Improvement of the Reliability and Robustness of Variance-Based Sensitivity Analysis of Final Repository Models by Application of Output Transformation

Dirk-Alexander Becker,
Gesellschaft fuer Anlagen- und Reaktorsicherheit (GRS) mbH, Braunschweig, Germany
dirk-alexander.becker@grs.de

Abstract: Long-time performance assessment models for final repositories for radioactive waste typically produce heavily tailed output distributions that extend over several orders of magnitude and under specific circumstances can even include a significant number of exact zeros. A variance-based sensitivity analysis gives a strong overweight to the typically very few values that are far away from the expected value of the distribution, which can lead to a low robustness of the evaluation. Moreover, while a variation of the model output, even over orders of magnitude, is of little interest if it happens on a radiologically irrelevant level, a mere factor of 2 near the permissible dose limits can be very important. Both types of problems can be mitigated by applying appropriate output transformations before performing the sensitivity analysis. The effects of different transformations on the sensitivity analysis results for typical final repository model systems are demonstrated.

Keywords: Variance-based sensitivity analysis, transformation, radioactive waste disposal

1. INTRODUCTION

The long-term performance of final repositories for radioactive waste has to be investigated using computational models describing the release of radionuclides from waste containers and their transport through the near field, the geosphere and the biosphere. Probabilistic sensitivity analysis is an important tool for improving the understanding of the model behavior as well as for identifying research needs. With realistic parameter ranges, however, such models typically produce heavily tailed output distributions that extend over several orders of magnitude and, under specific circumstances, can even include a significant number of exact zeros. This causes two types of problems:

a) A variance-based evaluation gives a strong overweight to values that are far away from the expected value of the distribution. Therefore, the total output variance is typically dominated by very few individual values, which can lead to low robustness of the evaluation.

b) In view of radiological safety, the highest output values are the most relevant ones. While a variation of the model output, even over orders of magnitude, is of little interest if it happens on a radiologically irrelevant level, a mere factor of 2 near the permissible dose limits can be very important. Usual sensitivity analysis methods do not by themselves take account of such non-mathematical asymmetries.

Both types of problems can be mitigated by applying appropriate output transformations before performing the sensitivity analysis. The goal pursued with such transformations is to map the model output, typically the effective annual dose to a human individual, to some magnitude that better represents the actual harmful effects to the environment on a scale that is well consistent with the mathematical method of evaluation.

In cases of model output values distributed over several orders of magnitude, it is sometimes recommended to analyze the logarithms of the output instead of the calculated values themselves. This, however, does not always solve the problem and can even make it worse. Zero values are mathematically excluded from a log-transformation, and very low values are highly overvalued. A change of, say, three orders of magnitude would be given the same weight, regardless of whether it occurs at a near-zero level, maybe due to some more or less arbitrary numerical specifics of the model, or at a numerically significant and radiologically relevant level. The results of a sensitivity analysis on a log-transformed model output can therefore become more or less useless.

To avoid this problem, we have to look for transformations that handle low values differently from high values. Of course, the essential precondition for such an approach is a sensible criterion that discriminates “low” from “high” values. One possibility is that we take a strictly mathematical point of view and define anything below the mean or the median of all values as low and anything above as high. Another, more physically motivated approach is to use a threshold value that is based on radiological considerations and does not take account of the actual distribution of the model output.

It seems sensible to select a continuous transformation that maps

- zero to zero,
- very low values to values near zero,
- the threshold value to one
- and high values to moderately increasing values above one, without giving them an undue overweight.

In this paper three such transformations are proposed and their effects on the results of a variance-based sensitivity analysis are demonstrated using two different final repository models.

2. THE MODELS

This chapter gives a short overview of the basic characteristics of the two models investigated here. Both models represent hypothetic final repositories in rock salt formations in Germany.

2.1. HLW Repository

The first model is based on a former planning for a possible final repository for high-level radioactive waste in Germany, which was foreseen to be installed in excavations specifically mined for this purpose in a salt formation with a high creeping capability. The concept envisaged boreholes for emplacement of canisters with vitrified waste from reprocessing as well as drifts for emplacement of containers with spent fuel elements. Each borehole or drift is sealed by a plug, and the same applies to the loading drifts. Since under the high temperatures in the vicinity of heat-producing waste the salt creeps rather fast, the wastes will, in most cases, be tightly included in the salt within several decades. Then each contact of brine with parts of the wastes is excluded for the future and no contaminant release is possible. This situation leads to a zero-output of the model. It is, however, possible that a brine inclusion in the salt formation opens during the early phase. Then the fluid can reach the waste containers and dissolve radionuclides. The creep-induced convergence of the remaining voids in the mine has then a disadvantageous effect, since it presses the contaminated brine through the seals and into the geosphere.

The calculation results analyzed here were produced in 2003 using the code package EMOS [1]. The model consists of parts for the near field, the geosphere and the biosphere, and it finally yields the hypothetical time-dependent dose rates to a human individual. 31 parameters were varied statistically according to their specific distributions. Due to the high probability of tight inclusion of all wastes, only some 15 % of the model runs yield a non-zero output.

The HLW repository model was calculated 3000 times with a random sample. Only 491 of the runs yielded a non-zero output, but some of them reach maxima above 10^{-4} Sv per year. The time curves of the six runs with the highest maxima are shown in Figure 1 (left).

2.2 LILW Repository

The second model investigated here represents a hypothetic final repository for low- and intermediate-level radioactive waste that is assumed to be installed in an abandoned salt production mine. Its main features are loosely based on a real site of this type in Germany, the Morsleben repository. The salt formation is inhomogeneous and has a low creeping capability. It is assumed that the mine openings are filled with brine from the overburden after some time. At that point in time, some short-lived wastes, which are disposed of in one of the openings, start to release contaminants. These are dissolved in the brine and pressed out to the geosphere by the convergence process. In order to protect the longer-lived and more radiotoxic wastes from the brine, the main waste emplacement area is isolated from the rest of the mine by a specific seal, which, however, can be chemically corroded by magnesium. Depending on the magnesium content of the brine and the initial permeability of the seal material, the seal can nearly suddenly fail at some point in time. This leads to a short-lasting decrease, followed by a fast, significant increase of the contaminant release. The decrease is due to the fact that after seal failure it takes a little time to fill up the emplacement area, during which the brine outflow from the mine is reduced.

Like the HLW model, the LILW model consists of three parts, describing the near field, the geosphere and the biosphere. The calculations were done with the software package RepoTREND [2], which contains modules for each of the three parts and is specifically designed for calculating the transport of radionuclides through and release from a repository system. The model output is the annual dose to a human individual. In the investigations presented here, 11 uncertain parameters, all pertaining to the near field, were varied according to appropriate distributions. In contrast to the HLW case the LILW model does not produce zero output runs.

Three different random samples, each containing 3000 parameter sets, were drawn. The right side of Figure 1 displays some typical time curves calculated by the LILW model, which show the seal failure and the subsequent increase of the dose rate. The highest maxima reach about 10^{-5} Sv/yr.

Figure 1: Typical Time Development Curves of Both Models (Left: HLW, Right: LILW)

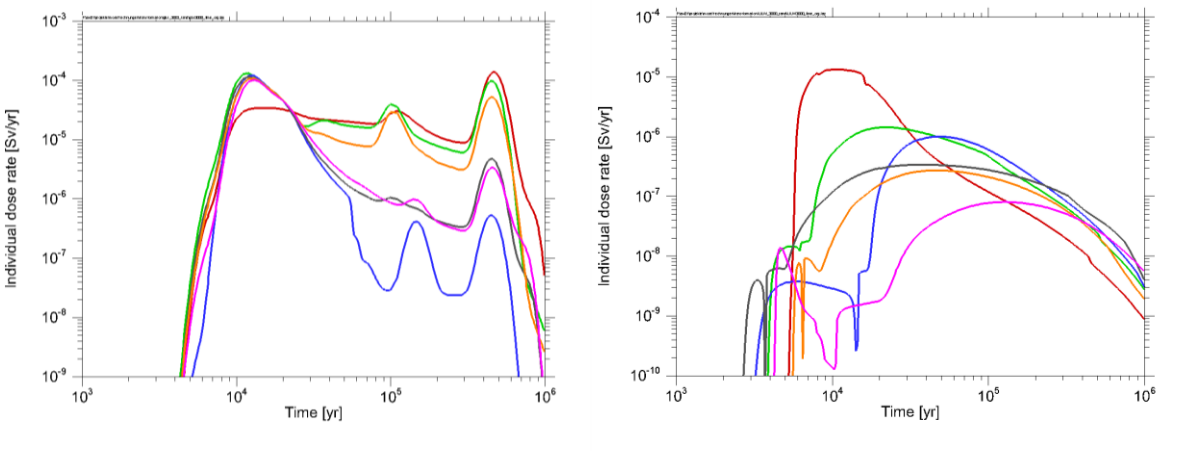

3. THE CONSIDERED OUTPUT TRANSFORMATIONS AND THEIR EFFECTS

In chapter 1, four requirements were formulated that should be fulfilled by an output transformation. According to these requirements, we chose the following three transformations for investigation in this paper (the model output value is denoted by *y*):

$$\text{transformation 1:} \quad y \mapsto \log_2(1 + y/a)\ ,$$
$$\text{transformation 2:} \quad y \mapsto (y/a)^{0.2}\,,$$
$$\text{transformation 3:} \quad y \mapsto (y/a)^{0.3}\,.$$

Figure 2 shows how these transformations act on the model output on a linear and on a logarithmic scale. All three transformations map low values to values near 0, the threshold value *a* to 1 and high values to values that seem “not too far above” 1. The threshold parameter *a* should be selected to discriminate “low” from “high” values in a sensible way. For the investigations presented here it seems reasonable to fix it about three orders of magnitude below the highest occurring dose rates. Therefore, we chose $a = 10^{-7}$ Sv/yr for the HLW model and $a = 10^{-8}$ Sv/yr for the LILW model. Additionally, the effects of transformation 1 are demonstrated for a threshold value *a* that is adaptively calculated as the median of the analyzed data for each point in time. Since the transformation is monotonic, it maps the median of the original data to the median of the transformed data, which is automatically equal to 1.

Figure 2: Effects of the Transformations ($a = 10^{-7}$)

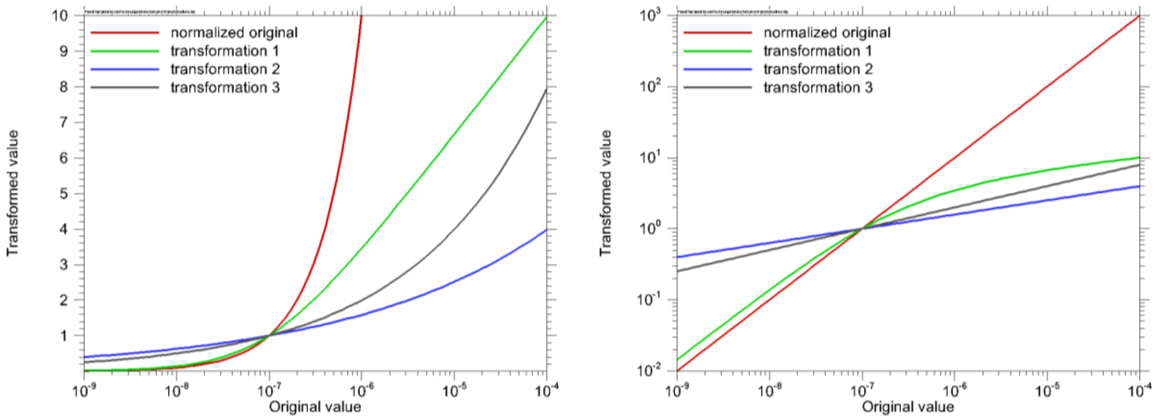

Figure 3 demonstrates how the three transformations affect the output of the HLW model in view of the variance. The histograms show the distribution of the peak values of all runs as red columns and the respective contributions to the total variance in blue. Due to the high number of zero runs the lowest bin is by far most populated. It can be seen that for the original data the figure looks rather unbalanced. While there are 2689 (89%) zeros or very low values with a common contribution to the total variance of about 5.5%, the two (0.067%) highest peak output values contribute 7.4% to the variance. The highest 25 values, that is 0.83%, are responsible for 50% of the variance. This disproportion is considerably mitigated by any of the three transformations. The figures show that in all cases the relation between the frequencies and the contributions to the variance is much more balanced than for the original data. The percentage of the highest values commonly responsible for 50% of the variance is about 4 to 5%. The highest contributions to the variance do no longer result from very few outliers at the upper end, but from the zeros and from the most populated bins in the region of higher values. It is therefore expected that a variance-based sensitivity analysis gives more robust and reliable results if performed on the transformed data.

Figure 3: Histograms for Frequency and Contribution to the Total Variance for the Original and the Transformed Data (Peak Dose Values)

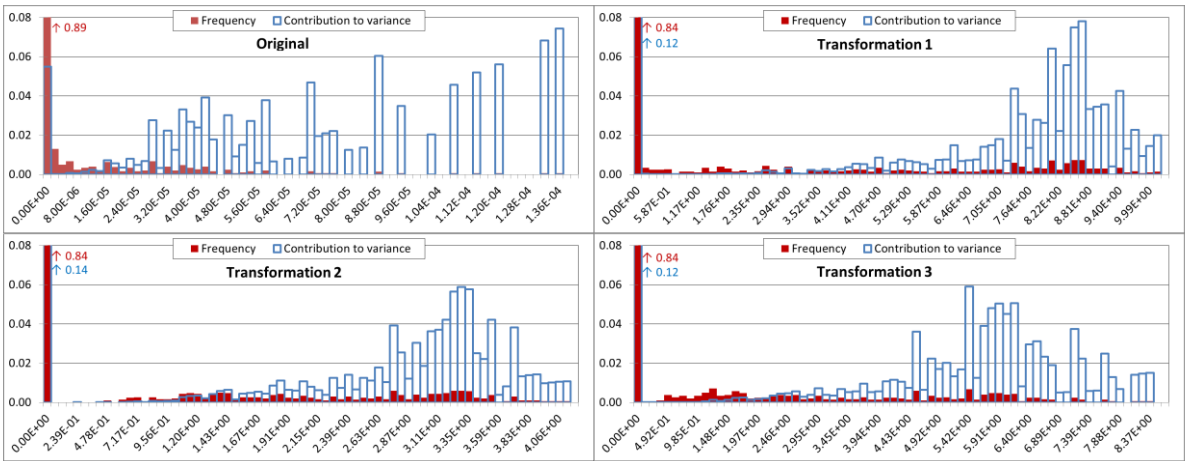

4. RESULTS

In the following some results are presented. With the HLW model 3000 runs were performed; with the LILW model three times 3000 runs were performed with different random samples. The transformations were applied to the time-dependent output of the calculations, and afterwards, a simple variance-based sensitivity analysis was performed. This means that the first-order sensitivity indices

$$\mathrm{SI1}_j = \frac{\mathrm{Var}(\mathrm{E}(Y \mid X_j))}{\mathrm{Var}(Y)}$$

were calculated for all input parameters X_j and 300 points in time. Y means the entirety of the model output values for a specific point in time, $\mathrm{E}(Y|X_j)$ is the expectation of Y under the condition that X_j is hold constant. For the details of the underlying theory see, e.g. [3]. For calculation of the sensitivity indices the EASI method [4] was applied using a MATLAB script by E. Plischke (available under http://ipsc.jrc.ec.europa.eu/?id=756). EASI is a simple effective algorithm for calculating global sensitivity indices of first order using Fast Fourier Transformation. It is very quick, can be applied with any kind of sample and seems to yield results of similar robustness and reliability as other, much more numerically expensive methods [5].

4.1. Results for the HLW Model

In Figure 4 the time-development of the SI1 values for the HLW model, calculated with EASI from the original model output as well as from the transformed data, is presented. The curves of all 31 parameters are plotted, but for clarity reasons, only the six most important ones are given in the legend.

Figure 4: Time Development of SI1 for the HLW Model, Calculated with EASI from the Original and the Transformed Data

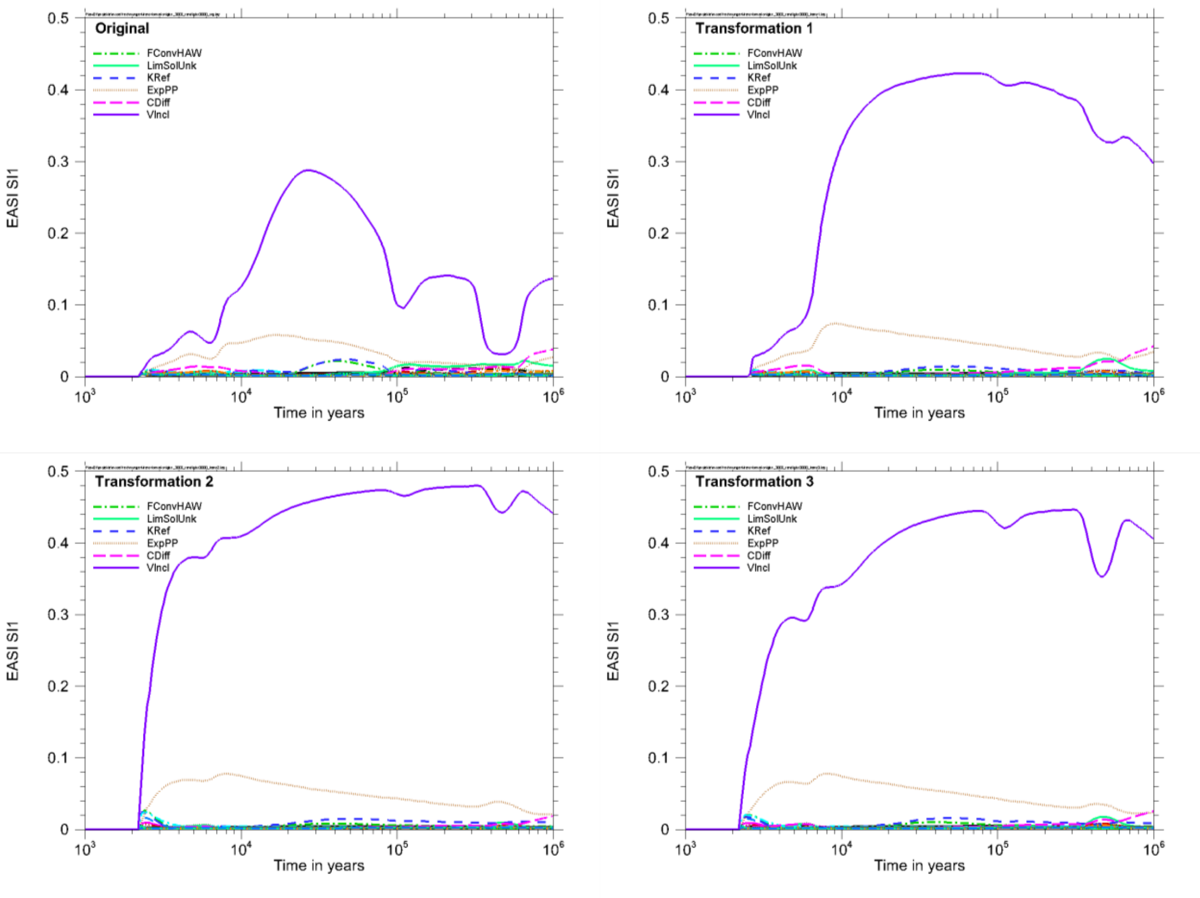

It can be seen from all figures that the model is clearly dominated at all times by the parameter *VIncl*, which is the volume of the brine inclusion in the salt formation. Additionally the parameter *ExpPP*, the exponent in the permeability-porosity relation, plays a certain role. Moreover, the curves of *CDiff* (Diffusion coefficient), *KRef* (reference convergence rate), *FConvHAW* (reduction of convergence in HLW fields) and *LimSolUnk* (solubility limits under unknown chemical conditions) reach, at least at some points in time, values that can be optically distinguished from the zero line.

While all four figures agree on these general facts, there are essential differences in the details. The dominance of *VIncl* is still more pronounced if transformed data are evaluated, no matter which of the transformations is applied. The calculated SI1 values for this parameter are at all times considerably higher for the transformed data than for the original output. The SI1 values of all other parameters but *ExpPP* seem to decrease, except at the very end of the simulation period. Looking at the time phase between about $4 \cdot 10^5$ and $6 \cdot 10^5$ years we see that, by reasons we do not discuss here, the SI1 of *VIncl*, if calculated for the original data, decreases below 0.04, which is only slightly above the values of the other parameters. From this figure alone we would conclude that during this period the parameter is nearly as insignificant for the output as all the other parameters. The figures for the transformed data, however, give a completely different impression. In all of them there is only a slight decrease of the SI1 of *VIncl* during the mentioned period and it remains between 0.32 and 0.44, depending on the applied transformation.

In the evaluation of the original data, the SI1 curve of *ExpPP* nearly vanishes in the mess of insignificant curves at about 10^5 years until the end of the simulation period. For the transformed output, however, it remains clearly silhouetted against the other curves and *ExpPP* is identified as the second most important parameter for nearly the entire simulation period. The small differences seem to be amplified by performing a transformation before the evaluation.

It is noticeable that in the early phase, up to about 8000 years, the results obtained with transformation 1 resemble more those of the original data than those of transformations 2 and 3. This is obviously due to the fact that during early times the model output is generally below the threshold of 10^{-7} Sv/yr, as can be seen in Figure 1 (left), and the logarithmic transformation leaves low values more or less unchanged, except from a factor (see Figure 2). In contrast, the power law transformations 2 and 3 specifically pronounce differences in values below the threshold, since their exponents are smaller than 1.

4.2. Results for the LILW Model

In Figure 5 the time-development of the SI1 values for the LILW model, calculated with EASI from the original model output as well as from the transformed data, is presented. All parameters are shown, distinguished by different colors. The results obtained using three different random samples are marked by different line styles.

It is predominantly conspicuous that the parameter *TBrine*, which does not produce considerably high SI1 values calculated from the original output, becomes much more important if transformed data are evaluated, especially in the early phase. For transformation 2 its SI1 reaches a maximum of 0.74. This parameter represents the point in time when the mine openings are filled with brine, which happens around 10 000 years with a log-normal distribution. Before this point in time there is no contaminant release at all, so that there are lots of zeros at early times. It is clear that *TBrine* is dominant in this phase, because it decides about zero or non-zero results. This dominance is much better reflected by the SI1 curves for the transformed output. Since in many cases, in the time phase between brine intrusion and seal failure the calculated dose rate remains below the threshold value of 10^{-8} Sv/yr, we have the same effect as described for the HLW model: the power law transformations 2 and 3 amplify variations in the range of low values and therefore additionally emphasize the dominance of the parameter *TBrine* in this phase.

Figure 5: Time Development of SI1 for the LILW Model, Calculated with EASI from the Original and the Transformed Data

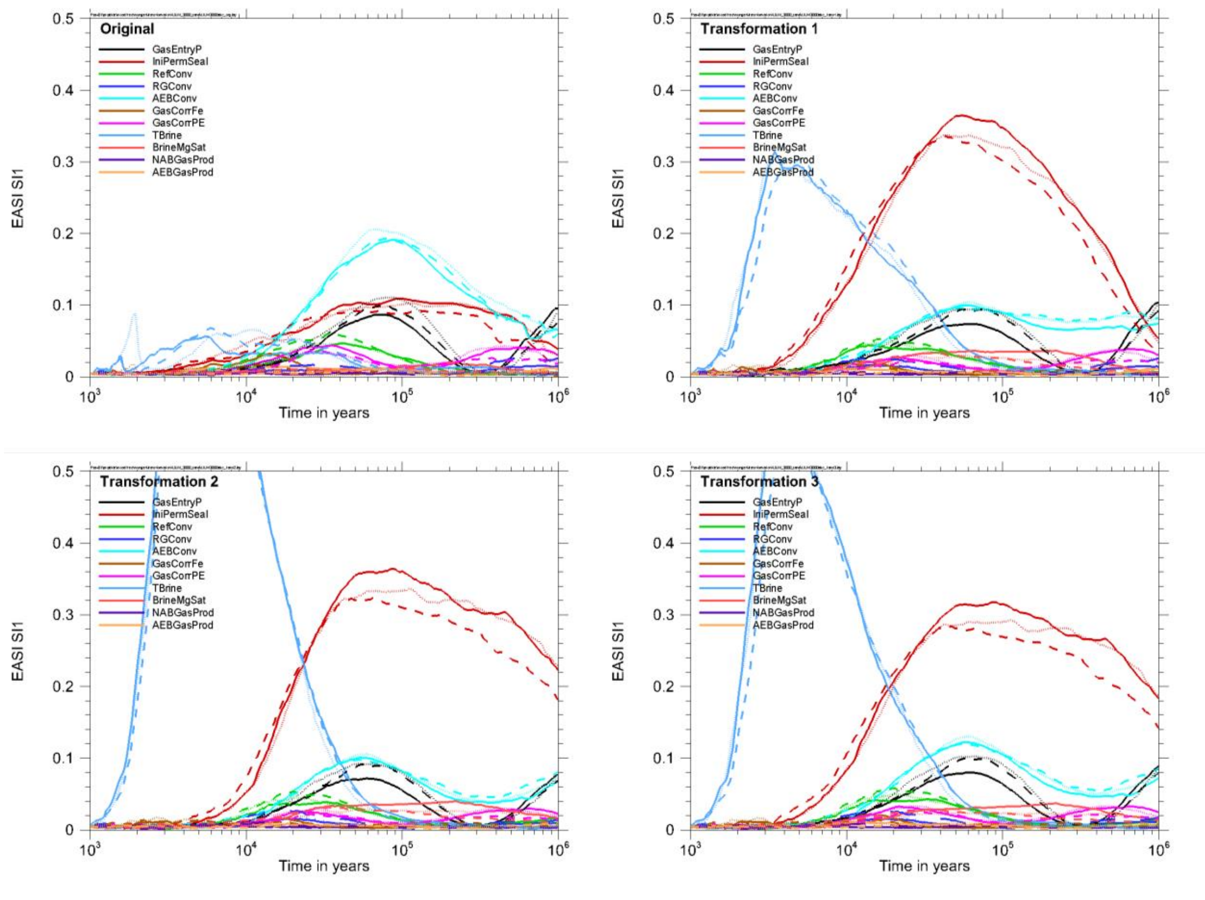

Apart from the eye-catching increase of the importance of *TBrine* due to the transformation, there is another noteworthy effect. While in the original data the parameter *AEBConv*, which represents the convergence rate of the sealed emplacement area, dominates for most of the simulation period, reaching a maximum SI1 of 0.2, this parameter is only of secondary importance in the transformed data with SI1 values no higher than 0.12. Instead, the parameter *IniPermSeal*, representing the initial permeability of the dissolving seal, assumes the dominating role with a maximum SI1 between 0.3 and 0.38, depending on the sample and the transformation. *IniPermSeal* determines the initial flow rate of corroding brine through the intact seal and is predominantly responsible for the time of seal failure. So, for each point in time, this parameter decides about whether the seal has already failed or not, which typically means a difference in dose rate of two or three orders of magnitude. We call this model behavior *quasi-discrete*, because it leads to separate congregations of output values. The parameter *BrineMgSat*, which represents the magnesium saturation of the brine and with it its corrosive potential, also has an influence on the time of seal failure and contributes to triggering this behavior. While, however, its SI1curve is nearly invisible in the evaluation of the original model output, it becomes more relevant if transformed data are evaluated. Obviously, transformations of the considered kind emphasize the sensitivity of the model against parameters that cause a quasi-discrete model behavior.

The SI1 curves of all other parameters progress close to the zero line and show only little differences between the four evaluations.

With the output of the LILW model an additional evaluation was performed using transformation 1 with a threshold value *a* that was determined adaptively. This means that, instead of using a more or less subjectively fixed threshold, for each point in time the value *a* was calculated independently as the median of all 3000 output values. On the one hand, this makes sure that the determination of *a* follows an objective procedure, always keeping one half of the data below and the other half above the threshold. As the transformation is monotonic and maps *a* to 1, the median of the transformed data is constantly equal to 1. On the other hand, a non-constant threshold between values considered as "low" or "high" is a bit hard to understand and must not lead to misinterpretation of results.

Figure 6 shows the time development of SI1, calculated from the adaptively transformed data, on the left side. As long as more than half of the data are zeros, the median is also exactly zero and cannot be used for *a*; therefore the curves start only at about 7000 years. On the right side of Figure 6, the time development of the median, which was used as the transformation parameter *a*, is presented for the three random samples.

Figure 6: Time Development of SI1, Calculated with EASI from the Adaptively Transformed Data, and of the Data Median

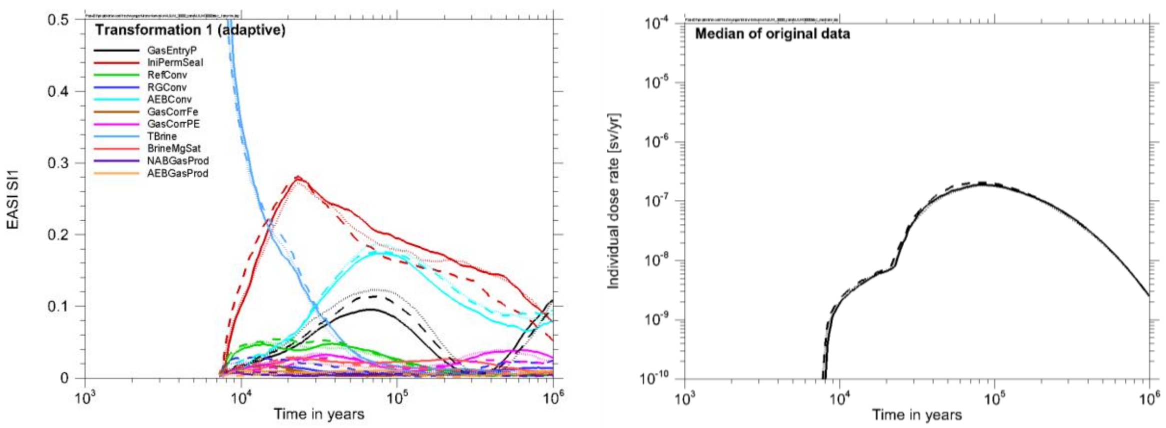

There are very little differences between the medians calculated for the three samples. For most of the time, the median of the model output is in the range between $2 \cdot 10^{-9}$ and $2 \cdot 10^{-7}$ Sv/yr. After about 22000 years it reaches the value of 10^{-8} Sv/yr and afterwards increases faster. At this time, in most of the cases the seal has failed, so that the higher results become dominating. This is also visible in the SI1 curve for *IniPermSeal*, which is the most relevant parameter for the time of seal failure. The curve reaches its maximum exactly at that time where the median curves show the bend. The then increasing threshold value obviously leads to a decrease of the SI1of *IniPermSeal*, but does not seem to have a comparable effect to the other parameters.

With the adaptive transformation it becomes clearly visible that in the very early phase below 10^4 years *TBrine* is by far the dominating parameter. The curve starts at a value of 0.95. The importance of *AEBConv* appears more pronounced with the adaptive transformation than it was the case with any of the fixed transformations. Its SI1 curve reaches a maximum of about 0.2 and resembles more that one calculated from the original data (see Figure 5). For all other parameters there are only little differences to the curves obtained from the original data or the other transformations.

5. SUMMARY AND CONCLUSIONS

We have applied three different transformations to the time-dependent output of two different models for long-term performance assessment of final repositories and calculated the variance-based first-order sensitivity indices using the EASI algorithm. All transformations make use of a threshold value *a*, which is used for normalization of the model output data and discriminates "low" from "high" values. Additionally, a transformation with adaptive determination of *a* was applied to the LILW model (this could not be done for the HLW model, because it produces 85% zero output so that the median is always zero and cannot be used as the threshold value). The transformations are monotonic and map zero to zero, *a* to one and "high" values to values in the range of 1 to about 10. The motivation for this was to project the widely distributed model output to a range that is better adequate for a variance-based evaluation and to reduce the overvaluation of a few high outliers. We investigated a logarithm-based transformation and two power transformations with different exponents smaller than 1.

Transformations of the considered type seem to amplify the differences and accentuate the relevant results of a variance-based sensitivity analysis. Generally, compared to an evaluation of the original model output, the SI1 values of the important parameters seem to increase if calculated from the transformed output, while those of less relevant parameters remain more or less unchanged or even decrease. Thus, by performing an adequate transformation prior to a variance-based sensitivity analysis, one can obtain more unique results.

While the log-based transformation does not significantly change low values and consequently, nearly the same SI1 values as for the non-transformed output are calculated if most of the data are below the threshold, the power-based transformations amplify differences in low values. Parameters that are specifically relevant for low values may therefore be better identified using a power-based transformation.

The threshold parameter a has to be selected adequately for the intended investigation. It can either be chosen from a radiological point of view and set to a value that is considered to be a threshold for radiological relevance, or according to more mathematical aspects. A possible choice of the latter type would be the overall median of all analyzed data. An adaptive transformation that uses the time-dependent median of the model output at each point in time as threshold value can also provide interesting results. In time phases with many exceptionally high or low values this approach inhibits over- or undervaluation of data. This kind of transformation can be of specific interest for the investigation of models that produce very different output distributions and medians at different times. When interpreting the results one should keep in mind, which threshold value has been used and why.

In a different investigation [6] we found that the SI1 calculated from the original model output yield results that qualitatively differ from those obtained with regression-based sensitivity analysis. This is no longer the case if the transformed output is analyzed, no matter which of the transformations is applied. The results of the variance-based sensitivity analysis of transformed data seem to be well in line with those of the regression-based evaluation. From this we conclude that such transformations might improve the reliability of a variance-based evaluation.

With the LILW model three different random samples were investigated. The differences between the SI1 time curves obtained with these three samples do not significantly decrease if a transformation is applied, except from the very early time phase, during which only a few parameter sets yield a non-zero model output. From this observation we conclude that transformations can improve the robustness of the variance-based evaluation if it is based on only a few non-zero data.

Acknowledgements

The work presented in this paper is financed by the German Federal Ministry for Economic Affairs and Energy (BMWi) under sign 02E10941.

References

[1] D. Buhmann, “*Das Programmpaket EMOS – Ein Instrumentarium zur Analyse der Langzeitsicherheit von Endlagern*”, GRS-159, Gesellschaft fuer Anlagen- und Reaktorsicherheit (GRS) mbH, 1999

[2] T. Reiche, D.-A. Becker, D. Buhmann and T. Lauke, “*Anpassung des Programmpakets EMOS an moderne Softwareanforderungen*”, GRS-A-3623, Gesellschaft fuer Anlagen- und Reaktorsicherheit (GRS) mbH, 2011

[3] A. Saltelli, K. Chan and E. M. Scott (ed.), “*Sensitivity Analysis*”, Wiley, 2000

[4] E. Plischke. "*An effective algorithm for computing global sensitivity indices (EASI)*", Reliability Engineering and System Safety 95, pp. 354–360, (2010)

[5] S. Spiessl and D.-A. Becker, "*Sensitivity Analysis of a Final Repository Model with Quasi-Discrete Behaviour Using Quasi-Random Sampling and a Metamodel Approach in Comparison to Other Variance-Based Techniques*", Reliability Engineering and System Safety (submitted for publication)

[6] S. Spiessl and D.-A. Becker, "*Investigation of Different Sampling and Sensitivity Analysis Methods Applied to a Complex Model for a Final Repository for Radioactive Waste*", PSAM12, Honolulu, (2014)