. This means that the pooled PD assigned to a risk bucket is a measure of the average value of the PDs of customers in that bucket.

Risk buckets	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
PD	0.276	0.198	0.181	0.090	0.085	0.065	0.037	0.034	0.023	0.012
N(frequency)	0.007	0.012	0.021	0.055	0.066	0.088	0.176	0.280	0.241	0.056

Table 1: Normalized PD and frequency of accounts in the portfolio across risk buckets.

[13] In case there is not enough data history in order to cover the whole economic cycle.

The normalized PDs vary between 0.01143 and 0.27638 and decrease from lower to higher risk ratings, although this decrease is moderate (Table 1). The distribution is highly skewed to the right (as expected) with mean and variance equal to 0.4177 and 0.08768, respectively. The frequency of accounts across buckets ranges from 0.70% to 28.89% (Table 1). The low PD buckets account for the majority of the portfolio, where approximately 75% of the total accounts have PDs lower than 0.04.

There are many factors that influence the direction and the extent of the relations between the developments in defaults and the risk factors influencing solvency. These comprise the macro indicators, market factors and idiosyncratic factors. For example, foreign clients tend to be more risky than domestic; the risk of default is expected to grow with increasing interest rates, unemployment rate or percentage deviation between the exchange rate level at which the individual loan was granted from the actual exchange rates.

7.2 Quantification of Model Risk

The question of understanding the impact of the uncertainty surrounding PD on the capital calculation can be approached by obtaining a mathematical representation of the underlying statistical distribution over the portfolio (see Figure 1). A change in the model assumptions or segmentation

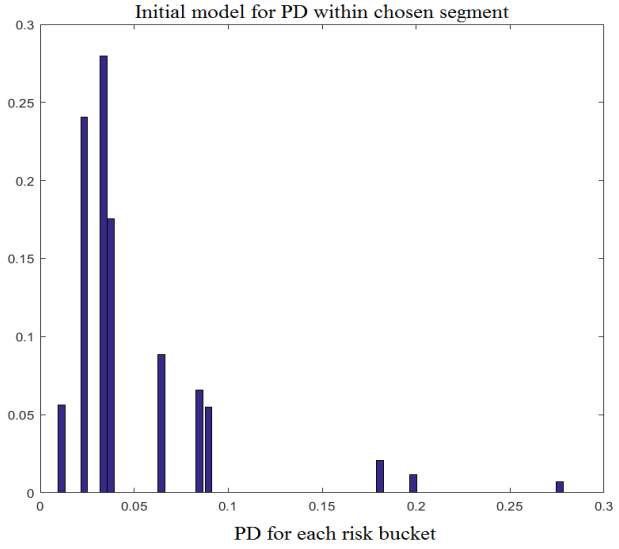

Figure 1: Distribution of the PD within the portfolio

would represent a shift in the parameters defining the statistical distribution. [14] We propose to predict

[14] Note that we can equally well work with empirical distributions (histograms) and use the Fisher–Rao distance between them [19]. However, to properly illustrate the aforementioned framework we approximate the portfolio by a parametric probability distribution.

the best future values of our given portfolio by using the 2–parameter Inverse Gaussian distribution (ING) given by

$$p(x, \lambda, \mu) = \left(\frac{\lambda}{2\pi x^3} \right)^{1/2} \cdot \exp\left(-\frac{\lambda}{2\mu^2 x}(x-\mu)^2 \right)$$

with mean μ and variance $\sigma^2 = \frac{\mu^3}{\lambda}$.[15] This means that we consider our model space to be a manifold of the ING distribution family, $\mathcal{M}$, parametrized by (λ, μ).

The initial model, $p_0 \in \mathcal{M}$, was estimated via Maximum Likelihood approach (MLE) with parameters $\lambda_0 = 0.5418$ and $\mu_0 = 4.8075$, i.e. $p_0 = ING(\lambda_0, \mu_0)$. With this representation we take into consideration how specific risk buckets perform within the default portfolio or, alternatively, the weight of default of each grade within the portfolio.

The PD model allows one to estimate the probability of default for a particular segment across different risk–buckets within the limits of defined parameters. The model represents a probability distribution of several correlated and mutually interacting events. Note that the parameters of the distribution are non–linear functions of the individual PD and N for each risk bucket, i.e. $\mu(PD, N)$ and $\lambda(PD, N)$. For the quantification of model risk, it is important to understand how the model integrates the information provided by the sample and how sensitive it is to changes in the sample characteristics (see Figure 2).

7.3 Identification of potential sources of model risk

To properly assess the model risk, we first need to identify and describe the potential sources of model risk, covering the model assumptions, weaknesses and arbitrariness in the development process. For the given PD model, the main risk may arise from the applied segmentation and risk scoring, bias toward historical experience, unexpected moves in the exchange rate, relevant changes in macro variables, missing values, incorrect structure and methodology, inappropriate factors used for the TTC adjustment, or due to the subjective selection of some parameter values.

The impact of these risk sources should be expressed quantitatively and then combined based on whether the separate model risks are independent, or some amplify each other, or counteract or absorb each other. In what follows and for the sake of illustration, we concentrate on the risk inherent in the applied segmentation arising from the PD estimation. The given model, p_0, is only one possible description of the data incorporating the expert knowledge and assumptions directed by the expected use, as well as implementation constraints. Any uncertainty and errors in the estimation of PD influence the required capital.

[15] The Inverse Gaussian density function represents a wide class of distributions, ranging from highly skewed distribution to a symmetrical one as λ/μ varies from 0 to ∞. For more details about the Inverse Gaussian distributions see for example [24].

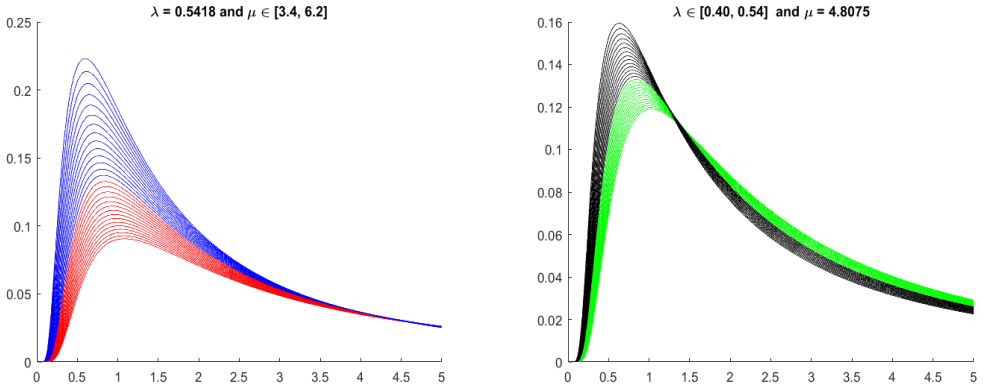

Figure 2: Probability density functions of Inverse Gaussian distributions: The left panel shows densities for different μ with $\lambda = 0.5418$. The right panel shows densities for different λ for $\mu = 4.8075$. The densities are unimodal with mode between 0 and μ. As μ/λ increases the distribution becomes more skewed to the right and the mode decreases relative to the mean.

7.4 Choosing a suitable Riemannian metric

We develop our example on the assumption that the model space should resemble the metric properties of a Riemannian manifold of negative curvature equipped with the Fisher–Rao information metric. Since there are many possible choices of metrics on a given differentiable manifold, it is important to consider the additional motivating properties of the Fisher–Rao information metric:

1. Fisher–Rao information metric is based on the notion of maximum entropy, which might be geometrically interpreted by the possible trajectory in a statistical manifold that describes its evolution. We can examine the degree of uncertainty by measuring the evolution of entropy in the model space.

2. The Fisher–Rao information may be thought of as the amount of information a sample supplies with respect to the problem of estimating the parameters. [16]

3. The Fisher–Rao information metric determines a constant negative Gaussian curvature for the ING distribution (for details see [7]). The metric is given by

$$g(\lambda, \mu) = \begin{bmatrix} \frac{1}{2\lambda^2} & 0 \\ 0 & \frac{\lambda}{\mu^3} \end{bmatrix}.$$

   Negative curvature guarantees that MLE estimates are well–defined and unique, [22].

In this metric, the model space has non–zero curvature and reflects model specification sensitivity accurately.

[16] Since the MLE is asymptotically unbiased, the inverse Fisher information represents the asymptotic fluctuations of the MLE around the true value.

7.5 Neighbourhood selection

Following Section 4, defining a proper neighbourhood around the given model is a trade–off between plausibility and severity in order to ensure that no harmful but plausible risks are missed and the irrelevant risk factors are not included. As aforementioned, we focus on the sensitivity of the chosen segmentation and the inherent uncertainty. Namely, we analyze the model risk arising from the uncertainty about the default counts, and so, we ask how different could the model be if the default counts were undercounted or overcounted under the assumption of the fixed number of risk buckets.

The default rates are calculated on the risk bucket level, and so the miss–estimation in each bucket is going to contribute to the overall uncertainty. Therefore, we consider all possible deviations within $[-0.3, 0.3]$ standard deviation changes in each of the ODF_i, $i = 1, \ldots, 10$. [17] We take into account all possible shifts in each bucket separately as well as all of the relevant combinations. The limit on the deviation was set to preserve the heterogeneity among risk buckets in terms of PD and to restrict the variability of nearby risk buckets, which could be influenced by the presence of outliers, noise or variation in the density of the points on the manifold. [18]

for the given manifold $\mathcal{M}$ and the point $p_0 \in \mathcal{M}$, in order to build the neighbourhood U we start by fitting probability distributions to all of the combinations of the maximal ± 0.3 standard deviation changes in the corresponding ODFs (i.e. 2^{10} distributions), and first to 1 000 000 and then to 10 000 000 random combinations within the interval $(-0.3, 0.3)$. The scale parameter varies in the interval $[\lambda_1, \lambda_2] = [0.517300, 0.565403]$ and the mean is in the interval $[\mu_1, \mu_2] = [4.58638, 4.95968]$. These fitted distributions form our chosen neighbourhood. [19] For illustration, Figure 3 represents the 3–dimensional embeddging of U, the set of alternative models around p_0, determined by the estimated ING distributions.

Within this chosen neighbourhood we can analyze and quantify the possible impact on the individual shifts but also all of their possible combinations on the model output that might be capable of causing material model movements. Notice the conservatism in the chosen neighbourhood.

Remark: Different adjusted data may result in the same distributions. Considering only movements in the distribution guarantees that the quantification of model risk will not depend on double count risk factors that are highly correlated with factors already included in the analysis.

7.6 Choosing an appropriate weight function

In what follows, we show how to estimate the weight function entirely from the data without imposing any additional constraint or particular structure. Roughly speaking, once we assume that

[17] Recall that ODFs refer to the maximum likelihood estimators of the binomial distributions for each risk bucket.

[18] The given portfolio (sample) is going to constrain the specificity and sensitivity possible.

[19] The key idea is to leave the data to determine the neighbourhood, instead of imposing one. With increasing number of sampled distributions we can see convergence also in the weight function (see Subsection 7.6) but also in the overall calculated Model Risk (see Subsection 7.7).

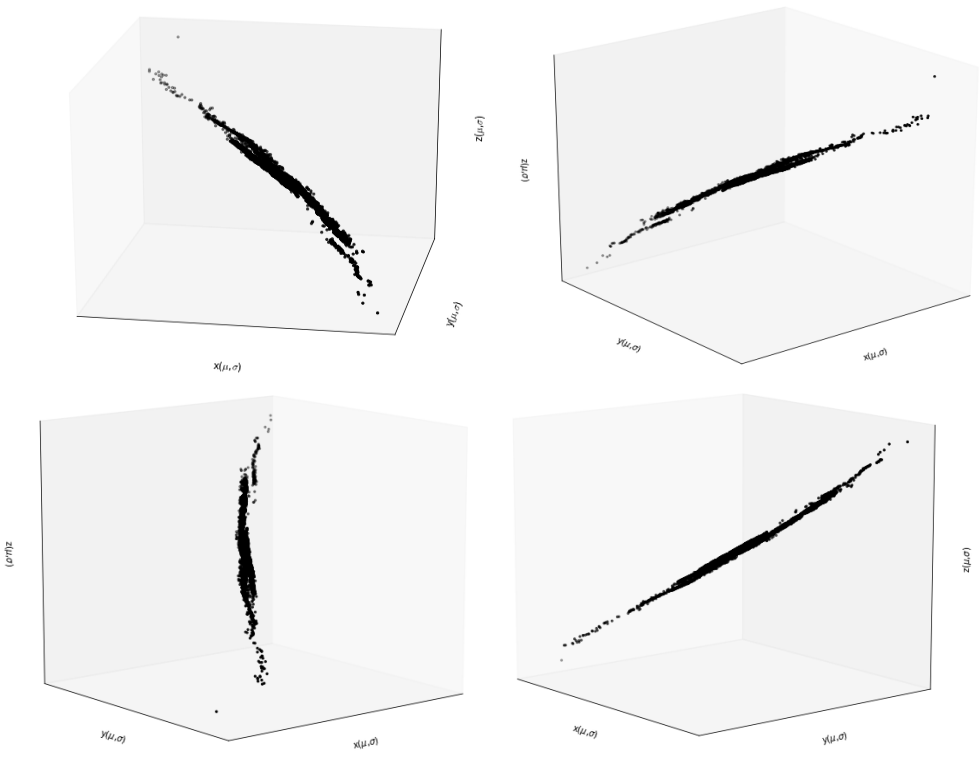

Figure 3: 3D multidimensional scaling embedding of U.

the alternative models are random objects valued on the Riemannian manifold, the weight function estimation consists of having the alternative models within U "contribute" to the estimate at a given point according to their distances from p_0.

More precisely, let $(\mathcal{M}, g)$ be a Riemannian manifold and let us consider $p_1, \ldots, p_n$ independent and identically distributed random points on $\mathcal{M}$ with density function $K(p)$. The estimate of $K(p)$ is then a positive function of the geodesic distance in U, which is then normalized by the volume density of (U, g) to take into account for curvature. These estimators are an average of the weights depending on the distance between p_i and p_0. Formally, using alternative models $p_1, \ldots, p_n$ we propose to define the weight function by a map $K_n : p \in U \to K_n(p) \in \mathbb{R}$ given by

$$K_n(p) = \frac{1}{\eta_{p_0}(p)}\, w_i \quad \forall i = 1, \ldots, n \tag{8}$$

where $\eta_{p_0}(p)$ denotes the volume density function and w_i are assigned weights, such that $\sum_{i=1}^{n} w_i = 1$.

The weights are determined as follows. First, we calculate the geodesic distance of all of the estimated distributions inside U from p_0. Next, we determine the number of levels m with respect to the maximal distance from p_0, i.e.

$$L(\lambda_i) = \{p \in U \mid \lambda_{i-1} < d(p_0, p) \leq \lambda_i\} \tag{9}$$

for a sequence $\lambda_1, \ldots, \lambda_m$ with $\lambda_m = \max d(p_0, p)$ and $\lambda_i - \lambda_{i-1} = \lambda_m/m$. We set $m = 5\,000$ in case of $1\,001\,024$ distributions inside U and $m = 50\,000$ for $10\,001\,024$ alternative models inside U. Next, we examine the number of distributions within all of these m level sets and based on the concentration within these levels we calculate normalized weights, i.e. w_i. The normalized average frequency of alternative models within U for $1\,001\,024$ and $10\,001\,024$ with respect to the distance from p_0 are illustrated in Figs. 4 and 5, respectively. The resulting normalized weights are then multiplied by the associated values of the volume density at $p \in U$.

Remark: The initial model and data single out a particular choice of a weight function through the concentration and variations of the fitted probability distributions on the manifold.

7.7 Measure of Model Risk

The last step in our proposed framework is the choice of a proper norm and the evaluation of the identified model risk described in Section 6. We suggest to use the $L^2(\mathcal{M})$ norm that in our particular example guarantees the consistency with the Maximum likelihood used for estimation and amplifies big changes in the outputs, i.e. capital level. With this choice, we obtain [20]:

$$Z\big(C(p), p_0\big) = \|C(p) - C(p_0)\|_2 = 2.35 \times 10^{-4}$$

[20] Model Risk calculated with $1\,001\,024$ distributions with the weight function based on $5\,000$ level sets equals $2.33499704056 \times 10^{-4}$ and Model Risk based on $10\,001\,024$ distributions inside U with $50\,000$ distance levels equals $2.35152164064 \times 10^{-4}$.

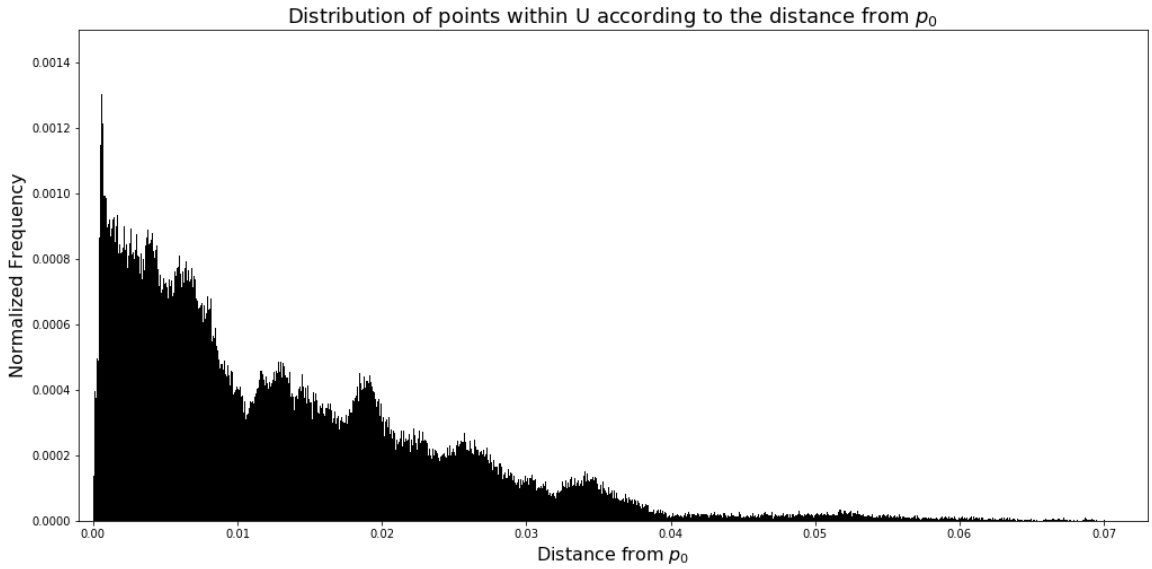

Figure 4: The figure represents the weight function given by the average concentration of the alternative models with respect to their geodesic distance from p_0 for $1\,000\,000 + 1024$ alternative models with $5\,000$ level sets.

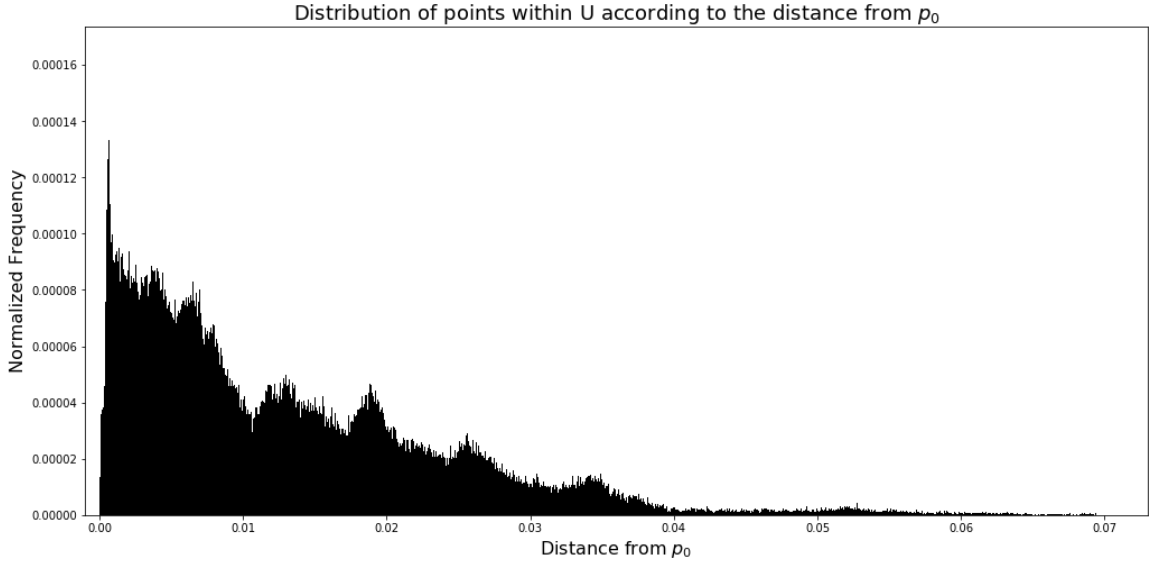

Figure 5: The figure represents the weight function given by the average concentration of the alternative models with respect to their geodesic distance from p_0 for $10\,001\,024$ alternative models with $50\,000$ level sets.

This means that the model risk from segmentation represents the 0.0235% variation of capital inside U.[21] The model risk is obtained in terms of capital, thus allowing an easy interpretation.

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 provide a sound mathematical definition of model risk using 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 complementary 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 financial institutions: 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 by naming just 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.

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 be a means to define the lognormal family).

[21] Note that this refers only to one source of model risk inherent in the model for capital calculation.

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.

Acknowledgements

The authors would like to express their sincere gratitude to José Carlos Colas Fuentes, Olivia Peraita Ezcurra and Inmaculada González Peréz from the Credit Risk models validation department at Banco Santander for providing insight and feedback that greatly assisted the example in Section 7.
This work has been partially funded by EU H2020-MSCA-ITN-2014 (WAKEUPCALL Grant Agreement 643045), Spanish MINECO (Grant MTM2016-76497-R) and by Xunta de Galicia (Grant CN2014/044).

References

[1] Amari, S.I.; Barndorff-Nielsen, O.E.; Kass, R.E.; Lauritzen, S.L., and Rao, C.R. *Differential geometry in statistical inference.* IMS Lecture Notes-Monograph Series 10, Hayward, California, USA (1987).

[2] Azzalini, A. "A class of distributions which includes the normal ones." *Scandinavian Journal of Statistics*, 12 (1985), 171–178.

[3] Bluhm, C.; Overbeck, L., and Wagner, C. *Introduction to credit risk modeling.* Florida, USA: Chapman & Hall/CRC Press (2016).

[4] Boucher, C. M.; Danelsson, J.; Kouontchou, P.S., and Maillet, B.B. "Risk models-at-risk." *Journal of Banking & Finance*, 44(1) (2014), 72–92.

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

[6] Chavel, I. *Riemannian geometry: a modern introduction.* Cambridge Studies in Advances Mathematics 98, Cambridge, UK: Cambridge University Press (2006).

[7] Chen, W.W.S. "Finding Gaussian curvature of lifespan distribution." *Applied Mathematics*, 5 (2014), 3392–3400.

[8] Christodoulakis, G., and Satchell, S. "The validity of credit risk model validation methods." In *The Analytics of Risk Model Validation*, G. Christodoulakis and S. Satchell editors, Burlington, MA: Academic Press (2008), 27-44. .

[9] Cont, R. "Model uncertainty and its impact on the pricing of derivative instruments." *Mathematical Finance*, 16 (2006), 519-547.

[10] 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, 117.

[11] Federer, H. *Geometric measure theory.* Berlin, Germany: Springer Verlag (2014).

[12] Gibson, R. *Model risk: concepts, calibration and pricing.* London, UK: Risk Books (2000).

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

[14] Joshi, S.; Srivastava, A., and Jermyn, I.H. "Riemannian analysis of probability density functions with applications in vision." In *2007 IEEE Conference on Computer Vision and Pattern Recognition (CVPR2017) Proceedings.* Piscataway, NJ, USA: IEEE (2007), 1664–1671..

[15] Kiefer, N. M., and Larson, C.E. "Biases in default estimation and capital allocations under Basel II." Unpublished Monograph, retrieved on November 13 (2012).

[16] Lang, S. *Fundamentals of differential geometry.* Graduate Texts in Mathematics, 191. New York, USA: Springer Verlag (2012).

[17] Morgan, F. "Manifolds with density." *Notices of the AMS*, 52 (2005), 853–858.

[18] Morini, M. *Understanding and Managing Model Risk: A practical guide for quants, traders and validators.* Chichester, West Sussex, UK: John Wiley & Sons (2011).

[19] Navab, N.; Pluim, J.P.W., and Viergever, M.A. *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2010).* Berlin, Germany: Springer-Verlag (2010).

[20] Murray, M. K., and Rice, J.W. *Differential geometry and statistics.* Monographs on Statistics & Applied Probability, 48. Florida, USA: CRC Press (1993).

[21] Rao, C. R. "Information and accuracy attainable in the estimation of statistical parameters." *Bull. Calcutta. Math. Soc.*, 37(3) (1945), 81–91.

[22] Said, S.; Bombrun, L.; Berthoumieu, Y., and Manton, J. H. "Riemannian Gaussian distributions on the space of symmetric positive definite matrices." *IEEE Transactions on Information Theory*, 63(4) (2017), 2153–2170.

[23] Saltelli, A.; Tarantola, S.; Campolongo, F., and Ratto, M. *Sensitivity analysis in practice: A guide to assessing scientific models.* Chichester, UK: John Wiley & Sons (2013).

[24] Tweedie, M. C. K. "Statistical properties of inverse Gaussian distributions I." *The Annals of Mathematical Statistics*, 28 (1957), 362–377.

[25] Willmore, T. *Riemannian geometry.* Oxford, UK: Oxford University Press (1993).

9 Appendix

In this Appendix we present the construction of the weight function and the associated proofs.

9.1 Weight Function Construction

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 p_0 and by the vectors tangent to the unique geodesics γ that pass through them. In order 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 $T_{p_0}\mathcal{M}$. The boundaries of C^n are

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

and represent the end points of the geodesics, i.e. $\partial_1 C^n$ will be transformed into p_0 and $\partial_2 C^n$ into ∂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 obtain a weight function satisfying all required properties $(K1) - (K3)$.

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. Since U is a subset of the normal neighbourhood of p_0, the exponential map is well–defined and it defines a local diffeomorphism from $T_{p_0}U$ to U. The geodesics γ are then given in these coordinates by rays emanating from the origin.

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

$$\begin{aligned} &P:\ [0,\infty) \times \mathcal{S}(p_0,1) \to T_{p_0}\mathcal{M}:\ \ P(t,v) = tv, \\ &P^{-1}:\ T_{p_0}\mathcal{M}\backslash\{0\} \to (0,\infty) \times \mathcal{S}(p_0,1):\ \ P^{-1}(v) = \Big(\|v\|, \frac{v}{\|v\|} \Big). \end{aligned} \tag{10}$$

In order to properly describe the neighbourhood U, we define the distance function to the points on the boundary of the neighbourhood along the geodesic $\gamma_v(t)$. For $p_0 \in \mathcal{M}$ and a unit vector $v \in \mathcal{S}(p_0, 1) \subset T_{p_0}\mathcal{M}$, we define the distance $\rho(v)$ by

$$\rho(v) = \sup_{t>0}\{tv \in \mathcal{U} \mid d(p_0, \gamma_v(t)) = t\},$$

as the distance from the origin to the boundary $\partial\mathcal{U}$ in direction $v \in \mathcal{S}(p_0, 1)$, i.e. $\mathcal{U} = \{tv : 0 \leq t \leq \rho(v), v \in S(p_0, 1)\}$. In other words, $\rho(v)$ is the maximal distance in direction v to the boundary ∂U. The point $\gamma_v(\rho(v))$ is the boundary point of U along the geodesic $\gamma_v(t)$. The geodesic $\gamma_v(t)$ minimizes the distance between p_0 and $\exp_{p_0}(tv)$ for all $t \in [0, \rho(v)]$. The distance ρ, considered as a real valued function on $\mathcal{S}(p_0, 1)$, is strictly positive and Lipschitz continuous on $\mathcal{S}(p_0, 1)$.

We now define the coordinate transformation

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

The mapping $\mathbb{S}^{n-1} \to \mathcal{S}(p_0, 1)$, $\nu \mapsto v$ 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 $v \in \mathcal{S}(p_0, 1) \subset T_{p_0}\mathcal{M}$. Since the distance $v \mapsto \rho(v)$ 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 $[0, \rho(v)] \times \mathcal{S}(p_0, 1)$.

The composition $\exp_{p_0} \circ P\Lambda_\rho$ defines a mapping from C^n onto U that maps $\partial_1 C^n$ onto the point $\{p_0\}$ and the right hand side boundary onto ∂U. Thus, the image of a continuous function f on U is also continuous on C^n, but not every function h on C^n is the image of a continuous function on U under the pull back operator $\Lambda_\rho^* P^* exp_{p_0}^*$. In order to identify a continuous function h on C^n with a continuous function f on U, the function h has to satisfy additional consistency conditions.

Definition 9.1 *A continuous function h defined on a cylinder C^n is called consistent with a continuous function f on U under the mapping $\exp_{p_0} \circ P\Lambda_\rho$ if $h = f(exp_{p_0} \circ P\Lambda_\rho)$. In this case, h satisfies the following consistency 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 $\partial_1 C^n$. When the function h on C^n is consistent with f, the constant value at $\partial_1 C^n$ 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 $\exp_p \circ P\Lambda_\rho$ 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 9.2 *The existence of the weight function K satisfying assumptions $(K1) - (K3)$ is equivalent to assuming the existence of a consistent function h defined on C^n with codomain $\mathbb{R}$ that satisfies the following properties:*

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

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

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

Proof of Lemma 9.2. In order to prove the stated equivalence, we need to show that the function h defined in Lemma 9.2 preserves the required properties of K under the continuous mapping $\exp_{p_0} \circ P\Lambda_\rho$. First, we show that the composition of three different mappings is well defined.

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

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

where $\mathbb{S}^{n-1} := \{\nu \in \mathbb{R}^n \mid \|\nu\|^2 = \nu_1^2 + \cdots + \nu_n^2 = 1\}$ 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

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

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 introduce a distance function to the points on the boundary of the neighbourhood along the geodesic $\gamma_v(t)$. For $p_0 \in \mathcal{M}$ and a unit vector $v \in \mathcal{S}(p_0,1) \subset T_{p_0}\mathcal{M}$, we define the distance $\rho(v)$ by

$$\rho(v) = \sup_{t>0}\{tv \in \mathcal{U} \mid d(p_0, \gamma_v(t)) = t\}$$

as the distance from the origin to the boundary $\partial\mathcal{U}$ in direction $v \in \mathcal{S}(p_0,1)$, i.e. $\mathcal{U} = \{tv : 0 \le t \le \rho(v), v \in S(p_0,1)\}$. In other words, $\rho(v)$ is the maximal distance in direction v to the boundary ∂U. The point $\gamma_v(\rho(v))$ is the boundary point of U along the geodesic $\gamma_v(t)$. The geodesic $\gamma_v(t)$ minimizes the distance between p_0 and $\exp_{p_0}(tv)$ for all $t \in [0,\rho(v)]$. The distance ρ, considered as a real valued function on $\mathcal{S}(p_0,1)$, is strictly positive and Lipschitz continuous on $\mathcal{S}(p_0,1)$.

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

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

with boundaries

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

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:

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

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 consider the polar transformation P defined by equation 10 which 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 $C^n_\rho \backslash \{\partial_1 C^n_\rho, \partial_2 C^n_\rho\}$ onto the open set $\mathcal{U}\backslash\{0, \partial\mathcal{U}\}$. Combining P with the exponential map $\exp_{p_0}$, we have

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

The composition $\exp_{p_0} P$ 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 [22]. Then the points $(t,v) \in C^n_\rho$ induce geodesic polar coordinates on $\mathbb{R}^n$.

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 $\exp_{p_0} \circ P\Lambda_\rho$ is a continuous mapping from C^n onto U. Moreover, $\exp_{p_0} \circ P\Lambda_\rho$ 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 $\exp_{p_0} \circ P\Lambda_\rho$ preserves connectedness and compactness and it a continuous mapping from C^n onto $\mathcal{M}$. Moreover, $\exp_{p_0} \circ P\Lambda_\rho$ 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 $f(\exp_{p_0} \circ P\Lambda_\rho)$ 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 9.1 is the image of a continuous function on $\mathcal{M}$ under the pull–back operator $\Lambda^*_\rho P^* \exp^*_{p_0}$.

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

The composition $\exp_{p_0} \circ P\Lambda_\rho$ applied to a function h that is continuous on C^n and satisfies the consistency conditions $(i) - (ii)$ of Def.9.1 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 $\exp_{p0} \circ P\Lambda_{\rho(v)}$.

- 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 $\int_{\mathcal{M}} d\zeta = 1$.

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

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 $\eta_{p_0}(v)$ represents the density function that is a positive and continuously differentiable function on $U \subset T_{p_0}\mathcal{M}$. Further, we have

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

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:

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

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

$$\begin{aligned} & \int_{S^{n-1}} \int_0^1 \frac{1}{\rho(\nu)} \frac{1}{\eta_{p_0}(\rho(\nu)t\nu)} t^{1-n}\rho(\nu)^{2-n} h\left(\frac{t}{\rho(\nu)}, \nu\right)\big(\rho(\nu)t\big)^{n-1}\eta_{p_0}(\rho(\nu)t\nu)dtd\mu(\nu) \\ = & \int_{S^{n-1}} \int_0^1 h\Big(t, \nu\Big)dtd\mu(\nu) = 1. \end{aligned}$$

□

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 9.3 *A continuous function h defined on C^n satisfying conditions $(H1)-(H5)$, determines a unique and well–defined weight function K satisfying $(K1)-(K3)$ on U. More precisely, the map $K: U \to \mathbb{R}$, which to each point p, associates the value $K(p)$ given by*

$$K(p) = \frac{1}{\eta_{p_0}(p)} t^{1-n} \rho(v)^{2-n} h\left(\frac{t}{\rho(v)}, v\right), \tag{11}$$

where $\eta_{p_0}(p)$ is the volume density with respect to p_0 [23]*, v is the tangent vector, $t \in [0, \rho(v)]$ is a scaling parameter, and $\rho(v)$ is the distance function defined above.*

Proof of Theorem 9.3. Note that the composition $\exp_{p_0} \circ P\Lambda_\rho$ induces a change of variables for integrable function f that yields to the following formula:

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

The volume density η_{p_0} is a well–defined, non–negative function in a neighborhood of p_0 (globally it may be defined using the Jacobi fields ([25], p.219)). Besides, η_{p_0} is a continuous and differentiable function on $\mathcal{M}$[24]. 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 $1/\rho(\nu)$. Therefore, the mapping Λ_ρ defines a bi–Lipschitzian mapping from C^n to C^n_ρ. Moreover, the composition $\exp_{p_0} \circ P$ defines a diffeomorphism from C^n_ρ $\{\partial_1 C^n_\rho, \partial_2 C^n_\rho\}$. Then, 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 $\exp_{p_0} \circ P\Lambda_\rho$ is an isomorphism. Then, for any h defined on C^n satisfying conditions $(i)-(ii)$ of Definition 9.1, the associated weight function K is well defined and continuous on U. The uniqueness of K follows after specifying a function h that satisfies properties $(H1)-(H5)$. □

[23] In terms of geodesic normal coordinates at p_0, $\eta_{p_0}(p)$ equals the square++root of the determinant of the metric g expressed in these coordinates at $\exp_{p_0}^{-1}(p)$. For p and q in a normal neighbourhood U of $\mathcal{M}$, we have $\eta_p(q) = \eta_q(p)$ [25].

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