Uncertainty Importance Evaluation in Probabilistic Risk Assessment Using Wasserstein Distance with Bias–Variance Decomposition

Takafumi Narukawa[a], **Louise Guichard**[b], **and Takashi Takata**[a]
[a] The University of Tokyo, Tokyo, Japan, narukawa@n.t.u-tokyo.ac.jp; takata_t@n.t.u-tokyo.ac.jp
[b] IMT Atlantique, Nantes, France, louise.guichard@imt-atlantique.net

Abstract: This study proposes a Wasserstein-based uncertainty importance evaluation method for probabilistic risk assessment (PRA). The method defines the Wasserstein Importance Measure as the primary measure of uncertainty importance, using the second-order Wasserstein distance, an optimal transport distance between probability distributions. To support interpretation, auxiliary measures are introduced by decomposing its squared value into the Wasserstein–Bias component, representing changes in mean location, and the Wasserstein–Variance component, representing residual changes in spread and shape after removing the mean. Numerical experiments using prescribed probability distributions showed that the proposed measure gives importance trends broadly consistent with the existing Borgonovo measure, while more clearly reflecting distributional changes involving large output deviations because it accounts for transport cost in the output-value direction. The auxiliary decomposition further identifies whether the distributional change is dominated by a location shift or by non-bias changes in spread and shape. These results indicate that the proposed method can support more interpretable prioritization of uncertainty reduction efforts in PRA.

1. INTRODUCTION

In probabilistic risk assessment (PRA), epistemic uncertainties [1] in input parameters associated with system components are propagated through the risk model and represented as a probability distribution of the output risk metric. Uncertainty importance measures quantify the influence of input uncertainties on output uncertainty [2]. They provide a rational basis for prioritizing data collection and model improvement efforts aimed at reducing epistemic uncertainty, as well as for supporting safety-related decision making.

Several uncertainty importance measures have been proposed for PRA, including variance-based measures and moment-independent measures [2–5]. However, these measures generally have difficulty distinguishing between changes in the location of the output distribution and changes in its shape caused by input uncertainty. From the viewpoint of uncertainty characterization, these two types of distributional change correspond to bias and variance, respectively, where “variance” is defined here in a broad sense as the non-bias component of the distributional difference, arising mainly from discrepancies in distributional shape. Distinguishing them is important for safety-related decision-making, because they may imply different priorities for uncertainty reduction and risk management.

To address this issue, this study proposes an uncertainty importance evaluation method based on the Wasserstein distance [6], an optimal transport metric between probability distributions. The method uses the second-order Wasserstein distance as the primary importance measure for changes in PRA output distributions and proposes auxiliary measures based on the bias–variance decomposition of the squared Wasserstein distance. These auxiliary measures help interpret whether the distributional change is mainly due to a shift in the mean location or to non-bias distributional changes after removing the mean. Through numerical analyses, we examine the basic behavior of the proposed measures and discuss their usefulness for prioritizing data acquisition and model improvements aimed at reducing epistemic uncertainty in PRA.

2. UNCERTAINTY IMPORTANCE MEASURES BASED ON WASSERSTEIN DISTANCE

2.1. Wasserstein Distance

The Wasserstein distance [6] is a probability metric that quantifies the discrepancy between two probability distributions in terms of the minimum cost required to transport probability mass from one distribution to the other.

Let ($\mathcal{X}$, d) be a Polish metric space, i.e., a complete and separable metric space, and let μ and v be two probability measures on $\mathcal{X}$. Let $\Pi(\mu,v)$ denote the set of all joint probability measures on $\mathcal{X}\times\mathcal{X}$, whose marginal distributions are μ and v, respectively. Then, for $p\in[1,\infty)$, the Wasserstein distance of order p between μ and v is defined as

$$W_p(\mu,v)=\left(\inf_{\pi\in\Pi(\mu,v)}\int_{\mathcal{X}\times\mathcal{X}} d(x,y)\,d\pi(x,y)\right)^{1/p}. \tag{1}$$

Here, $\pi\in\Pi(\mu,v)$ represents a transport plan, which specifies how probability mass associated with a point x under the distribution μ is coupled with probability mass associated with a point y under the distribution v. Thus, the integral in the above equation represents the total transportation cost under a given transport plan π. In general, even when the two distributions μ and v are fixed, the admissible transport plan between them is not unique. The Wasserstein distance is therefore obtained by minimizing the total transportation cost over all admissible transport plans. The use of the infimum “inf” reflects this minimization over all admissible transport plans.

In the context of uncertainty importance evaluation considered in this study, the second-order Wasserstein distance is used to evaluate changes in the distribution of the PRA output Y. Specifically, the distance is evaluated between the reference probability distribution of Y and the conditional probability distribution of Y given X_i.

Let $F_Y(y)$ denote the cumulative distribution function of the reference distribution, and $F_{Y|X_i}(y)$ denote the cumulative distribution function of the conditional distribution. When the output Y is a one-dimensional random variable, the second-order Wasserstein distance can be expressed in terms of the corresponding quantile functions as

$$W_2\left(F_Y(y),\,F_{Y|X_i}(y)\right)=\left(\int_0^1\left\{F_Y^{-1}(u)-F_{Y|X_i}^{-1}(u)\right\}^2 du\right)^{1/2}. \tag{2}$$

Here, $F_Y^{-1}(u)$ and $F_{Y|X_i}^{-1}(u)$ are the quantile functions corresponding to $F_Y(y)$ and $F_{Y|X_i}(y)$, respectively, and $u\in[0,1]$ denotes the cumulative probability level.

Figure 1 schematically illustrates these quantile functions. At each cumulative probability level u, the vertical distance between the two curves represents the transport distance between the corresponding quantiles. The shaded area in the figure visually highlights these pointwise differences over the range of u. The second-order Wasserstein distance W_2 is obtained by first integrating the square of this distance over $u\in[0,1]$ and then taking the square root of the resulting value.

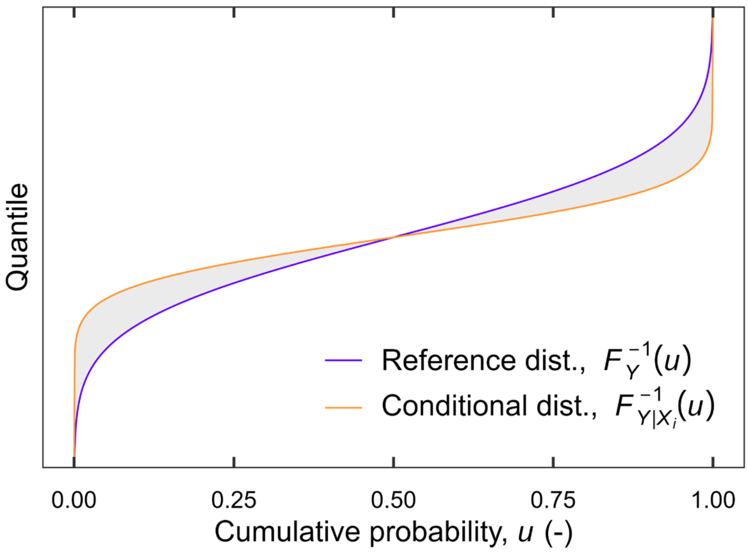

Figure 1: Schematic Illustration of the Wasserstein Distance with Quantile Functions

2.2. Decomposition of Wasserstein Distance

The square of the second-order Wasserstein distance W_2, defined in Eq. (2), is given by

$$W_2^2\left(F_Y(y),\, F_{Y|X_i}(y)\right) = \int_0^1 \left\{F_Y^{-1}(u) - F_{Y|X_i}^{-1}(u)\right\}^2 du. \tag{3}$$

Here, the means of the reference distribution $F_Y(y)$ and the conditional distribution $F_{Y|X_i}(y)$ are defined as

$$\mu_Y = \int_0^1 F_Y^{-1}(u)\,du, \tag{4}$$

$$\mu_{Y|X_i=x_i} = \int_0^1 F_{Y|X_i}^{-1}(u)\,du. \tag{5}$$

The centered quantile functions are then defined by subtracting the corresponding mean from each quantile function:

$$\delta_Y = F_Y^{-1}(u) - \mu_Y, \tag{6}$$

$$\delta_{Y|X_i} = F_{Y|X_i}^{-1}(u) - \mu_{Y|X_i}. \tag{7}$$

Using these definitions, the difference between the two quantile functions can be decomposed into the difference between the means and the difference between the centered quantile functions as follows:

$$F_Y^{-1}(u) - F_{Y|X_i}^{-1}(u) = \left(\mu_Y - \mu_{Y|X_i}\right) + \left[\delta_Y(u) - \delta_{Y|X_i}(u)\right]. \tag{8}$$

Substituting this expression into Eq. (2) yields the following decomposition. The cross term vanishes because the integrals of the centered quantile functions over $u \in [0,1]$ are zero:

$$W_2^2\left(F_Y(y),\, F_{Y|X_i}(y)\right) = \left(\mu_Y - \mu_{Y|X_i}\right)^2 + \int_0^1 \left\{\delta_Y(u) - \delta_{Y|X_i}(u)\right\}^2 du. \tag{9}$$

The first term on the right-hand side is the squared difference between the means of the reference distribution $F_Y(y)$ and the conditional distribution $F_{Y|X_i}(y)$. It represents the extent to which the location of the output distribution is shifted by conditioning on the input parameter. In this study, this term is referred to as the bias component. It indicates the change in the representative value of the PRA output *Y* caused by conditioning on the input parameter X_i. The second term on the right-hand side represents the difference between the quantile functions after removing the difference in their mean locations. This term captures changes in the distribution that cannot be explained by the mean shift, including changes in spread, skewness, and tail behavior. In this study, this term is referred to as the variance component for convenience. More precisely, however, it should be interpreted as a component representing changes in the spread and shape of the distribution after removing the mean.

This decomposition allows the distributional change measured by the Wasserstein distance to be separated into a bias component arising from a shift in the mean location and a component arising from changes in the distributional shape after removing the mean. This is useful in uncertainty importance analysis because it helps identify whether the influence of an input parameter on the PRA output distribution is mainly due to a shift in the mean or to distributional changes after removing the mean.

2.3. Proposed Measures

Based on the decomposition of the squared second-order Wasserstein distance shown in Eq. (9), this study proposes a Wasserstein-based uncertainty importance evaluation method consisting of one primary measure and two auxiliary measures. In the proposed evaluation method, the overall distributional change is first quantified by a primary measure based on the second-order Wasserstein distance. The decomposed components of the squared distance are then used as auxiliary measures to interpret the origin of the overall change, namely whether it is mainly associated with a shift in the mean location or with distributional changes after removing the mean.

Wasserstein Importance Measure

Let $F_{Y|X_i}(y)$ denote the conditional cumulative distribution function of the output *Y* given X_i. The Wasserstein importance measure is defined as the expectation, with respect to the distribution of X_i, of the second-order Wasserstein distance between the reference output distribution $F_Y(y)$ and the conditional output distribution $F_{Y|X_i}(y)$:

$$W_i := \mathbb{E}_{X_i}\left[W_2\left(F_Y(y),\, F_{Y|X_i}(y)\right)\right], \tag{10}$$

where $\mathbb{E}_{X_i}[\cdot]$ denotes the expectation over X_i. This measure represents the overall effect of the uncertainty in the input parameter X_i on the entire distribution of the PRA output *Y*.

Wasserstein–Bias Importance Measure

Based on the decomposition in Eq. (9), the squared second-order Wasserstein distance is separated into a component corresponding to the change in the mean location of the output distribution and a component corresponding to the change in the distributional shape after removing the mean. Using the former component, the Wasserstein–Bias importance measure is defined as

$$W_i^B := \mathbb{E}_{X_i}\left[\left(\mu_Y - \mu_{Y|X_i}\right)^2\right]. \tag{11}$$

This measure represents the extent to which the uncertainty in X_i changes the mean location of the distribution of the PRA output *Y*.

Wasserstein–Variance Importance Measure

Based on the component corresponding to the change in the distributional shape after removing the mean in Eq. (9), the Wasserstein–Variance importance measure is defined as

$$W_i^V := \mathbb{E}_{X_i}\left[\int_0^1 \left\{\delta_Y(u) - \delta_{Y|X_i}(u)\right\}^2 du\right]. \tag{12}$$

This measure evaluates changes in the spread and shape of the distribution that remain after removing the difference in the mean location.

The Wasserstein–Bias and Wasserstein–Variance importance measures are derived from the decomposition of the squared second-order Wasserstein distance and are used as auxiliary measures for interpreting the primary Wasserstein importance measure. Specifically, they indicate whether the overall distributional change is mainly associated with a shift in the mean location or with distributional changes after removing the mean.

3. SENSITIVITY STUDY OF PROPOSED MEASURES

To examine the fundamental behavior of the proposed Wasserstein-based uncertainty importance measures, sensitivity studies were performed using prescribed probability distributions. The analysis considered three representative cases: bias shifts, variance changes, and various distributional shapes. For comparison, the sensitivity of the Borgonovo measure [5] was also evaluated.

The Borgonovo measure σ_i is a moment-independent uncertainty importance measure that quantifies the difference between the reference and conditional probability density functions of the output *Y*. It is defined as

$$s(X_i) = \int | f_Y(y) - f_{Y|X_i}(y) | \, dy, \tag{13}$$

$$\sigma_i := \frac{1}{2}\mathbb{E}_{X_i}[s(X_i)], \tag{14}$$

where $f_Y(y)$ is the reference probability density function of *Y*, and $f_{Y|X_i}(y)$ is the conditional probability density function of *Y* given X_i.

3.1. Sensitivity to Bias Shifts

First, the sensitivity of the proposed measures to bias changes was evaluated using two normal distributions with the same variance but different means.

In this sensitivity study, prescribed probability distributions were used to represent the reference output distribution and the conditional output distribution corresponding to a fixed value of $X_i = x_i$, where x_i is a realization of the uncertain input parameter X_i. The reference output distribution was assumed to follow the standard normal distribution, $N(0,1^2)$, whereas the conditional output distribution was assumed to follow $N(m,1^2)$, whereas *m* was varied from 0.5 to 10. This setting was used to examine the response of the uncertainty importance measures to a pure bias shift between the two distributions.

Figure 2 shows the sensitivity of the Wasserstein and Borgonovo importance measures (W_i, σ_i) to the mean shift of the conditional distribution. W_i increases almost linearly with the magnitude of the mean shift, reflecting its direct dependence on the displacement between the two distributions. In contrast, σ_i increases more gradually and eventually tends to saturate. This is because the Borgonovo measure is based on the integrated absolute difference between the two probability density functions and approaches its upper bound of 1 as the distributions become increasingly separated. This indicates that the Borgonovo measure becomes less sensitive to further increases in large location shifts.

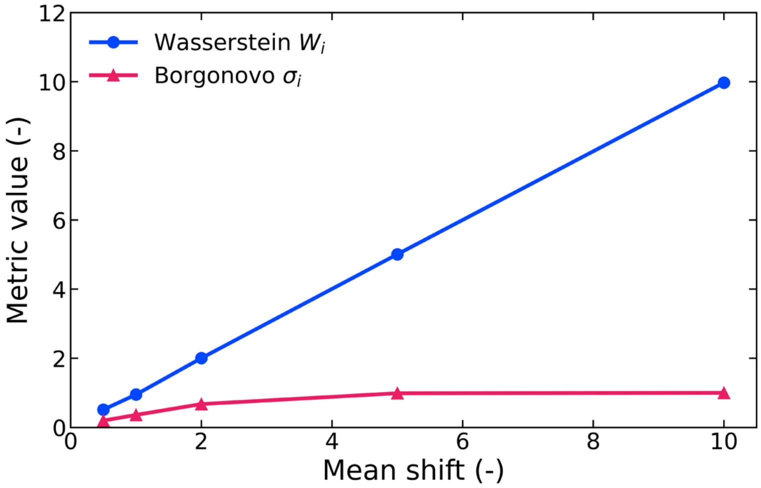

Figure 2: Sensitivity of Uncertainty Importance Measures to Mean Shifts

3.2. Sensitivity to Variance Changes

Next, the sensitivity of the proposed measures to variance changes was evaluated using two normal distributions with the same mean but different variances.

As in Section 3.1, the reference distribution was set to the standard normal distribution, $N\left(0,1^2\right)$. The conditional distribution was also assumed to be $N\left(0,\tau^2\right)$, where τ^2 was varied from 1 to 20. This setting was used to examine the response of the uncertainty importance measures to a pure variance change between the two distributions.

Figure 3 shows the sensitivity of the Wasserstein and Borgonovo importance measures (W_i, σ_i) to changes in the variance of the conditional distribution. The figure shows that W_i increases monotonically as the variance of the conditional distribution increases. This is because a larger spread increases the transport distance required to transform the reference distribution into the conditional distribution, even when the two distributions have the same mean. Thus, W_i is sensitive not only to location shifts but also to changes in distributional spread.

σ_i also increases with the variance, but more gradually, and its rate of increase decreases for large variance values. As in the bias-shift case, this tendency reflects the bounded nature of the Borgonovo measure: when the difference between the two distributions becomes sufficiently large, σ_i approaches its upper bound of 1 and becomes less sensitive to further increases in distributional difference.

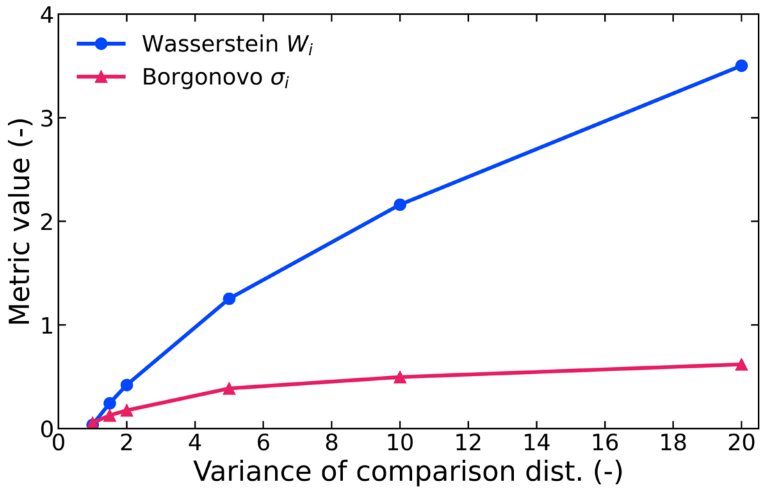

Figure 3: Sensitivity of Uncertainty Importance Measures to Variance Changes

3.3. Sensitivity to Various Distributional Shapes

Finally, the sensitivity of the proposed measures to different distributional shapes was evaluated. In Sections 3.1 and 3.2, the response of the measures was examined using normal distributions in which only the mean or variance was varied. In this section, more general changes in distributional shape are considered, including heavy-tailed behavior, tail shifts, multimodality, and variance mixtures.

As in Sections 3.1 and 3.2, the reference distribution was set to the standard normal distribution, $N(0,1^2)$. The following distributions were then used as conditional output distributions:

Bias change:

$$Y \mid X_i = x_i \sim N(0.7,1^2), \tag{15}$$

Variance change:

$$Y \mid X_i = x_i \sim N(0,1.8^2), \tag{16}$$

Added mode (bimodal):

$$Y \mid X_i = x_i \sim 0.5N(0,1^2)+0.5N(3,1^2), \tag{17}$$

Extreme tail shift:

$$Y \mid X_i = x_i \sim 0.99N(0,1^2)+0.01N(5,1^2), \tag{18}$$

Added minor mode:

$$Y \mid X_i = x_i \sim 0.95N(0,1^2)+0.05N(3,0.1^2), \tag{19}$$

Variance mixture:

$$Y \mid X_i = x_i \sim 0.5N(0,0.5^2)+0.5N(0,2.0^2). \tag{20}$$

The histograms of the distributions defined above are shown in Figure 4.

Figure 5 compares the calculated Wasserstein and Borgonovo importance measures for these distributional changes. In this figure, each measure is normalized by its maximum value to facilitate comparison of the relative responses across different measures. The relative importance trends for the different distributional changes are generally consistent between the Wasserstein Importance Measure and the Borgonovo measure.

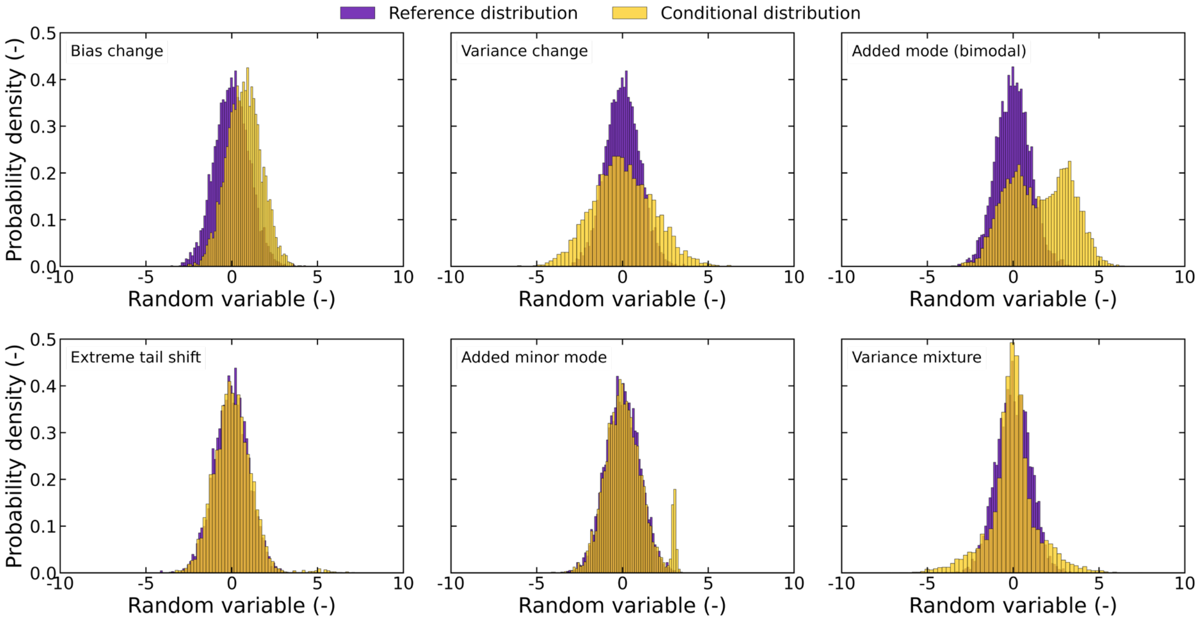

Figure 4: Histograms of the Reference and Conditional Distributions Used to Examine Sensitivity to Various Distributional Shapes

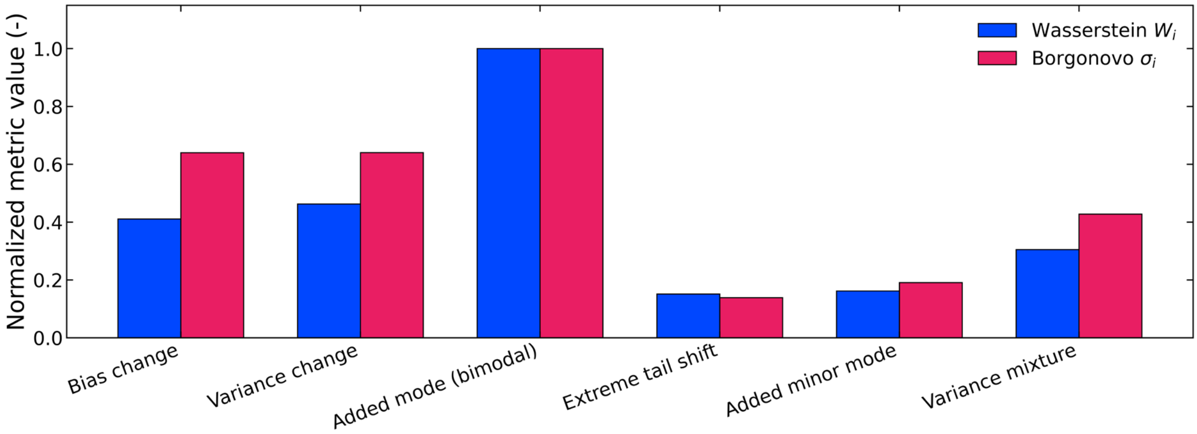

Figure 5: Sensitivity of Uncertainty Importance Measures to Various Distributional Shapes

Figure 6 shows the Wasserstein–Bias and Wasserstein–Variance components. The vertical axis is normalized by the maximum value of the squared second-order Wasserstein distance, i.e., the maximum value of the sum of the two components. For the bias change case, only the Wasserstein–Bias component appears, indicating that the proposed decomposition correctly identifies the change as a shift in the mean location of the distribution. In contrast, for the variance change case, the Wasserstein–Variance component is dominant, showing that the change is attributed to the spread of the distribution rather than to a shift in the mean. For the case with an added mode resulting in a bimodal distribution, the Wasserstein–Bias component is large, while the Wasserstein–Variance component also contributes. This indicates that the added mode changes both the mean location and the shape of the distribution.

For the extreme tail shift, the added minor mode, and the variance mixture cases, the Wasserstein–Variance component is dominant, suggesting that these changes are mainly associated with distributional shape, such as tail behavior, modal structure, or dispersion, rather than with the mean location.

These results demonstrate that the proposed decomposition enables the effect of distributional changes to be separated into a component associated with the mean location and a component associated with the distributional shape after removing the mean. This distinction is difficult to obtain from existing uncertainty importance measures that evaluate only the overall discrepancy between unconditional and conditional output distributions.

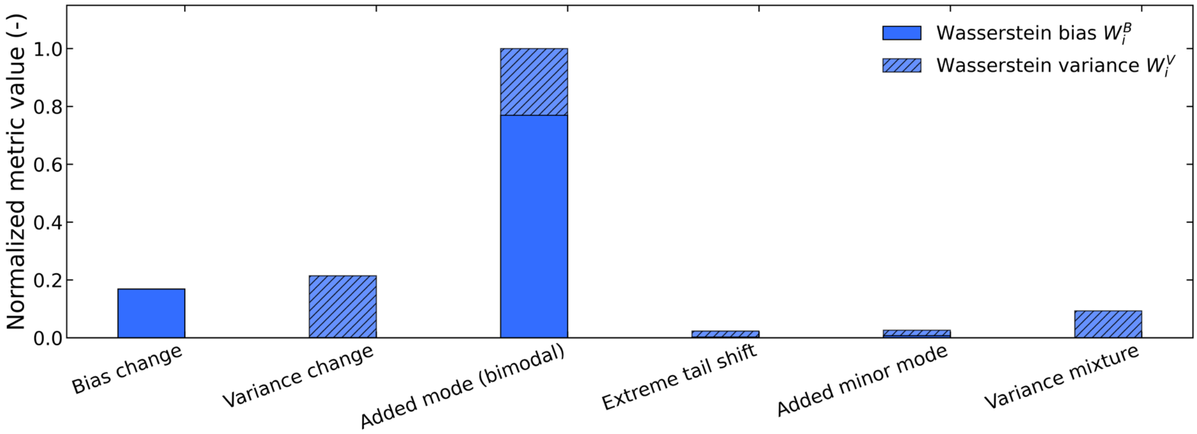

Figure 6: Sensitivity of Wasserstein–Bias and Variance Components to Various Distributional Changes

4. CONCLUSION

This study proposed a Wasserstein-based method for uncertainty importance evaluation to assess the influence of input parameter uncertainty on the overall uncertainty of PRA outputs. The basic applicability and usefulness of the proposed method was examined through numerical experiments using prescribed probability distributions.

In the proposed method, the second-order Wasserstein distance between the reference output distribution and the conditional output distribution is used to define the Wasserstein Importance Measure as the primary measure for quantifying the overall influence of input uncertainty on the output distribution. In addition, the squared second-order Wasserstein distance for one-dimensional distributions is decomposed into a Wasserstein–Bias component, which corresponds to changes in the mean location, and a Wasserstein–Variance component, which corresponds to distributional changes after removing the mean. Although the latter component is referred to as the variance component in this study, it should be interpreted more broadly as representing changes in spread, tail behavior, multimodality, and other shape-related features of the distribution.

Numerical experiments were conducted for bias changes, variance changes, and several types of distributional shape changes. The results showed that the relative importance trends obtained by the Wasserstein Importance Measure were generally consistent with those obtained by the Borgonovo measure. At the same time, because the Wasserstein Importance Measure is based on the transport distance between distributions, it showed clear sensitivity to distributional changes involving large deviations in output values, such as large mean shifts, tail shifts, and extreme-value shifts.

The results also demonstrated that the auxiliary components provide information that is difficult to obtain from conventional uncertainty importance measures. Specifically, the Wasserstein–Bias and

Wasserstein–Variance components make it possible to interpret whether the effect of input uncertainty is mainly caused by a shift in the mean location or by distributional changes after removing the mean.

These findings indicate that the proposed method can provide relative importance evaluations that are broadly consistent with existing measures, while also offering additional interpretability regarding the nature of output distribution changes. This feature may be useful for prioritizing data collection and model improvement efforts aimed at reducing epistemic uncertainty, because it provides not only the magnitude of importance but also insight into what aspect of the output distribution change should be targeted.

Future work will apply the proposed method to practical PRA models and examine its applicability to cases involving multiple uncertain input parameters. Extensions to account for dependencies among input parameters should also be investigated.

References

[1] J. C. Helton and D. E. Burmaster, "*Guest editorial: treatment of aleatory and epistemic uncertainty in performance assessments for complex systems*," Reliability Engineering and System Safety, 54, pp. 91–94 (1996).
[2] E. Zio, "*Risk importance measures*," Safety and Risk Modeling and Its Applications, pp. 151–196 (2011).
[3] T. Aven and T. E. Nøkland, "*On the use of uncertainty importance measures in reliability and risk analysis*," Reliability Engineering & System Safety, 95, pp. 127–133 (2010).
[4] R. L. Iman, "*A matrix-based approach to uncertainty and sensitivity analysis for fault trees*," Risk Analysis, 7, pp. 21–33 (1987).
[5] E. Borgonovo, "*A new uncertainty importance measure*," Reliability Engineering & System Safety, 92, pp. 771–784 (2007).
[6] C. Villani, "*Optimal transport: old and new*," Springer, Berlin (2009).