Why insurance regulators need to require sensitivity settings of internal models for their approval

Emanuele Borgonovo [a], Gian Paolo Clemente [b], Giovanni Rabitti [c,*]

[a] *Department of Decision Sciences, Bocconi University, Milan, Italy*
[b] *Department of Mathematics for Economics, Financial and Actuarial Sciences, Università Cattolica del Sacro Cuore, Milan, Italy*
[c] *Department of Actuarial Mathematics and Statistics, Heriot-Watt University, and Maxwell Institute for Mathematical Sciences, Edinburgh, UK*

A R T I C L E   I N F O

A B S T R A C T

According to the Solvency II directive, insurers can use internal models for solvency assessment, but regulators must approve these models. Sensitivity analysis is a crucial part of the approval process. However, the directive lacks clarity on the required sensitivity analysis. Various techniques exist in literature to assess the impact of model assumptions on output, each revealing different aspects of model behaviour. In this letter, we suggest a minimum standard for regulators to ensure model quality. We propose complementary sensitivity settings for internal model development, governance, and approval. Implementing these settings enhances the explainability of approved models and their reliability.

1. Introduction

In recent years, scenario analyses and stress testing have gained increasing interest in the insurance industry (CRO Forum, 2023). The Solvency II regulation (European Commission, 2009) emphasizes the importance of assessing the impact of changes in key assumptions or variables on an insurance company's solvency capital requirement (SCR). The SCR calculation relies on multiple assumptions and inputs, such as the company's risk profile, business nature, and operating economic conditions. These assumptions and inputs are subject to uncertainty, and proper analyses must be developed to identify which factors have the greatest impact on the SCR and their magnitude. Additionally, a key component of the framework is the Own Risk and Solvency Assessment (ORSA), a forward-looking risk management process that requires insurance companies to assess their overall risk profile and solvency position. In particular, assessing the impact of adverse events on financial position aids firms in effective risk management and informed decision-making about capital and investment strategies.

The assessment of a company's overall risk profile involves simulation models with inputs such as economic, financial, and actuarial variables. However, as these models grow in complexity, their internal workings often become opaque, resembling black-boxes. Consequently, policymaking based on model outputs become questionable. It is also worth mentioning that the use of machine learning and artificial intelligence is gaining significant importance in the development and implementation of internal models under Solvency II. This trend is expected to grow in the future, as companies increasingly adopt a business model characterized by fintech and insurtech components. These methodologies offer significant enhancements to internal models by enabling advanced predictive modelling and capturing complex relationships among diverse risk factors. Yet, the integration of machine learning algorithms introduces additional complexities and challenges to the model validation process. The utilization of such models amplifies the black-box effect of internal models, making it difficult to understand the reasoning behind their

---

predictions. Thus, exploring and validating the underlying factors and assumptions influencing models' outputs become essential before than making inference from these models. This step, known as interpretability and explainability analysis, ensures alignment between decisions, risk management objectives and compliance with regulatory requirements, as, for instance, remarked by the Solvency II directive (European Commission, 2015a) stressing profit and loss attribution in the internal model validation. Indeed, the purpose of this test is to ensure that the internal model captures the main sources of profit and loss that firms face.

In this context, to make the models transparent and interpretable to the decision-maker, best practice recommends to adopt proper uncertainty quantification and sensitivity analysis for model interpretation and explanation (Saltelli et al., 2000, 2008, 2020). Sensitivity analysis can support model explanation by quantifying how variations in the input variables affect the model response. Understanding which inputs affect the output and how they do so is crucial for model interpretability.

The evolving methods for analysing complex models aim to quantify assumption impacts and allocate output uncertainty to specific drivers. Uncertainty quantification involves propagating variability in model output stemming from uncertainty in input variables and assumptions. Sensitivity analysis complements uncertainty quantification by identifying the primary causes of variations in model output (Saltelli et al., 2019). Instead, global sensitivity analysis allocates uncertainty to individual inputs or interactions, while local sensitivity analysis investigates input variables responsible for output variations across scenarios. Recent literature has introduced alternative approaches, as we will outline in this letter. The literature on sensitivity analysis is, indeed, primarily scattered across the fields of environmental sciences, statistics, and operational research, with most papers typically concentrating on specific aspects of a model (Borgonovo and Plischke, 2016; Razavi et al., 2021). Instead of being redundant when combined, they produce complementary insights. Hence, uncertainty analysis and sensitivity analysis are essential elements for an effective risk management and for a proper monitoring of the company's solvency ability. These settings can be adopted by the financial industry, regulators and other fields too.

Nonetheless, these terms are sometime misunderstood and conflated in the academic literature (Saltelli et al., 2019). This happens in the actuarial literature as well. For instance, Booth et al. (2005) write that *"sensitivity analysis, in order to gain some insight into the range of possible outcomes. This involves making different model [...] projections on the basis of a range of assumptions for the underlying parameters. [...] The assumed variations would aim to encompass the main range of future outcomes that the actuary would consider as feasible or probably under certain conditions. The sensitivity analysis brings a number of additional benefits [...] It will give the idea of the increasing uncertainty of any projection the further into the future the projection is made [...] It enables the actuary to assess the relative importance of each factor to the solvency of the office"* (Booth et al., 2005, page 356). However, they call this approach *"deterministic sensitivity assessment"*, mixing the notions of uncertainty analysis, local (deterministic) sensitivity analysis and global sensitivity analysis.

At the same time, insurance authorities have yet to establish a formal, comprehensive framework for sensitivity analysis techniques. While the Solvency II regulation prescribes the use of sensitivity analysis, it lacks specific best-practices, letting individual companies responsible for developing their own sensitivity assessment. The absence of a unified common framework contributes to market fragmentation, resulting in non-uniform standards of internal model quality and reliability. The lack of a regulatory guidance coupled with the complexity of internal models exacerbate this issue.

Implementing sensitivity analysis techniques provides substantial benefits for model validators. It helps them to understand and evaluate the inputs, assumptions, calculations, and outputs of internal models, ensuring the accuracy, robustness, and adherence to regulatory requirements. Using explainability methods, validators can assess whether the model aligns with the intended risk measurement objectives and capital calculation. Therefore, validating internal models, particularly those integrating machine learning, requires a multifaceted approach. This involves combining statistical analysis, stress testing, sensitivity analysis, and interpretability assessment. To this end, in this letter we present up-to-date sensitivity indices and describe how each sheds light on specific aspects of the model's behaviour. This framework aims to propose a unified benchmark for sensitivity analysis in internal models, built upon complementary sensitivity settings.

2. Internal models and regulators

The Directive Solvency II (European Commission, 2009) allows insurance and reinsurance companies to adopt internal models for the assessment of their own capital requirement, rather than relying solely on the standard formula provided by the regulator authorities (European Commission, 2015a).

The use of internal models in Solvency II should provide insurance companies with a more flexible and accurate way to manage their risk and allocate capital, but, at the same time, it requires significant investment in data, systems, and expertise, and is subject to strict regulatory oversight. Indeed, once the internal model has been developed, the insurance company must seek approval from the regulatory authorities. This involves providing detailed documentation of the model and its underlying assumptions (EIOPA, 2016), as well as demonstrating its accuracy and reliability through backtesting and validation (see EIOPA (2022)).

Therefore, a validation process must be developed (see European Commission (2015b)) in order to ensure that the up-to-date risk profile has been considered in the modelling and to verify the process for choosing assumptions and using expert judgment (EIOPA, 2022). Validation is an iterative process of identification of model limitations and the implementation of possible improvements.

The validation policy must set out the validation tools that are used, i.e. proper approaches designed to evaluate that the internal model is appropriate and reliable. Regulatory technical standards (RTS) and implementing technical standards (ITS), under the level 2 guidance, define as mandatory the inclusion of sensitivity analysis and backtesting for each material risk in order to validate the model. Hence, sensitivity analysis is a critical component of the internal model used for Solvency II compliance (Loisel and Nisipasu, 2016). It involves testing the sensitivity of the model's output to changes in the underlying assumptions and parameters, in order to

assess the model's robustness and reliability. This includes assessing the impact of variations in key risk factors, such as interest rates, market prices, technical assumptions or mortality rates on the model's output (European Commission, 2009). However, the regulator does not prescribe a specific method for sensitivity analysis. Instead, Solvency II provides general guidance on the requirements for sensitivity analysis and the factors that should be considered when performing such analysis. Some general indications are also given in the documents regarding the system of governance (see EIOPA (2014)), where it is required that insurance companies should have a clear methodology for sensitivity analysis, which should be proportionate to the complexity of the internal model and the risks being modelled.

It is evident that when the solvency position is assessed using complex models, it becomes essential to understand the impact that the assumptions used in the model construction have on the reliability of the model. This information is delivered by uncertainty and sensitivity analyses, which constitute important elements of the model building process. Neglecting a proper sensitivity analysis might lead to overconfident attitude by the model builders as well as to underestimation of the uncertainty (Saltelli et al., 2000). The importance of stressing both single risk drivers and the effect of their dependencies is indeed a crucial aspect that several papers are pointing out (see, e.g., CRO Forum 2013, Dacorogna 2017, International Actuarial Association 2013, Lloyds 2022, SCOR 2008). Art. 242 of Delegated Regulation (European Commission, 2015a) remarks indeed the relevance of testing at the level of single outputs as well as at the level of aggregated results. Additionally, the statistical process for validating the internal model shall include a reverse stress test, identifying the most probable stresses that would threaten the viability of the insurance or reinsurance undertaking. Although there is not much literature on internal insurance model validation, the importance of this process is well stressed in Dacorogna (2023).

3. Uncertainty and sensitivity settings

The issue of explaining complex models is a subject of ongoing debate in the literature, with numerous papers consistently proposing new strategies to shed light on the internal mechanisms of these models. However, a significant number of these approaches lack theoretical justification, while others offer only partial explanations. Only a limited number of works present a comprehensive strategy for investigating the model effectively, making best use of available techniques (Borgonovo and Plischke, 2016). In this letter we illustrate and integrate several up-to-date best practices in uncertainty and sensitivity analysis. As Dacorogna writes, "*The sensitivity analysis [of internal models] is important. It is not possible to base management decisions on results that could drastically change if some unimportant parameters are modified in the input. Unfortunately, note that this statement contains the adjective "unimportant", which is hard to define*" (Dacorogna, 2017, page 11). Depending on the context of analysis, the notion of "importance" can be formalized. There are two main settings for investigating the model behaviour:

1. scenario analysis,
2. uncertainty quantification,

The scenario analysis mode is also known as local sensitivity analysis and it involves the inspection of the model behaviour at a limited number of locations in the parameter space. The uncertainty quantification setting aims at exploring the response of the model for variations of the input parameters across the entire parameter space. The two modes are not mutually exclusive as we are to see. However, the analysis needs to be clear up front about these aspects.

Within these two settings, sensitivity measures can be used to answers sensitivity analysis goals (Borgonovo and Plischke, 2016; Borgonovo, 2023). The goals of a sensitivity analysis exercise should be specified up front, so that, on the one hand, the sensitivity analysis question is clearly specified and, on the other hand, the analyst is able to identify the sensitivity method that is most appropriate to answer the sensitivity question. The literature has identified alternative goals, we summarize the three main ones in this note:

1. factor prioritization,
2. trend identification,
3. interaction quantification.

Factor prioritization refers to the identification of the key-drivers of the model response, either in a local or global context. Trend identification refers to understanding the direction of change in the model response following a change in the inputs. Interaction quantification refers to understanding whether relevant interactions are present in the model response.

3.1. Local sensitivity analysis

Let us start with a scenario analysis. We write the internal model as $y = f(\mathbf{x})$, where y is the model output and $\mathbf{x} \in \mathbb{R}^d$ is the d-dimensional input vector representing economic, financial and actuarial variables. In general, y is the quantity of interest to the decision-maker. In this case, it is intended as a real number that can represent either the output of a deterministic model or a key indicator obtained by summarizing the distribution given by a stochastic model. For instance, the own funds, the SCR or the Solvency Ratio can represent possible outcomes of an internal model for an insurance firm.

In this analysis the analyst typically fixes a reference or base case scenario. In this scenario, the inputs are assigned a value $\mathbf{x}^0 = (x_1^0, x_2^0, \ldots, x_d^0)$. The internal model is then evaluated at this reference scenario. We denote with y^0 the base-case value of the model, that $y^0 = f(\mathbf{x^0})$. The stressed scenario is an alternative scenario in which the inputs are fixed at $\mathbf{x}^1 = (x_1^1, x_2^1, \ldots, x_d^1)$. The

value of the model output at the alternative scenario is then $y^1 = f(\mathbf{x}^1)$. In general, the quantity of interest undergoes a change across these two points, which we denote as

$$\Delta y = g(\mathbf{x}^1) - g(\mathbf{x}^0) \tag{1}$$

It turns out that this change can be exactly apportioned to the variation in each of the inputs as (Rabitz and Alis, 1999; Borgonovo, 2010)

$$\Delta y = g(\mathbf{x}^1) - g(\mathbf{x}^0) = \sum_{i=1}^{n} \phi_i^{\mathbf{x}^0 \to \mathbf{x}^1} + \sum_{i<j} \phi_{i,j}^{\mathbf{x}^0 \to \mathbf{x}^1} + \cdots + \phi_{1,2,\dots,n}. \tag{2}$$

In Eq. (2), the change in model response is decomposed in 2^n terms. The first order terms are given by

$$\phi_i^{\mathbf{x}^0 \to \mathbf{x}^1} = f(x_i^1 : \mathbf{x}_{-i}^0) - f(\mathbf{x}^0), \tag{3}$$

where $f(x_i^1 : \mathbf{x}_{-i}^0)$ denotes the model evaluated at the point having all coordinates equal to $\mathbf{x}^0$ except the ith one equal to x_i^1. These terms account for the individual effects of the input. We call them main effects in the remainder. They convey the information about what is the effect of having changed input x_i, while keeping all the others at the base-case scenario. They are frequently visualized by Tornado diagrams (Borgonovo and Rabitti, 2023). A Tornado diagram is a graphical representation of the local sensitivity results, showing the relative importance of a variable as a bar. Main effects are often used to quantify importance of x_i in driving the difference $f(\mathbf{x}^1) - f(\mathbf{x}^0)$ of the model output from the base-case scenario to the stressed scenario (Borgonovo, 2010). A sensitivity analysis based on main effects is also called a one-factor-at-a-time analysis, because all variables are shifted individually from their base-case values to the stressed ones. If we normalize the change in quantity of interest by the change in input variation, we obtain

$$v_i^{\mathbf{x}^0 \to \mathbf{x}^1} = \frac{\phi_i^{\mathbf{x}^0 \to \mathbf{x}^1}}{(x_i^1 - x_i^0)}. \tag{4}$$

When the difference between the base-case scenario $\mathbf{x}^0$ and the stressed scenario $\mathbf{x}^1$ is very small, the local indices $v_i^{\mathbf{x}^0 \to \mathbf{x}^1}$ tend to the partial derivatives $\partial g / \partial x_i$, which are known as sensitivities in the financial literature. In terms of goals of a sensitivity analysis, we can answer questions concerning factor prioritization and direction of change. Indications on factor prioritization are gained by considering the magnitude of $\phi_i^{\mathbf{x}^0 \to \mathbf{x}^1}$. Indications on the direction of change are obtained looking at $v_i^{\mathbf{x}^0 \to \mathbf{x}^1}$ (see Borgonovo and Rabitti (2023) among others for additional details). The first order indices do not deliver insights on interactions. These can be gained by considering higher order terms of the expansion in Eq. (2). For example, it is sometimes convenient to consider also the total finite-change effect $\phi_i^T = \phi_i^{\mathbf{x}^0 \to \mathbf{x}^1} + \sum_j \phi_{i,j}^{\mathbf{x}^0 \to \mathbf{x}^1} + \cdots + \phi_{1,2,\dots,n}$ which includes all the terms in which i contributes.

3.2. Uncertainty quantification and global sensitivity analysis

One-at-a-time sensitivity analysis is popular, intuitive and it is often computationally inexpensive. However, the indications it yields are local. The validity of any conclusion is limited to changes across the two scenarios. Insights from this deterministic approach are completely disregarding the uncertainty on future scenarios and provide only a partial exploration of the model input space.

In order to account for uncertainty, we need to move to the uncertainty quantification setting. In this setting, the random inputs are typically assigned probability distributions. They become random variables and, consequently, the internal model response $Y = g(\mathbf{X})$ becomes a random variable. The output Y is random because $\mathbf{X}$ is. Let us denote the probability distribution of Y by $\mathbb{P}_Y$, its cumulative distribution function by P_Y, its density function by p_Y, its variance by $\mathbb{V}[Y]$ and its mean value by $\mathbb{E}[Y]$. We then need to propagate uncertainty from Y to $\mathbf{X}$, in order to assess how uncertainty in the inputs reverberates on the output. This is typically performed by Monte-Carlo simulations, by bootstrapping market data or by simulating new economic scenarios (Flaig and Junike, 2022). Quantifying uncertainty is essential to avoid solvency decisions based on model point estimates.

We can then address the sensitivity analysis goals within an uncertainty quantification framework. Let us start with factor prioritization. When uncertainty in the inputs is considered, then factor prioritization means to identify the key-drivers of uncertainty in the model response. Towards achieving this first goal, we can start with the question: if we are informed that $X_i = x_i$ what is the effect of this information on the distribution of Y? Will the new distribution $\mathbb{P}_{Y|X_i}$ change significantly or will it be close to the current distribution of Y? If the distribution of Y does not change, then the variable is uninfluential (statistically independent). Thus, we are interested in quantifying the separation between the original distribution $\mathbb{P}_Y$ and the conditional distribution $\mathbb{P}_{Y|X_i}$. Such a separation can be written as $\zeta(X_i) = \zeta(\mathbb{P}_{Y|X_i}, \mathbb{P}_Y)$, where $\zeta(\mathbb{P}_{Y|X_i}, \mathbb{P}_Y)$ is some distance or divergence between probability distributions (let us leave it general for the moment, we will provide examples shortly). Because X_i is uncertain, we average over the possible values of X_i obtaining the global sensitivity index (Borgonovo et al., 2016)

$$\xi_i = \mathbb{E}[\zeta(\mathbb{P}_{Y|X_i}, \mathbb{P}_Y)]. \tag{5}$$

The global sensitivity index ξ_i can be used for comparing the influence of variables: if $\xi_i < \xi_j$, then X_i is less influential than X_j. There are several ways in which we can specify the separation measurement. An important aspect here is that the selection needs to match the statistical property of Y that the analyst or the decision-maker is considering in the risk assessment. For instance, the

decision-maker might be considering the mean, a quantile or the entire distribution of Y. Then, the separation measurement should lead to a sensitivity measure that matches the statistical property of the quantity of interest.

One popular way to consider the influence of X_i on Y is to quantify the separation between $\mathbb{P}_{Y|X_i}$ and $\mathbb{P}_Y$ as a reduction in the variance of Y. Let us consider the effect of learning X_i on the mean of Y. After learning X_i, the expected value of Y becomes $\mathbb{E}[Y|X_i]$. Pearson proposes to consider the variance of this conditional expectation, introducing the correlation ratio

$$S_i = \frac{\mathbb{V}(\mathbb{E}[Y|X_i])}{\mathbb{V}(Y)} = \frac{\mathbb{V}(Y) - \mathbb{E}[\mathbb{V}(Y|X_i)]}{\mathbb{V}(Y)} \tag{6}$$

This ratio has later on become the well-known variance-based importance measure, also known as the first order Sobol' sensitivity index (Sobol', 1993; Saltelli et al., 2008). The variance-based importance index S_i quantifies the contribution of the individual variable X_i to the output variance, as the expected fractional reduction in variance that we obtain if we get to know X_i. It is possible to extend the index S_i for the variable X_i to S_u for the importance of a group u of variables $\mathbf{X}_u = \{X_i : i \in u\}$, for any subset of variables $u \subseteq \{1, 2, \ldots, d\}$ (see Saltelli et al. (2008), Borgonovo and Plischke (2016) for technical details).

Under variable independence, it becomes meaningful to consider the total importance index

$$S_{T_i} = \frac{\mathbb{E}_{\mathbf{X}_{-i}}\left(\mathbb{V}_{X_i}(Y|\mathbf{X}_{-i})\right)}{\mathbb{V}(Y)} \tag{7}$$

which quantifies the contribution of X_i to the variance of Y including its main and interaction effects. The technical justification lies in the functional ANOVA expansion of the model f (see Saltelli et al. (2008) for more details). When the X_i's are independent, we have $S_i \leq S_{T_i}$ for all i, meaning that the total effect of X_i is always greater than its individual effect. The value $S_i \approx S_{T_i}$ can be interpreted in terms of a low relevance of interactions in the model response involving X_i. Moreover, if $S_{T_i} \approx 0$ then X_i can be frozen at a nominal value since it is not influencing Y. This helps the analyst to reduce the dimensionality of the model f.

When the variables $\mathbf{X}$ are assumed to be dependent, the Sobol' indices can be still defined but their interpretation in terms of explained variance is no more valid due to spurious correlations. For example, total effects can be even smaller than their corresponding individual indices. Two avoid the issues in the presence of dependent variables, two main approaches have been developed in the literature. A first approach is to use the Shapley effect for variable X_i defined as (Owen, 2014)

$$Sh_i = \sum_{u \subseteq \{1,2,\ldots,d\}\setminus\{j\}} \frac{|u|!(d-|u|-1)!}{d!}\left(S_{u \cup \{i\}} - S_u\right). \tag{8}$$

The Shapley effects are based on game theory (Shapley, 1953) and satisfy the properties that $Sh_i \geq 0$ for all i and $\sum_{i=1}^d Sh_i = \mathbb{V}[Y]$ for any dependence structure of $\mathbf{X}$. Moreover, if $Sh_i = 0$, then the variable X_i is not affecting the model behaviour neither in terms of functional relationship nor in terms of dependence with other variables. When variables are independent, the Shapley effects can be bounded between the first-order and the total Sobol' indices

$$S_i \leq Sh_i \leq S_{T_i} \tag{9}$$

for all $i = 1, 2, \ldots, d$. Hence, in this case it is more convenient to estimate the total indices only. A drawback of Shapley effects is that their computation becomes costly as the dimension d increases. Eq. (8) indeed requires the computation of all 2^d possible Sobol' indices S_u. For $d > 20$, this number exceeds 1 million.

An alternative approach relies on using global sensitivity measures that draw from the literature on measures of statistical association. The definitions and properties of these global sensitivity measures are not affected by the presence of input dependence. The increase in size and dimensionality of modern datasets has renewed interest in the definition of measures of statistical association that can be conveniently computed in the case of large dimensional data. A recent example is the new correlation coefficient by Chatterjee (2021). Chatterjee correlation coefficient can be defined as a global sensitivity measure in Eq. (5)

$$\xi_i^{CvM} = \mathbb{E}\left[\int (P_Y(y) - P_{Y|X_i}(y))^2 dP_Y(y)\right], \tag{10}$$

where the separation measurement is the Cramer-von-Mises distance between cumulative distribution functions. One early example of this class of importance measures is the well-known mutual information, which can be written as

$$\xi_i^{\mathrm{MI}} = \mathbb{E}[\int f_Y(y) ln \frac{f_Y(y)}{f_{Y|X_i}(y)} dy] \tag{11}$$

where $\zeta^{\mathrm{KL}}(X_i) = \int f_Y(y) ln \frac{f_Y(y)}{f_{Y|X_i}(y)} dy$ is the well-known Kullback–Leibler divergence (Kullback and Leibler, 1951). It turns out that if one selects the separation measurement appropriately, one obtains a global sensitivity measure that possesses relevant properties towards understanding variable importance. The first is zero-independence: the sensitivity measure is null if and only if Y is independent of X_i. This property is useful to analysts for avoiding false negatives, that is the risk of judging a variable as unimportant when, in fact, it plays a role in the problem. The indices S_i in Eq. (6) do not possess this property. The second property is max-functionality: The global sensitivity is maximal if and only if Y is a noiseless function of X_i, that is ξ_i takes on its maximum value if and only if $Y = g(X_i)$. Note that, in this case, fixing X_i eliminates uncertainty in Y completely. The third property is monotonicity. Suppose that instead of being informed about X_i we are informed about a variable Z_i which is related to X_i, but through a non-monotonic transformation, (i.e., $Z_i = h(X_i)$, where X_i is a non-monotonic function. Monotonicity guarantees that

the importance of Z_i is lower than the importance of X_i. Global sensitivity measures in (5) can be estimated from a given dataset, with a minimization of the computational cost which is then independent of the number of variables d (Borgonovo et al., 2016; Chatterjee, 2021).

Let us now come to the second goal, trend identification. The analyst aims to understand the behaviour of Y as a function of one or more input variables. One straightforward approach is to examine the signs of the Newton Ratios in Eq. (4) across multiple locations within the input domain. Additionally, an alternative approach to uncovering such trends involves considering the conditional regression functions, denoted as $r_i(x_i) = \mathbb{E}[Y|X_i = x_i]$, when the variables are assumed to be independent. In particular, insights on the monotonicity of the risk profile Y with respect to X_i can be assessed from this function when inputs are independent. An input variable X_i with monotonic effect might compromise the risk profile of the company when it gets to its high values, and need to be properly controlled.

The third goal, interaction quantification, can also be addressed on a global scale. There are different indices for interactions. It might be useful to first estimate the overall impact of interactions due to a single variable. Under variable independence, one can define the total interaction index of variable X_i is (Borgonovo and Rabitti, 2023)

$$S_{I_i} = S_{T_i} - S_i \tag{12}$$

which is the difference between the total effect of X_i and its individual effect. If the analyst wants to investigate the relevance of specific interactions (for instance between variables X_i and X_j), they should consider the difference $S_{i,j} - S_i - S_j$ where the second-order Sobol' index is explicitly given by

$$S_{i,j} = \frac{\mathbb{V}_{X_{i,j}}\left(\mathbb{E}_{\mathbf{X}_{-\{i,j\}}}(Y|X_{i,j})\right)}{\mathbb{V}(Y)}. \tag{13}$$

The sensitivity settings proposed and discussed above allow the insurance company to effectively, transparently and responsibly analyse its risk profile when this is evaluated through a complex internal model. The model design and building phases could benefit too using these settings.

The adoption of these sensitivity settings, each with a clear focus on the model behaviour, can provide complementary insights. For example, from the computational point of view (Borgonovo and Rabitti, 2023) show that an estimate of S_{T_i} is

$$\hat{S}_{T_i} = \frac{\frac{1}{N}\sum_{k=0}^{N-1}\left(\phi_i^{\mathbf{x}^k \to \mathbf{x}^{k+1}}\right)^2}{2\hat{\sigma}^2} \tag{14}$$

where $\mathbf{x}^l$ is the lth scenario, for $l = 0, 1, \ldots, N$ and $\hat{\sigma}^2$ is the estimate of the variance of Y. Eq. (14) shows that the total-order index of variable i can be estimated as a squared mean of local effects $\phi_i^{\mathbf{x}^k \to \mathbf{x}^{k+1}}$ computed considering multiple scenarios $\{\mathbf{x}^0, \mathbf{x}^1, \ldots, \mathbf{x}^N\}$. Thus, having at hand not only the local sensitivity indices for two scenarios but for many more, the analyst can

1. obtain many local effects and gain insights on model output drivers for several scenarios;
2. estimate the total-order indices.

The intuition is that, if the terms $\phi_i^{\mathbf{x}^k \to \mathbf{x}^{k+1}}$ are always small across multiple scenarios, the numerator of Eq. (14) is negligible and the ith variable can be considered unimportant from the local and the global point of view (Borgonovo and Rabitti, 2023). This information strongly supports the understating of the model behaviour globally as well as across deterministic scenarios.

In common actuarial practice, internal models might contain large amounts of variables. For example, Dacorogna (2017) writes that the SCOR's model is based on few thousand risks. A possible strategy to screen large internal models is using Eq. (14) as done in Borgonovo and Rabitti (2023). When the dimension d is very large, investigating the local importance with the one-at-a-time method across multiple scenarios can become computationally expensive for estimating Eq. (14), especially when the internal model is slow. In such cases, it might be advantageous to allocate the entire computational budget to a Monte Carlo sample analysis, which not only enables uncertainty quantification but also allows for the utilization of statistical association measures in Eq. (5) computed from given data. Additionally, the use of a Monte Carlo simulation allows the analyst, on the one hand, to inspect the model input space more thoroughly and, on the other hand, to fit a faster metamodel (i.e. a machine learning model) which can significantly reduce computational time (Krah et al., 2020b,a).

In general, we recommend employing an ensemble of sensitivity measures. Table 1 provides a summary of the characteristics of various sensitivity indices and the information they provide.

Table 1 presents a comparison of sensitivity indices based on type of insights, presence of discrete variables, impact of interactions, variables uncertainty and variables dependence. For instance, if the analyst is interested in the interaction between the i−th and the j−th variables, the partial derivative $\partial^2 f/\partial x_i\partial x_j$ can be adopted but the analyst needs to remember that it depends on the specific location where it is evaluated and it does not take into account the uncertainty of X_i and X_j nor their possible statistical dependence. Moreover, this partial derivative cannot be defined when variables are discrete. In this case, finite-change sensitivity indices can be used. The global impact of interactions can be measured using Sobol' indices, but their interpretation is not clear when variables are dependent. Research on global measures for quantifying interactions under dependent variables is still active. When variables are independent, the analyst using a combination of total Sobol' indices and conditional regression functions can understand the most important variables and their effect on the model output in terms of increase or decrease.

The integration of these sensitivity settings can provide insightful information to the model analysts in insurance companies and regulators. For instance, Floryszczak et al. (2016) develop and study an internal model under the Solvency II framework to study

Table 1

Comparison of sensitivity indices in terms of delivered insights; their well-definiteness in presence of discrete variables; their extension to interaction quantification; whether they account for variables uncertainty; their well-definiteness in presence of variables dependence.

	Analysis	Discrete variables	Impact of interactions	Variables uncertainty	Variables dependence
Partial derivatives	Scenario-specific	No	Yes	No	No
Finite-change indices	Scenario-specific	Yes	Yes	No	No
Conditional regression functions	Global	Yes	No	Yes	No
Sobol' indices	Global	Yes	Yes	Yes	No
Shapley effects	Global	Yes	No	Yes	Yes
Distribution-based	Global	Yes	No	Yes	Yes

the company's net asset value (NAV) and the ratio NAV/SCR. However, Floryszczak et al. (2016) adopt only the trend identification setting for the sensitivity of the company's indicators NAV and NAV/SCR with respect to financial and actuarial variables (namely the interest rate risk, equity risk, mortality risk, surrender risk and specific insurance contract parameters). While this setting allows these authors to identify and visualize the marginal impact of these variables on the company's indicators, it provides no insights on their importance ranking nor on their interaction effects.

4. Numerical illustration

To illustrate the proposed approaches, we have developed a simplified partial internal model in order to quantify the solvency capital requirement (SCR) for premium risk in a framework consistent with Solvency II regulation. To this end we focus on a multi-line non-life insurance company, where the total aggregate claim amount A of the portfolio is described by a frequency-severity approach (Daykin et al., 1994) as follows:

$$A = \sum_{i=1}^{l} A_i = \sum_{i=1}^{l} \sum_{j=1}^{N_i} Z_{j,i} \tag{15}$$

with $A_i = 0$ if $N_i = 0$ and where:

- l is the number of lines of business (LoBs), here represented by Motor Third-Party Liability (MTPL), General Third-Party Liability (GTPL), Property and Motor Own Damages (Mod).
- the aggregate claim amount of each LoB A_i is modelled by a frequency-severity approach
- the number of claims of each LoB N_i is described by the Poisson law $N_i \sim Poi(n_i \cdot Q_i)$, with an expected claim count n_i subject to a parameter uncertainty modelled by a structural variable Q_i.
- Q_i serves as a mixing variable (or contagion parameter) and represents the parameter uncertainty associated with the number of claims. It is assumed that $E(Q_i) = 1$ and that the random variable is defined only for positive values. Following a common approach, we assume that Q_i follows a Gamma distribution, resulting in N_i being distributed as a Negative Binomial.
- the severity is described by random variables $Z_{i,j}$ using LogNormal distributions with a mean m_i and coefficient of variation c_{Z_i}.

We consider classical assumptions of independence between N_i and $Z_{i,h}$ and that the random variables $Z_{i,h}$ are independent and identically distributed. For the sake of simplicity, the dependence between LoBs is here modelled via a Gaussian copula, whose parameters are calibrated using the correlation coefficients provided by Solvency II standard formula (European Commission, 2015a).

The model is based on the parameters described in Table 2 in the Appendix and the main parameters have been calibrated considering a representative insurer in the Italian market. The SCR has been quantified using a Value at Risk at 99.5% confidence level. We are aware that the model can be extended in order to consider other relevant aspects as the inclusion of other sources of risks, modelling tail dependence, the effect of reinsurance treaties, the separate modelling between attritional and large claims, the volatility of expenses, the use of alternative risk measures, but the aim is to provide a realistic toy example that allows to appreciate the alternative approaches provided in the previous section.

For any input parameter combination **x**, we conduct one million Monte Carlo replicates, and from the resulting distribution, we calculate the aggregate claim amount A and the SCR ratio, i.e., the ratio between SCR and the volume of gross premiums. We then validate the model using the available data obtaining a good convergence between the exact and the simulated moments. We estimate a resulting SCR equal to roughly 29.6% of premiums.

As a second step, to introduce uncertainty, we have perturbed the value of the parameters ±10% with respect to their nominal values in Table 2 in the Appendix. We sampled 1,000 observations of **X** using Monte Carlo and estimated the SCR ratio using the original internal model, which took 1.54 days. Subsequently, using this sample we constructed a neural network metamodel with two layers (with 10 and 1 neurons respectively). We compared the linear association between the original simulated value and its predicted counterpart using the same input parameters. The R^2 accuracy scores were 0.98 on the training set (70% of the observations), 0.96 on the validation set (15% of the observations), and 0.96 on the test set (15% of the observations). We then extended our sampling to 10,000 values of **X** and employed the metamodel to compute the SCR ratio. Using the metamodel, this process took only 1.65 s, demonstrating a significant improvement in computational speed.
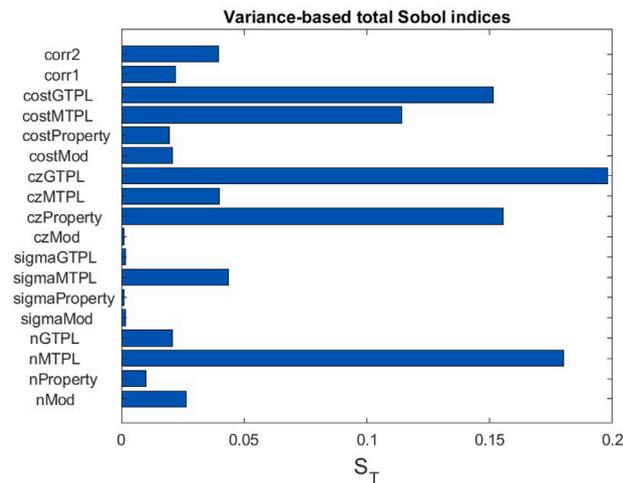
Fig. 1. Global sensitivity indices.

For the sake of space, we restrict the sensitivity analysis to Sobol' total effects, finite changes, and conditional regression functions.

Let us delve into the total Sobol' indices, as illustrated in Fig. 1. It is intriguing to observe how the distinct pooling effects within the Lines of Business (LoBs) influence the ranking of the most significant variables. Notably, the expected claim count of MTPL emerges as a pivotal parameter. This is primarily due to MTPL being the largest segment in the firm's portfolio, where the parameters of the frequency distributions exert a predominant influence on severity. In contrast, GTPL and Property exhibit a notably low number of claims and a highly volatile severity distribution. Given the limited diversification in the portfolio, the coefficient of variation c_{Z_i} assumes a paramount role for these LoBs. According to Mod, the LoB displays a more balanced significance between frequency and severity.

In Fig. 2, we examine the finite-change indices for the difference in SCR ratio between the base-case scenario (in Table 2) and the stressed case, where all variables have been increased by 10%. Fig. 2 substantiates the significance of certain variables, namely the coefficient of variations of the most volatile LoBs (such as GTPL and Property). As anticipated, these characteristics of the severity distribution exert an increasing influence on the SCR ratio. Conversely, it is noteworthy that, for MTPL, both a higher expected number of claims and a greater average cost lead to a reduction in the SCR ratio. This can be justified by the fact that the SCR tends to increase at a rate that is less than proportional to the premiums. Notably, this effect is attenuated when considering the total-finite change due to the compensatory impact of other variables that operate in an opposing manner. We highlight the distinction between the insights obtained from Figs. 1 and 2. In Fig. 1, a high bar implies a significant contribution of a variable to the variability of the SCR ratio. Conversely, a very high (or very low) bar in Fig. 2 indicates that a variable has a substantial impact on increasing (or decreasing) the SCR ratio when transitioning from the base-case scenario to the stressed scenario. This latter analysis is local as it assesses variable importance between two scenarios only.

To investigate the trend of the impact of the four globally most important variables, we consider the conditional regression functions in Fig. 3. As depicted in Fig. 3, these variables exhibit clear monotonic effects, aligning with actuarial expectations. Specifically, a higher volatility of the severity (c_{Z_i}) leads to an increase in the Solvency Capital Requirement (SCR), and in this instance, a rise in the SCR ratio, given that the loading percentages applied by the firm remain unchanged. Conversely, an upswing in the expected number of claims (n_i) has the opposite effect on the SCR ratio. It is a well-established fact that for premium risk, the SCR tends to increase at a slower rate than n_i, resulting in a reduction of the SCR ratio.

In terms of average cost, this parameter typically holds significance for LoB (such as GTPL in our example) characterized by a notable distribution volatility. In this context, although the increased average cost is offset by higher premiums, we observe a more pronounced growth in the standard deviation, consequently leading to an increase in the SCR ratio.

5. Discussion: Benefits and challenges for insurance companies and regulators

We suggest the use of an ensemble of sensitivity measures to conduct a proper sensitivity analysis, incorporating a complex and unified approach that integrates different settings. This integrated sensitivity settings approach combines various techniques, such as one-factor-at-a-time analysis, scenario testing, and global analysis, to provide a comprehensive assessment of the internal model's performance. By adopting this approach, we can thoroughly evaluate the model, assess the firm's risk profile, and identify potential areas for improvement. These insights are invaluable for supporting risk management decisions and ensuring a nuanced understanding of the company's risk profile across different scenarios.

In addition, considering the sensitivity approaches together enhances the understanding and communication of solvency results to model auditors and regulators. Insurance companies gain a greater awareness of how risk assumptions and dependencies affect
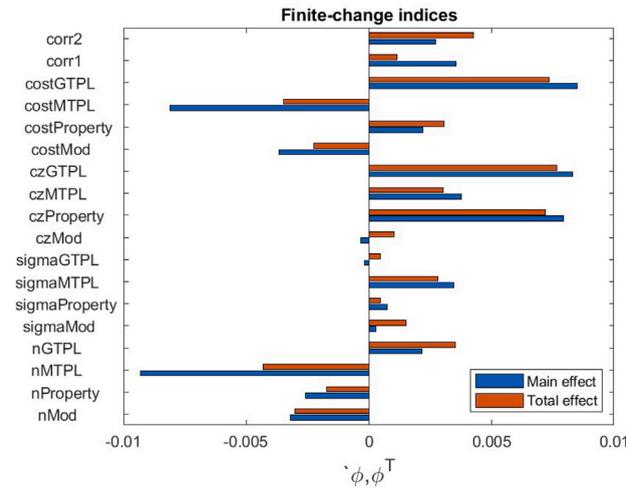
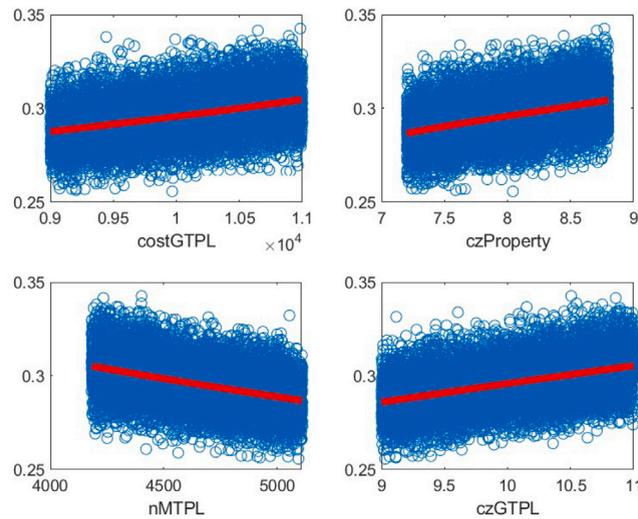**Fig. 2.** Tornado diagram with finite-change sensitivity indices.

Fig. 3. Scatterplots (in blue) with conditional regression functions (in red) for the four most important variables.

model response, while experts can make more informed decisions based on internal model predictions. Moreover, these sensitivity settings can serve as a benchmark for sensitivity analysis approaches in analysing and validating internal models. It is recommended that companies utilize comparable statistical indices to perform sensitivity and uncertainty analysis of their models. Establishing a common minimum requirement based on these settings would enable regulators to consistently and objectively compare sensitivity analyses across different companies' internal models. This approach would also promote a statistical culture within companies, fostering improved model construction, validation, and uncertainty analysis.

Integrating alternative sensitivity approaches in the validation of internal models is of significant importance in the realm of risk management and decision-making. Traditional validation methods, which often rely on single-point estimates and assumptions, might not fully capture the complexity of risks associated with complex systems. By incorporating alternative sensitivity approaches, organizations can enhance their understanding of potential vulnerabilities, identify hidden risks, and make more informed decisions. These approaches allow for assessing the model's robustness and resilience, gaining insights into inputs, assumptions, and key risk drivers, and detecting potential limitations and biases.

Moreover, considering alternative sensitivity approaches promotes a culture of continuous improvement and learning within organizations. By acknowledging the uncertainty and complexity of internal models, stakeholders are encouraged to challenge assumptions, explore alternative perspectives, and refine their understanding of risks. This iterative process fosters innovation and the development of more accurate and robust models over time.

Ultimately, the integration of alternative sensitivity analyses enables organizations to better understand and manage risks, enhances the resilience of decision-making processes, and contributes to the overall stability and sustainability of the organization.

By harnessing the power of these approaches, organizations can proactively adapt to changing circumstances, make more informed decisions, and navigate complex environments with greater confidence.

To conclude, the proper integration of sensitivity analysis techniques and the adoption of alternative sensitivity approaches play a crucial role in validating internal models, enhancing risk management practices, and improving decision-making processes. By considering multiple sensitivity settings and utilizing statistical measures, companies can gain valuable insights into their models' performance, assess their risk profiles, and identify opportunities for improvement. Regulatory bodies should establish common minimum requirements based on these sensitivity settings to facilitate consistent and objective comparisons among different companies' internal models. Through the integration of alternative sensitivity approaches, organizations can foster a culture of continuous learning, refine their understanding of risks, and enhance their resilience in a dynamic and uncertain business environment.

CRediT authorship contribution statement

Emanuele Borgonovo: Conceptualization, Methodology, Supervision, Writing – review & editing, Writing – original draft. **Gian Paolo Clemente:** Conceptualization, Data curation, Methodology, Supervision, Writing – original draft, Writing – review & editing, Software. **Giovanni Rabitti:** Conceptualization, Data curation, Formal analysis, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Appendix. Table with parameters in the internal model

See Table 2.

Table 2
We display the parameters used for modelling the aggregate claim amount of each LoB, the correlation matrix used for calibrating the Gaussian Copula and the parameters that have been used only for the assessment of gross premiums. The central portion of the table reports the base-case values of the correlation parameters. We denote with $corr_1$ the values of the correlations equal to 0.25 and with $corr_2$ the values of the correlations equal to 0.50. In Figs. 1, 2, and 3 the inputs m_i are denoted by $cost$. Parameters λ_i, $i = 1, 2, 3, 4$, represent the safety loadings, expressed as percentages of fair premiums. Parameters η_i, $i = 1, 2, 3, 4$, represent the expenses loadings, expressed as percentages of gross premiums.

Parameters used for A_i computation				
Parameters	Mod	Property	MTPL	GTPL
n_i	1,806	1,167	4,646	306
σ_{Q_i}	0.06	0.09	0.08	0.18
m_i	2,500	6,000	4,800	10,000
c_{Z_i}	2	8	4	10
Correlation matrix				
	Mod	Property	MTPL	GTPL
Mod	1	$corr_1 = 0.25$	$corr_2 = 0.5$	$corr_1 = 0.25$
Property	$corr_1 = 0.25$	1	$corr_1 = 0.25$	$corr_1 = 0.25$
MTPL	$corr_2 = 0.5$	$corr_1 = 0.25$	1	$corr_2 = 0.5$
GTPL	$corr_1 = 0.25$	$corr_1 = 0.25$	$corr_2 = 0.5$	1
Parameters used for gross premiums computation				
λ_i	13%	1%	1%	15%
η_i	32%	33%	22%	33%

References

Booth, Philip, Chadburn, Robert, Haberman, Steven, James, Dewi, Khorasanee, Zaki, Plumb, Robert H., Rickayzen, Ben, 2005. Modern Actuarial Theory and Practice, second ed. CRC Press, Boca Raton.

Borgonovo, Emanuele, 2010. Sensitivity analysis with finite changes: An application to modified EOQ models. European J. Oper. Res. 200 (1), 127–138.

Borgonovo, Emanuele, 2023. Sensitivity analysis. Tutor. Oper. Res. 5, 2–81.

Borgonovo, Emanuele, Hazen, Gordon B., Plischke, Elmar, 2016. A common rationale for global sensitivity measures and their estimation. Risk Anal. 36 (10), 1871–1895.

Borgonovo, Emanuele, Plischke, Elmar, 2016. Sensitivity analysis: A review of recent advances. European J. Oper. Res. 248 (3), 869–887.

Borgonovo, Emanuele, Rabitti, Giovanni, 2023. Screening: From tornado diagrams to effective dimensions. European J. Oper. Res. 304 (3), 1200–1211.

Chatterjee, Sourav, 2021. A new coefficient of correlation. J. Amer. Statist. Assoc. (116).

CRO Forum, 2013. Stress test and scenario setting.

CRO Forum, 2023. ORSA stress and scenario testing. Best practice for assessing risks.

Dacorogna, Michel, 2017. Approaches and techniques to validate internal model results. SSRN.

Dacorogna, Michel, 2023. How to gain confidence in the results of internal risk models? Approaches and techniques for validation. Risks 11 (5).

Daykin, Chris D., Pentikáinen, Teivo, Pesonen, Martti, 1994. Practical Risk Theory for Actuaries. Chapman & Hall.

EIOPA, 2014. Guidelines on system of governance.

EIOPA, 2016. Guidelines on submission of information to national competent authorities.

EIOPA, 2022. Guidelines on the use of internal models.

European Commission, 2009. Directive 2009/138/EC of the European Parliament and of the council of 25 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II).

European Commission, 2015a. Commission Delegated Regulation (EU) 2015/35 of 10 2014 supplementing Directive 2009/138/EC of the European Parliament and of the Council on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II).

European Commission, 2015b. Commission Implementing Regulation (EU) 2015/460 laying down implementing technical standards with regard to the procedure concerning the approval of an internal model in accordance with Directive 2009/138/EC of the European Parliament and of the Council.

Flaig, Solveig, Junike, Gero, 2022. Scenario generation for market risk models using generative neural networks. Risks 10 (11), 199.

Floryszczak, Anthony, Courtois, Olivier Le, Majri, Mohamed, 2016. Inside the Solvency 2 Black Box: Net Asset Values and Solvency Capital Requirements with a least-squares Monte-Carlo approach. Insurance Math. Econom. 71, 15–26.

International Actuarial Association, 2013. Stress testing and scenario analysis.

Krah, Anne-Sophie, Nikolic, Zoran, Korn, Ralf, 2020a. Least-squares Monte Carlo for proxy modeling in life insurance: Neural networks. Risks 8 (4), 116.

Krah, Anne-Sophie, Nikolic, Zoran, Korn, Ralf, 2020b. Machine learning in least-squares Monte Carlo proxy modeling of life insurance companies. Risks 8 (1), 21.

Kullback, Salomone, Leibler, Richard A., 1951. On information and sufficiency. Ann. Math. Stat. 22 (1), 79–86.

Lloyds, 2022. Internal model validation guidance.

Loisel, Stephan, Nisipasu, Kati, 2016. Ex-ante model validation and back-testing. In: Laurent, J.P., Norberg, R., Planchet, F. (Eds.), Modelling in Life Insurance - A Managemenent Perspective. In: European Actuarial Academy Series, Springer, pp. 151–160.

Owen, Art B., 2014. Sobol' indices and Shapley value. SIAM/ASA J. Uncertain. Quantif. 2 (1), 245–251.

Rabitz, Herschel, Alis, Omer F., 1999. General foundations of high-dimensional model representations. J. Math. Chem. 25, 197–233.

Razavi, Saman, Jakeman, Anthony, Saltelli, Andrea, Prieur, Clementine, Iooss, Bertrand, Borgonovo, Emanuele, Plischke, Elmar, Piano, Samuele Lo, Iwanaga, Takuya, Becker, William, Tarantola, Stefano, Guillaume, Joseph H.A., Jakeman, John, Gupta, Hoshin, Melillo, Nicola, Rabitti, Giovanni, Chabridon, Vincent, Duan, Qingyun, Sun, Xifu, Smith, Stefán, Sheikholeslami, Razi, Hosseini, Nasim, Asadzadeh, Masoud, Puy, Arnald, Kucherenko, Sergei, Maier, Holger R., 2021. The Future of Sensitivity Analysis: An essential discipline for systems modeling and policy support. Environ. Model. Softw. 137.

Saltelli, Andrea, Aleksankina, Ksenia, Becker, William, Fennell, Pamela, Ferretti, Federico, Holst, Niels, Li, Sushan, Wu, Qiongli, 2019. Why so many published sensitivity analyses are false: A systematic review of sensitivity analysis practices. Environ. Model. Softw. 114, 29–39.

Saltelli, Andrea, Bammer, Gabriele, Bruno, Isabelle, Charters, Erica, Fiore, Monica Di, Didier, Emmanuel, Espeland, Wendy Nelson, Kay, John, Piano, Samuele Lo, Mayo, Deborah, et al., 2020. Five ways to ensure that models serve society: a manifesto. Nature 582, 482–484.

Saltelli, Andrea, Ratto, Marco, Andres, Terry, Campolongo, Francesca, Cariboni, Jessica, Gatelli, Debora, Saisana, Michela, Tarantola, Stefano, 2008. Global sensitivity analysis. In: The Primer. John Wiley & Sons.

Saltelli, Andrea, Tarantola, Stefano, Campolongo, Francesca, 2000. Sensitivity analysis as an ingredient of modeling. Statist. Sci. 15 (4), 377–395.

SCOR, 2008. From principle-based risk management to solvency requirements. In: Analytical Framework for the Swiss Solvency Test.

Shapley, Lloyd, 1953. A value for n-person games. In: Kuhn, H.W., Tucker, A.W. (Eds.), Ann. Math. Study 28, Contributions to the Theory of Games. pp. 307–317.

Sobol', Ilya M., 1993. Sensitivity estimates for nonlinear mathematical models. Math. Model. Comput. Exp. 4, 407–414.