Selecting technologies in aircraft conceptual design using probabilistic inversion

Roelofs, Martijn N.; Kurowicka, Dorota; Vos, Roelof

DOI
10.2514/1.C035985

Publication date
2021

Document Version
Final published version

Published in
Journal of Aircraft

Citation (APA)
Roelofs, M. N., Kurowicka, D., & Vos, R. (2021). Selecting technologies in aircraft conceptual design using probabilistic inversion. *Journal of Aircraft*, *58*(2), 347-362. https://doi.org/10.2514/1.C035985

JOURNAL OF AIRCRAFT
Vol. 58, No. 2, March–April 2021

Selecting Technologies in Aircraft Conceptual Design Using Probabilistic Inversion

Martijn N. Roelofs,* Dorota Kurowicka,† and Roelof Vos‡
Delft University of Technology, 2600AA Delft, Zuid-Holland, The Netherlands

https://doi.org/10.2514/1.C035985

Technology forecasting is an essential starting point for conceptual design of any complex engineering system. In fact, many research projects are focused on developing a small set of promising technologies to a suitable readiness level. However, selecting a set of technologies from a larger pool is a nontrivial task, opposed by uncertainty and subjective tradeoffs. This paper proposes a probabilistic method to represent technologies and quantify their effects, while accounting for uncertainty. Using probabilistic inversion, technologies can be selected from a larger set to meet a certain combination of requirements. Several test cases illustrate the method and how it may be used in conceptual design projects. It is concluded that probabilistic inversion enables answering technology development and selection queries, which would be challenging to answer with traditional deterministic approaches, or purely forward uncertainty propagation approaches.

Nomenclature

e	=	Oswald factor
F	=	cumulative distribution function
$\mathbb{I}$	=	indicator function
k	=	impact factor
M_F	=	fuel mass
M_{MTO}	=	maximum takeoff mass
M_{OE}	=	operating empty mass
M_P	=	payload mass
M_{ZF}	=	zero-fuel mass
$\mathcal{N}$	=	Gaussian distribution function
$p(\cdot)$	=	probability
q	=	interquantile
R	=	range
s	=	sample
$\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}$	=	random vectors
$\boldsymbol{x}$	=	sample vector, optimization vector
α	=	Copula shape parameter
Δ	=	difference
δ	=	deflection angle
μ	=	mean
σ	=	standard deviation
τ	=	Kendall's tau

Subscripts

cl	=	climb
cr	=	cruise
des	=	descent
e	=	Oswald factor
f	=	flap
i	=	ith entry
M	=	mass
max	=	maximum
OEM	=	operating empty mass
req	=	requirement
tot	=	total
0	=	initial

Received 1 April 2020; revision received 21 July 2020; accepted for publication 22 July 2020; published online 12 October 2020.

*Ph.D. Candidate, Faculty of Aerospace Engineering, Kluyverweg 1.
†Associate Professor, Delft Institute of Applied Mathematics, Van Mourik Broekmanweg 6 2628XE.
‡Assistant Professor, Faculty of Aerospace Engineering, Kluyverweg 1. Associate Fellow AIAA.

I. Introduction

DURING conceptual design, decisions are made regarding which configurations and technologies to include in an engineering system. For example, questions arise on which wing movables to employ and where to position them, what material and structural layout is optimal for the wing box, whether to use a hydraulic or an electric actuator system, and so on. We call such questions inverse (design) queries because they involve working from a computed quantity of interest back to the input variables and figuring out how the latter should change to satisfy certain requirements. Typically, such design decisions are supported by deterministic investigations; the different options are represented in some simulation model, and quantities of interest (QOI) are computed for a specific or small set of designs or missions. Such an approach provides poor generalization of the conclusions because these are specific to the design or mission chosen. Therefore, very little can be concluded about the performance of a certain technology or configuration across the board.

Many design queries can be answered with optimization approaches. When including uncertainty, one can use either robust design optimization or reliability-based design optimization [1–3]. These approaches usually convert a probability distribution into scalar measures like mean and variance. Moreover, they tend to focus on specific designs and/or missions. However, technology selection takes place when the precise designs and missions are still unknown. Design space exploration is a solution that avoids this issue. The high-dimensional design space can be visualized with cobweb plots. These display the joint distribution of the variables and can subsequently be used to single out regions (i.e., value ranges of variables) of interest [4]. Cobweb plots can only be used to analyze the current design space and show which ranges of variables influence other variables. Figuring out how the design space should change to achieve certain goals is not possible and requires other methods.

More robust and reliable decisions may be made when uncertainty is included in the analyses [5]. Chakraborty and Mavris [6] confirm that statement in practice, by comparing more-electric system architectures with conventional ones probabilistically. Technology impact forecasting (TIF) [7,8] generalizes the workflow of probabilistic evaluation of technologies. Only forward uncertainty propagation is performed, however, so inverse problems cannot be directly addressed.

After the technology impacts are evaluated, different approaches to ranking and selecting the technologies (or portfolios) are available. Gatian and Mavris [9–11] use an approach very similar to TIF and perform portfolio selection using a probability of success and a signal-to-noise ratio metric, which summarize cumulative distribution

function (CDFs) into a scalar value, such that alternatives can be compared. Although these measures aim to capture as much of the CDF as possible, information is still lost. Instead, Jimenez and Mavris [12] determine Pareto-optimal solutions, which are the set of technologies that lie on a Pareto front of the objectives. A similar approach [13] uses stochastic dominance to remove design alternatives (i.e., technology portfolios) from the complete set, leaving decision makers with only optimal alternatives, to reduce the risk of incorrect decisions. While these approaches do retain the full information of the CDFs, it becomes harder to rank the technologies because crossing CDFs do not exhibit stochastic dominance and subjective judgment is needed to prefer one solution over another.

Even though all of the mentioned research uses probability theory, there are alternatives: Bayesian theory, possibility theory [14], interval analysis [15,16], and Dempster–Schafer (D-S) theory [17–19]. Akram et al. [20–24] use D-S theory for a technology valuation process–technology portfolio analysis. While there are strong arguments against theories other than (subjective) probability theory [25,26], the most relevant argument in the current context is that only Bayesian theory supports the inverse problem, besides probability theory. Unfortunately, the Bayesian inverse techniques are computationally expensive for continuous or larger problems.

This paper investigates probabilistic inversion (PI) to support technology development and selection. The hypothesis is that this approach circumvents the drawbacks associated with optimization and design space exploration approaches and is a useful means to prioritize technologies when their specification and effects are uncertain. Forward uncertainty propagation is performed through sampling (similar to TIF), and PI is employed to answer inverse (design) queries. Probabilistic inversion is explained in Sec. II. The application consisting of the model, the input variables, and their distributions are discussed in Sec. III. Two test cases are performed on this application. First, a single technology is analyzed with both forward uncertain quantification (UQ) and PI in Sec. IV. Second, a set of three technologies is assessed, and PI is used to infer which technologies are most suitable to meet user-imposed requirements in Sec. V. The merits and drawbacks of the method are discussed in Sec. VI, and the paper is concluded in Sec. VII by establishing that probabilistic inversion is a powerful method for technology prioritization for selection during conceptual design and offers possibilities that conventional optimization or design space exploration techniques do not.

II. Methodology

PI is employed in this paper to solve inverse queries. The aim is to find the values for some input values that achieve certain goals defined on the output variables. In this section, the principles behind PI are discussed, and the algorithms to solve PI problems are presented.

Before diving into PI itself, though, an overview of how PI fits in with uncertainty quantification and propagation is provided in Fig. 1. The first part of the method, uncertainty quantification and propagation, is common to other approaches [7–11]. This part entails setting up the input variables and their (joint) distributions. Not all existing approaches employ joint distributions, but it is not unique to PI. In this paper, the joint distributions are specified using copulas, which are explained in Sec. II.B. All the variable distributions are then sampled N times, and each sample is propagated through the analysis models that compute the QOI. The sampling procedure is indifferent to the type of model, so any black box input–output model works. Therefore, the models can be anything from physics-based simulations to design codes to surrogate models. The output samples are used to construct CDFs on the QOI. Further metrics may be derived from these CDFs to support decision making. When the output does not satisfy certain goals, traditional approaches require an iterative procedure, indicated by the dashed line in Fig. 1. Rather simply, the input is changed iteratively to obtain the desired output.

This is where PI comes in. Instead of adopting an iterative procedure, PI is one additional step that is conducted after the uncertainty propagation. Constraints are imposed on any of the variables (input and output), and PI figures out how the distributions of the other variables change to satisfy those constraints. The updated distributions are used to answer design queries or technology prioritization queries.

A. Probabilistic Inversion

Probabilistic design, in the most general form, requires finding input random vector $\boldsymbol{X} \in \mathbb{R}^N$ (with some properties) to a particular model $G{:}\mathbb{R}^N \mapsto \mathbb{R}^M$ that yields output vector $\boldsymbol{Y} \in \mathbb{R}^M$ with a prespecified distribution. In other words, we want to find $\boldsymbol{X}$ such that $G(\boldsymbol{X}) \sim \boldsymbol{Y}$ (where $\sim$ denotes having the same distribution); hence, invert the function G over random variables. This process is called PI [27].

For some simpler types of models, PI can be carried out with Bayesian methods. In many applications, however, the function G is very complicated, so the distribution of $G(\boldsymbol{X})$ is found via simulations. In such cases, PI can be performed via sample reweighting [28], which avoids inverting the function G, by reweighting the samples of inputs and outputs of the model to satisfy specified constraints. The sample reweighting process is as follows. For a set of samples $(x_1^{(i)}, \ldots, x_N^{(i)})$ generated by sampling from $\boldsymbol{X}$, a set of output samples is computed as

$$\left(y_1^{(i)}, \ldots, y_M^{(i)}\right) = G\left(x_1^{(i)}, \ldots, x_N^{(i)}\right)$$

The sample file $s_i = (x_1^{(i)}, \ldots, x_N^{(i)}, y_1^{(i)}, \ldots, y_M^{(i)})$, $i = 1, \ldots K$ is obtained (shown in Fig. 1), where each sample has the same probability $p(s_i) = 1/K$. The idea is to find a different than uniform distribution over the samples. That is, the weights w_i have to be found, such that resampling with these weights leads to a new set of samples that satisfies the requirements on $\boldsymbol{Y}$ and $\boldsymbol{X}$.

To make this problem computationally tractable, the requirements on $\boldsymbol{Y}$ and $\boldsymbol{X}$ are in the form of quantiles or percentiles of variables in $\boldsymbol{Y}$ and $\boldsymbol{X}$ or functions of these variables. In this simplified form, PI via sample reweighting can be viewed as an optimization problem to find weights $(w_1, \ldots, w_K)$ which minimize relative information with respect to the uniform distribution subject to linear constraints corresponding to percentiles of variables in $\boldsymbol{Y}$ and $\boldsymbol{X}$ or their functions. In theory, there are many distributions that could satisfy the imposed constraints, which is why the minimum relative information gives the distribution closest to the starting distribution, in other words, the uniform distribution. If we were only interested in computing the mean, then updating the sample weights and computing $E(X) = \sum w_i \cdot x_i$ would suffice. However, we are interested in the full distributions, which requires resampling with the new weights to construct them.

In principle, this problem can be solved with help of optimization software. In this paper, however, the PI problem via sample reweighting is solved with the Iterative Proportional Fitting (IPF) algorithm,

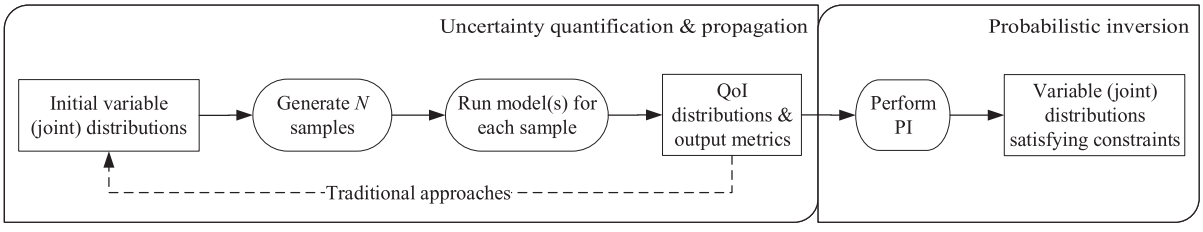

Fig. 1 Overview of the method in three steps: uncertainty quantification, propagation and probabilistic inversion. The information from PI can be used for subsequent technology selection.

introduced in Ref. [29]. This algorithm is much faster than optimization approaches. IPF starts with a uniform distribution over the samples and then iteratively reweights them to satisfy the imposed constraints, one constraint at the time. If the optimization problem is feasible, IPF converges and provides a minimally informative solution to the problem with respect to the starting distribution [30]. IPF will not converge when the problem does not have a solution. In this case, algorithms that minimize infeasibility are proposed, such as PARFUM [31] and PREJUDICE [32,33].

1. Iterative Proportional Fitting

IPF iteratively reweights samples to satisfy the imposed constraints. Those constraints are provided as percentiles for a set of variables. Even if the complete sample set contains all variables $\boldsymbol{X}$, IPF only requires the constrained variables $\boldsymbol{Z} \subseteq \boldsymbol{X} \cup \boldsymbol{Y} \cup \boldsymbol{H}(\boldsymbol{X}, \boldsymbol{Y})$. Thus, $\boldsymbol{Z}$ is a subset of $\boldsymbol{X}$, $\boldsymbol{Y}$, and functions of these two. The samples of $\boldsymbol{Z}$ are presented as a matrix,

$$\boldsymbol{Z} = \begin{bmatrix} z_{11} & \cdots & z_{1M} \\ \vdots & \ddots & \vdots \\ z_{K1} & \cdots & z_{KM} \end{bmatrix} \tag{1}$$

where K is the number of samples and M is the number of variables in $\boldsymbol{Z}$.

For each variable Z_m, a vector of percentiles is specified. From these percentiles, the interpercentiles q_m are derived. For example, the 5, 50, and 95% percentiles are specified for Z_m: then, $q_m = [0.05, 0.45, 0.45, 0.05]$. All vectors q_m are stored in the set $\mathcal{Q} = \{q_m\}$. It is defined as a set of vectors rather than a matrix because this allows for different amounts of percentiles to be specified for different variables.

For each percentile, the value of the corresponding variable has to be specified. Thus, let $Q_m = |q_m|$ be the amount of quantile constraints q_m on the variable Z_m. Then, r_m has length $|r_m| = Q_m - 1$, where each entry r_{mj} is computed as

$$\forall\ j = 1, \ldots, Q_m - 1 : r_{mj} \equiv P(Z_m \leq r_{mj}) = \sum_{i=1}^{j} q_i \tag{2}$$

For all variables, the vectors r_m are collected into a set $\mathcal{R} = \{r_m\}$.

The matrix Z, the interpercentile set $\mathcal{C}$, and the constraint values $\mathcal{R}$ together are the inputs to the IPF algorithm. The algorithm starts by creating a set of indicator matrices, which specify which samples belong to which interpercentiles. Therefore, for each variable, there is a matrix $A^m \in \mathbb{R}^{K \times Q_m}$ given by

$$A^m_{ij} = \begin{cases} Z_{im} \in (-\infty, r_{m,1}] & \text{if } j = 1 \\ Z_{im} \in (r_{m,Q_m-1}, \infty) & \text{if } j = Q_m \\ Z_{im} \in (r_{m,j-1}, r_{m,j}) & \text{otherwise} \end{cases} \tag{3}$$

Thus, each column corresponds to an interpercentile $q_{m,j}$, and each row evaluates whether the value of that sample in $\boldsymbol{Z}$ is within the range specified by r_m for that interpercentile. The set $\mathcal{A} = \{A^m\}$ is the last piece of information IPF requires.

IPF starts with an initial vector of sample probabilities p, which is the uniform distribution over the samples. Thus, each values $p_i = p(s_i) = 1/K$. Then, an outer loop runs through a prescribed number of iterations. During each iteration, the vector $\boldsymbol{p}$ is updated for each variable separately. Therefore, an inner loop runs over the variables Z_m and performs the following operation:

$$p_i' = \sum_{j=1}^{Q_m} A^m_{ij} \frac{p_i q_j}{\sum_{i \in A^m_j} p_i} \tag{4}$$

This process is shown in Fig. 2. After each inner loop iteration, p_i is set to p_i'. As such, IPF updates $\boldsymbol{p}$ to satisfy each quantile constraint on each variable one at a time, while possibly violating the constraints over a previous variable. However, as the iterations proceed (and if the problem has a solution), the constraints for each variable will be satisfied [27,30].

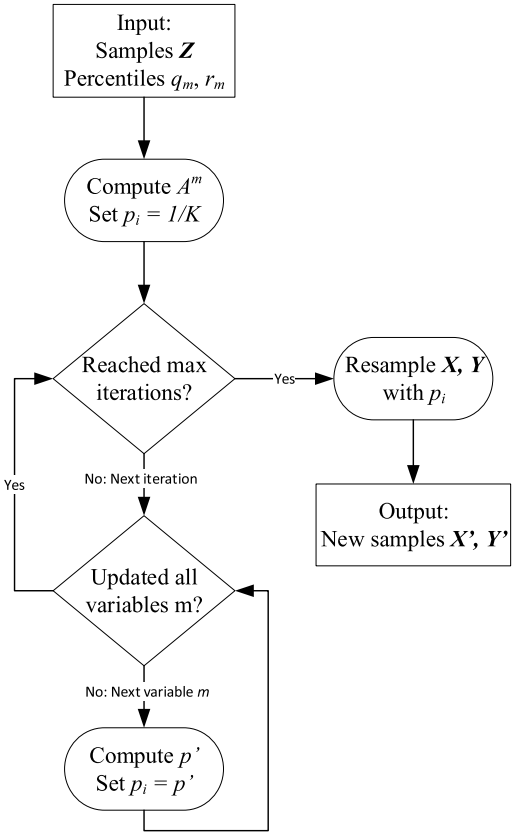

Fig. 2 Flow chart of the IPF algorithm. The outer loop runs over a prescribed maximum number of iterations, while the inner loop iterates over all variables m.

2. Error of IPF

The outer-loop of IPF runs over a prescribed number of iterations, but it could happen that IPF reaches satisfactory convergence before that limit is reached, or instead does not converge on time. With the constraints and the vector $\boldsymbol{p}$, the achieved percentiles can be computed and compared to the specified percentiles. This provides a measure of the error of IPF at any given iteration.

The constraints, specified by the combination of the quantiles $\mathcal{Q}$ and their values $\mathcal{R}$, can be written as a linear combination of the sample probability vector $\boldsymbol{p}$. They form the constraint set $\mathcal{C}$, which contains for all $m = 1, \ldots, M$ and $j = 1, \ldots, Q_m$ the following equality:

$$C_{j,m} : \sum_{i=1}^{K} p_i A^m_{ij} = q_{j,m} \tag{5}$$

Essentially, each constraint counts the number of samples that would fall within a certain interpercentile, given the new probabilities p_i. Thus, there are $|\mathcal{C}| = \sum_{m=1}^{M} Q_m$ constraints: one for each combination of percentile and variable in $\boldsymbol{Z}$. The set $\mathcal{C}$ can now be written as a set of linear equations $\boldsymbol{C} \cdot \boldsymbol{p} = \boldsymbol{q} \cdot \boldsymbol{C}$ is a $|\mathcal{C}| \times K$ matrix, with each entry $C_{(j,m),i} = A^m_{ij}$. The vector $\boldsymbol{p}$ of length K contains p_i, and $\boldsymbol{q}$ are the percentiles, reshaped into a vector of length $|\mathcal{C}|$. The vector $\boldsymbol{q}$ should have the same values as the corresponding entries in $\mathcal{Q}$. As a measure of the error, the maximum absolute value of $\boldsymbol{q} - \boldsymbol{C} \cdot \boldsymbol{p}$ is taken.

3. PARFUM

IPF may not always converge, especially when the provided constraints make the problem infeasible. In such cases, an approximate

algorithm may be used. PARFUM is such an algorithm [31], which works similarly to IPF. Unlike IPF, however, PARFUM updates each variable separately on each iteration [identical to Eq. (4)],

$$p_i^m = \sum_{j=1}^{Q_m} A_{ij}^m \frac{p_i q_j}{\sum_{i \in A_j^m} p_i} \tag{6}$$

after which a mean is taken of the weights for each sample. In this case, the geometric mean is used:

$$p_i = \left(\prod_{m=1}^{M} p_i^m \right)^{\frac{1}{M}} \tag{7}$$

When the problem is feasible, this gives the same result as IPF. When the problem is infeasible, PARFUM provides a minimally informative solution that minimizes the distance to the constraints [27].

4. *Grouped Sample Reweighting*

For reasons that will become apparent in Sec. V, it is sometimes desired to only reweight groups of samples simultaneously, instead of individually for each sample as IPF and PARFUM do. This way, the distributions over the variables in a group of samples are left unaltered. To achieve this, the sample set is divided into P disjoint groups. After PI is performed on the full vector $\boldsymbol{p}$, each sample has its own weight. However, now each sample in a group should receive the same weight. This is done by simply taking the mean of the sample weights for each group and assigning that weight to each sample in the group:

$$\forall s_i \in \mathcal{P}_j, j = 1, \ldots, P : p(s_i) = \frac{\sum_{s_k \in \mathcal{P}_j} p(s_k)}{|\mathcal{P}_j|} \tag{8}$$

Afterward, $\boldsymbol{p}$ is scaled to sum to 1. Note that, because solving PI problems with grouped sample reweighting often is infeasible, the PARFUM algorithm should be used instead of IPF. However, we have not yet been able to prove that PARFUM with this additional step converges.

B. Copulas

In many practical applications, multivariate distributions have to be specified over the input variables, instead of just univariate distributions. This is the case when two or more input variables are codependent. Copulas offer a general framework for specifying such multivariate distributions using any univariate marginals and a copula function C that specifies the dependency structure between the marginals [34]. The copula function C relates a set of real random variables $U_1, \ldots, U_N$ with standard uniform[§] margins as

$$C(U_1, \ldots, U_N) = P(U_1 \leq u_1, \ldots, U_N \leq u_N) \tag{9}$$

The significance of copulas is a result of Sklar's theorem [35] that states that any multivariate distribution H can be represented as a copula function of its marginals,

$$H(X_1, \ldots, X_N) = C(F_1(X_1), \ldots, F_N(X_N)) \tag{10}$$

where $F_i(X_i) = P(X_i \leq x_i)$; in other words, it is the cumulative distribution function of the variable X_i and represents its marginal distribution. Because $F_i{:}X_i \mapsto [0, 1]$, they are the U_i in Eq. (9).

To specify a multivariate distribution H, a copula should be specified that captures the dependency as required. Because it has standard uniform margins, sampling from a copula is straightforward. Then, knowing for each variable X_i the CDF F_i and its inverse F_i^{-1}, each sampled variable U_i can be transformed into the values for X_i as follows:

$$X_i = F_i^{-1}(U_i) \tag{11}$$

The copula of a set of samples of X_i can be obtained by transforming each X_i into copula space, in other words, its corresponding U_i. When the CDF F_i is known, this is trivial. However, when it is not known, it may be estimated through kernel density smoothing, for example. Alternatively, the ranks of the values in X_i can be computed and divided by the amount of samples in X_i plus 1:

$$U_i \sim \frac{\text{rank}(X_i)}{|X_i| + 1} \tag{12}$$

Not any function is suitable as the copula function C. The ones that have been developed are roughly divided into two families: Gaussian and Archimedean. The Archimedean family contains, among others, the Clayton, Gumbel, and Frank copulas. It is not relevant for this text to further elucidate these copulas and their differences. However, they are so-called parametric copulas because they are governed by a parameter that indicates the dependence strength. In the following, this parameter is called α and can be computed from common correlation coefficients, such as Kendall's τ.

C. Example of Applying PI

Before continuing with a more realistic problem setting, a simple example is discussed to show how probabilistic inversion works in practice. Consider the mass breakdown of the maximum zero fuel mass M_{ZF} computed as

$$M_{\text{ZF}} = M_{\text{OE}} \cdot (1 - k_{\text{OEM}}) + M_P \tag{13}$$

where M_{OE} is the operating empty mass, k_{OEM} is a percentage mass reduction, and M_P is the passenger or payload mass. This mass breakdown is not representative of the real world and only is constructed for a simple demonstration of PI.

We assume that k_{OEM} is uniform on the interval [0, 0.3] and is independent of M_{OE} and M_P. The margins of M_{OE} and M_P are assumed uniform on $[1e4, 2e5]$ and $[5.5e3, 7.7e4]$, respectively. These values are purely notional. In a real case, they would have to be estimated from data or expert elicitation. Furthermore, from engineering insight, it is logical that M_{OE} and M_P are correlated because a higher payload mass leads to higher structural mass. Conversely, no more structural mass is present than strictly necessary for a given payload mass. This dependency is modeled with a Frank copula with $\alpha = 18.1915$ corresponding to Kendall's τ of 0.8. The Frank copula is chosen because it is symmetric and has no tail dependency. The value of 0.8 is chosen arbitrarily and should in real applications be estimated from data or with an expert solicitation procedure. The resulting scatter plot of samples from the joint distribution of M_{OE} and M_P is shown in Fig. 3a.

Now, a given set of aircraft is considered for which M_{OE} and M_P follow the uniform distributions as specified in the preceding paragraph, and the percentage mass reduction is k_{OEM}. Suppose that for this range of aircraft some new technology can be considered that allows to alter the distribution of k_{OEM}. The question is how much percentage M_{OE} reduction is required for a given reduction in M_{ZF}.

The requirements on the distribution of M_{ZF} are specified in the form of 5th, 10th, 30th, 50th, 70th, 90th, 95th, and 99th percentiles, which are equal to 0.2428, 0.3353, 0.7440, 1.1605, 1.5740, 1.9721, 2.1083, and 2.2932 (all $\times 10^5$), respectively. These values are 11% smaller than the original distribution of M_{ZF}.

PI requires the problem to be translated into a set of constraints. Recall that constraints take the form of the quantiles $\mathcal{Q}$ and their values $\mathcal{R}$. Essentially, these are points on the CDFs that the practitioner desires to obtain after performing PI.

We distinguish two different functions for constraints: either a constraint fixes a variable's distribution so PI cannot alter it or it reflects a new distribution that is desired by the designer. For the first function, the percentiles of the original distributions can be specified as constraints. However, that does not mean that the distributions are not changed by PI at all. In case of a joint distribution, all linear combinations of the

[§]A standard uniform distribution is a uniform distribution on [0, 1].

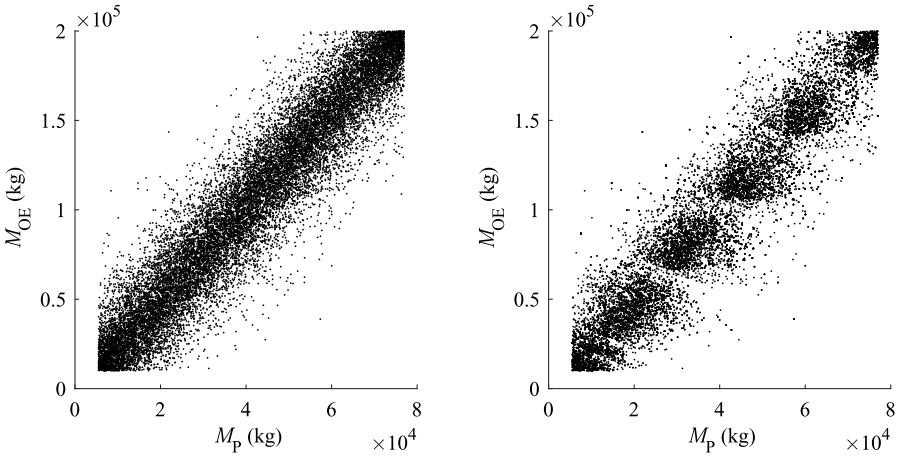

a) Original distribution **b) Distribution after PI**

Fig. 3 Joint distribution between M_{OE} and M_P, modeled using a Frank copula, before and after PI.

variables would have to be constrained. Thus, a distribution can never be fully constrained in a way that keeps it completely unaltered during PI. Nonetheless, these constraints do limit the extent to which a distribution is affected and are therefore sufficient. Examples of how and why certain variables or their joint distributions are constrained can be found later in Secs. V.A and V.B. The second constraint function requires the percentiles of a new distribution. These percentiles may be obtained from data or from expert elicitation or policy makers using existing techniques. Naturally, they may be assumed to get an indication of the resulting responses.

For this example, a fixed set of aircraft is considered. Thus, M_{OE} and M_P need to be kept relatively unchanged, and only the distribution of k_{OEM} should be adjusted to the requirements on the distribution of M_{ZF}. To that end, M_{OE} and M_P are constrained to keep their 5th, 10th, 30th, 50th, 70th, 90th, 95th, and 99th percentiles close to the original. Thus, for each of these variables, each of the percentiles is specified with the value of the original uniform distribution. Furthermore, to keep the dependency between M_{OE} and M_P similar, the distribution of the sum of M_{OE} and M_P is also constrained on the same eight percentiles. Therefore, the dependency imposed by Fig. 3a should remain identical. Dropping the constraints on the margins or dependence of M_{OE} and M_P would permit PI to alter their distributions, which would consequently reflect a different set of aircraft.

With the aforementioned constraints, the IPF algorithm is run, and a new sample set is generated. The CDFs for the new samples are shown in Fig. 4 (as dashed, red lines), with the original CDFs for reference (as blue, solid lines). Notice that the CDFs of M_{OE} and M_P hardly have changed due to the constraints, while k_{OEM} becomes skewed to the right as expected.

To better illustrate what PI does, the probability density functions (PDFs) corresponding to the CDFs are shown in Fig. 5. Again, the red, dashed lines show the PDFs after PI, while the blue, solid lines are the original PDFs. The distribution over k_{OEM} has shifted all the way to its largest value. Observe the oscillatory nature of the PDFs for M_{ZF}, M_{OE}, and M_P. This is an important effect of using percentiles rather than complete distributions as constraints; thus, it is inherent to IPF. PI picks samples within each interpercentile and reweights those equally to meet a certain percentile constraint. Therefore, the distributions only match at the specified percentiles, and may vary in between.

A similar observation can be made for the bivariate distribution of M_{OE} and M_P in Fig. 3. After PI (Fig. 3b), the scatter plot is different than the original one. Clear discontinuities can be seen in the samples corresponding to the oscillations of the PDFs shown in Fig. 5. Nonetheless, the dependency between M_{OE} and M_P, measured as the Pearson coefficient, remains almost identical to the initial value of 0.95, due to the constraint on the sum of these two variables.

If the constraints on M_{OE}, M_P, and their sum are not included, the results of PI look very different. These are shown in Figs. 4 and 5 as dotted, red lines. From both figures, it becomes clear that k_{OEM} is hardly affected, while the distributions of M_{OE} and M_P show a larger deviation form the original. This is easily understood as PI picks the most influential variables to achieve the constraints. In this case, obviously, M_{ZF} is most easily reduced by reducing either M_{OE} or M_P or both. From Fig. 5, we can furthermore observe that the PDFs are smoother as PI has more freedom with a constraint only on M_{ZF}. Thus, weights are redistributed more evenly over samples, and the resulting PDFs and CDFs after resampling are smoother. Finally, the correlation between M_{OE} and M_P also changes from 0.95 to somewhat below 0.94. That is not a big difference but shows that for more complex problems and/or more constraints PI alters the dependency between variables.

III. Application

The mission analysis method that is used to showcase PI in an aircraft technology evaluation and selection setting is discussed in

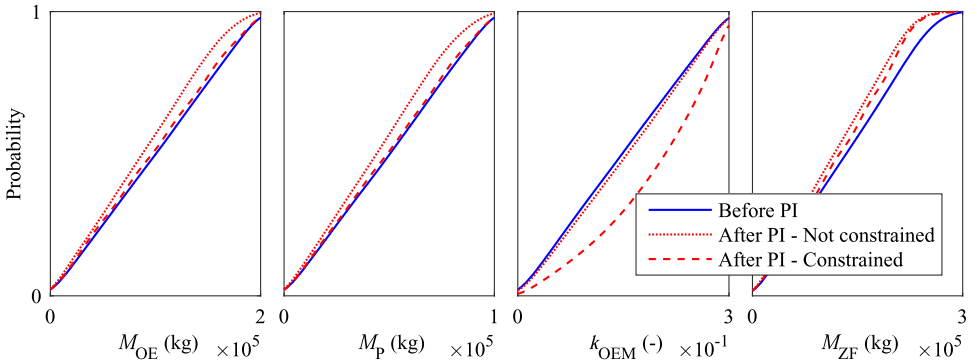

Fig. 4 Cumulative distribution functions before and after PI.

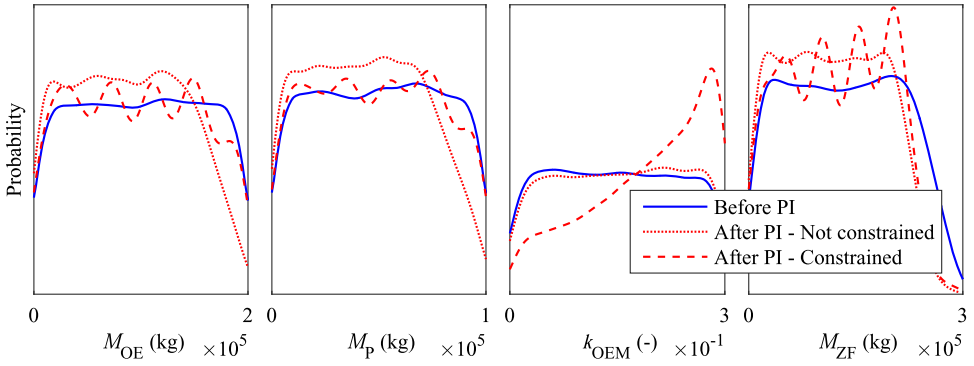

Fig. 5 Probability density functions before and after PI.

this section. Consecutively, the inputs to this method are presented together with their probability distributions. Finally, the technologies that are under investigation and how they are modeled are detailed.

A. Mission Analysis Method

A simplified mission is simulated and starts at zero altitude and takeoff speed V_0 as well as an assumed fuel mass M_{F_0}, then climbs to cruise altitude h_{cr} and accelerates to cruise speed V_{cr}. The range flown during climb is now computed, but the range flown during cruise or descent is unknown. Therefore, descent is first computed backward (because the final point is known), such that the range flown during descent becomes known. The final point is at zero altitude and V_0, with zero fuel mass. When both climb and descent have been computed, the range flown in cruise is known and is assumed to be flown at constant altitude and speed. In its entirety, the mission analysis function can be described as

$$(M_F, R) = \mathrm{MA}(M_{F_0}, \boldsymbol{x}) \tag{14}$$

where $\boldsymbol{x}$ is a sample that holds a value for each of the input variables (presented in Sec. III.C) other than the initial fuel mass estimate M_{F_0}. The function MA returns both the consumed fuel mass M_F and range flown R. At any point, the assumed initial fuel mass M_{F_0} may turn out to be insufficient ($R < R_{req}$). Conversely, when the entire mission is flown ($R \geq R_{req}$), residual fuel may remain ($M_{F_0} > M_F$). Therefore, an outer iteration aims to find the particular value for M_{F_0} that is adequate to fly the specified range. That value gives the quantity of interest: fuel burn. The outer iteration is implemented using a minimization-within-bounds algorithm, which minimizes the error in fuel mass and range

$$\epsilon(M_{F_0}, R) = \mathrm{abs}(M_F/M_{F_0} - 1) + \mathrm{abs}(R/R_{req} - 1) \tag{15}$$

by adjusting M_{F_0} as

$$\mathrm{argmin}_{M_{F_0}} \epsilon(M_{F_0}, R) \tag{16}$$

For a fair comparison between aircraft and technologies, the block fuel burn M_F is not a suitable metric; heavier, long-range aircraft consume more fuel, even if they are more efficient than lighter, regional aircraft. Therefore, the payload-range energy efficiency (PREE) [36] is used. However, this metric is based on energy consumption and is defined as $R \cdot M_P \cdot E^{-1}$, where E is the energy consumed during an entire mission. Thus, this metric measures efficiency, which has to be increased. However, in this paper, the QOI has to be minimized, so the inverse of PREE is taken, and the energy consumption is replaced with block fuel burn, resulting in the PRE^{-1} metric. Thus, PRE^{-1} is the block fuel normalized with range and payload mass, in other words, $M_F \cdot (R \cdot M_P)^{-1}$.

The mission analysis computations assume steady flight. Furthermore, thrust during climb is assumed to be a constant fraction T_{cl} of the maximum thrust T_{max}, and during descent, a constant fraction T_{des} is assumed. Finally, a speed–altitude profile is assumed, given by

$$\frac{dV}{dh} = \frac{V_{cr} - V}{h_{cr} - h} \tag{17}$$

during climb and the negative of that during descent. This profile is not necessarily realistic but suffices for the current application, where only differences between different inputs matter.

B. Cost Computation

No detailed cost module is present in the current method. However, a cost metric is included in order to have a second objective that conflicts with PRE^{-1}. It can be interpreted as any kind of cost, for example, recurring cost, nonrecurring cost, and/or direct operating cost. The baseline aircraft is considered to have a normally distributed cost with mean $\mu = 1$ and standard variance $\sigma = 0.05$. Each technology adds some cost measure to this baseline distribution. Depending on these technology costs, the final cost for each portfolio may be smaller or larger than the baseline.

C. Input Variables

In this section, the distributions and dependencies for the input variables of the mission analysis are set up. Most researchers that include uncertainty do not focus on how the probability distributions should be obtained [37]. Characterizing those distributions for subsequent analysis is a challenging task, which often relies on expert judgment. Even though that is labor intensive and subjective, there is no viable alternative presently, and Fig. 1 shows it is adopted here as well. However, an alleviating remark is made by Cook and Jarrett [2], who address the question of "How important is the choice of how to represent input uncertainties mathematically in robust airfoil optimization?" as an example of what effect the specific choice of uncertainty distribution has on the outcome. The answer is that the difference between probabilistic results (i.e., different probability distributions) is insignificant, albeit their difference with the deterministic cases was large. In one case, even, the probabilistic optimum was Pareto dominant with respect to the deterministic one. Thus, including uncertainty is important. However, not just any distribution works, and especially dependencies between variables have to be taken into account.

Table 1 shows all the inputs to the mission analysis method [x in Eq. (14)]. Most of these variables are assigned a uniform distribution, to reflect a wide range of aircraft and missions. Thrust-specific energy consumption (TSEC) and specific energy (SE) are assigned a triangle distribution because their mean is derived from data and the triangle distribution allocates more probability mass around this value, while having finite bounds. Finally, some variables are represented with scalar values because they are aircraft independent and this way the design space is reduced.

D. Technologies

In the ensuing case studies, three technologies are used to showcase PI. The first, which is also investigated individually, is a second-degree-of-freedom (2nd DOF) flap. This flap has two actuators that independently control the extension and rotation of the flap, to allow it to provide maneuver load alleviation (MLA) and camber optimization.

Table 1 Input variables specification

Variable name	Symbol	Distribution	Dependency
Cruise altitude, m	h_{cr}	Uniform	Independent
Cruise speed, m/s	V_{cr}	Uniform	Independent
Range, m	R	Uniform	Independent
Payload mass, kg	M_P	Uniform	Correlated with M_{OE}
Takeoff and landing speed, m/s	V_0	Scalar (45)	Independent
Wing loading, N/m^2	W/S	Uniform	Determines wing area S
Wing aspect ratio	A	Uniform	Determines wing span b
Thrust-to-weight ratio	T/W	Uniform	Independent
Climb throttle	T_{cl}	Scalar (0.85)	Independent
Descent throttle	T_{des}	Scalar (0.05)	Independent
Operating empty mass, kg	M_{OE}	Uniform	Correlated with M_P
Clean zero-lift drag	C_{D_0}	Uniform (0.01, 0.02)	Independent
Thrust specific energy consumption, J/(N · s)	TSEC	Triangle (600,750,900)	Independent
Oswald factor	e	Scalar (0.7)	Independent
Fuel specific energy, J/kg	SE	Triangle ($45 \cdot 10^6$, $46 \cdot 10^6$, $47 \cdot 10^6$)	Independent
Time step, s	Δt	Scalar (30)	N/A

Therefore, it is modeled using three variables: $\delta_{f,\max}$, $k_{M,\text{flap}}$, and k_e. The first is the maximum deflection of the flap, either upward or downward. The second is the effect of maneuver load alleviation on the structural weight, measured as a percentage of M_{OE}. (A lighter wing box can be designed when MLA effectively reduces the maximum load factor [38–40]). The third models the effect of camber optimization [41] by being added to the aircraft's Oswald factor. Its distribution is estimated from Ref. [42]. Both $\delta_{f,\max}$ and k_e influence the aircraft lift–drag polar as follows:

$$C_D = C_{D_0} + 5 \times 10^{-5} \cdot \delta_f + \frac{C_L^2}{\pi A(e + k_e \cdot \delta_f)} \tag{18}$$

In this equation, δ_f is measured in degrees; thus, k_e is measured in $1/\deg$. The effect of δ_f for a fixed value of k_e is shown in Fig. 6. Essentially, an increasing flap deflection shifts the polar to the right, while reducing induced drag, thus making the polar less steep.

In the mission analysis, the L/D ratio is optimized at each point during the mission by varying the flap deflection angle δ_f in the range $[0, \delta_{f,\max}]$. During climb and cruise, L/D is maximized by the flap for a given C_L, as

$$\text{argmin}_{\delta_f \in [0,\delta_{f,\max}]} C_D \tag{19}$$

while during descent, it is minimized for the steepest path:

$$\text{argmax}_{\delta_f \in [0,\delta_{f,\max}]} C_D \tag{20}$$

All three technology-defining variables and their distributions are shown in Table 2. It also shows the dependency imposed on $\delta_{f,\max}$ and $k_{M,\text{flap}}$. This follows the insight that the higher the flap deflection is, the more maneuver load alleviation is achieved. However, for increasingly large deflections, more uncertainty is present in the mass reduction that is achieved. To model this effect, a rotated Clayton copula is used, with $\alpha = 8$, computed from a Kendall's τ of 0.8. This copula is shown in Fig. 7 in the leftmost plot.

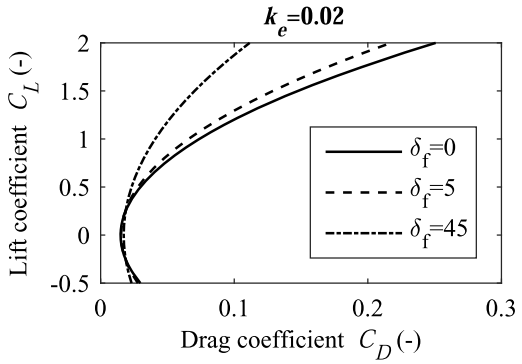

Fig. 6 Lift-to-drag polars affected by flap deflection for notional 2nd DOF flap.

The second technology is a rotating winglet downer, which is a rotating element originating from the winglet's base and pointing downward. It deflects to alter the lift distribution around the wing tip and consequently offers maneuver and gust load alleviation [43]. It is modeled using a single variable $k_{M,\text{downer}}$ which also is a percentage of M_{OE}. When both this and the flap technology are present, the combined effect is not simply the multiplication of both mass reduction factors. Instead, $k_{M,\text{downer}}$ is expected to be closer to 1 when $k_{M,\text{flap}}$ is low and vice versa. Therefore, a negative dependency has to be imposed on these two variables. Furthermore, it is assumed that it is impossible to reach either variable's minimum when both are present. For these reasons, this dependency is modeled using a Clayton copula with $\alpha = -0.75$, computed from a Kendall's τ of -0.3. The resulting joint distribution is shown in the center plot in Fig. 7. The rightmost plot shows the dependency between $\delta_{f,\max}$ and $k_{M,\text{downer}}$, resulting from the two previously mentioned dependencies. The combined effect of the two technologies $k_{\text{tot}} = k_{M,\text{flap}} \cdot k_{M,\text{downer}}$. Thus, drawing from the center distribution in Fig. 7, and applying this multiplication to obtain k_{tot}, the distributions in Fig. 8 are obtained. Observe that k_{tot} (in some cases) achieves more reduction than either $k_{M,\text{flap}}$ or $k_{M,\text{downer}}$ alone, although not as much as the multiplication of their minima (i.e., $0.8 \cdot 0.9 = 0.72$). Instead, the minimum of k_{tot} lies around 0.77.

In summary, the 2nd DOF flap and winglet downer offer a mass reduction as follows:

Table 2 Technology variables specification

Variable name	Symbol	Distribution	Dependency
Maximum flap deflection, deg	$\delta_{f,\max}$	Uniform (0, 30)	Correlated with $k_{M,\text{flap}}$
Flap mass impact	$k_{M,\text{flap}}$	Uniform (0.8, 1)	Correlated with $\delta_{f,\max}$
Flap Oswald efficiency, deg^{-1}	k_e	Uniform (0.012, 0.036)	Independent
Downer mass impact	$k_{M,\text{downer}}$	Uniform (0.9, 1)	Independent
Engine TSEC impact	k_{TSEC}	Uniform (0.5, 0.99)	Independent

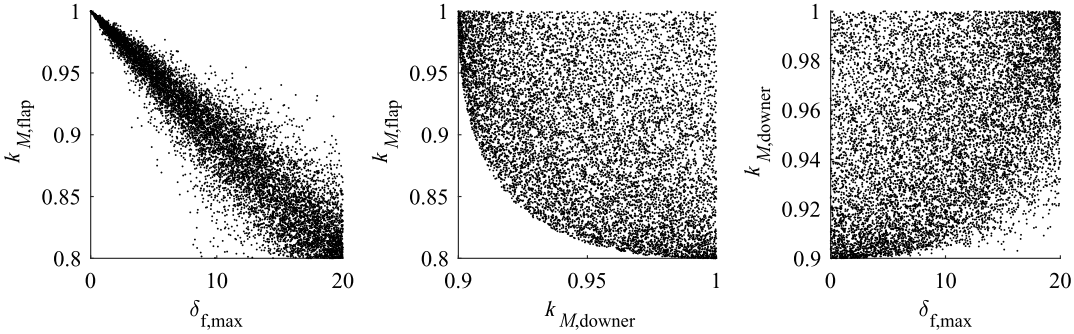

Fig. 7 Dependencies between the three variables $\delta_{f,\max}$, $k_{M,\text{flap}}$, and $k_{M,\text{downer}}$ characterizing the 2nd DOF flap and winglet downer.

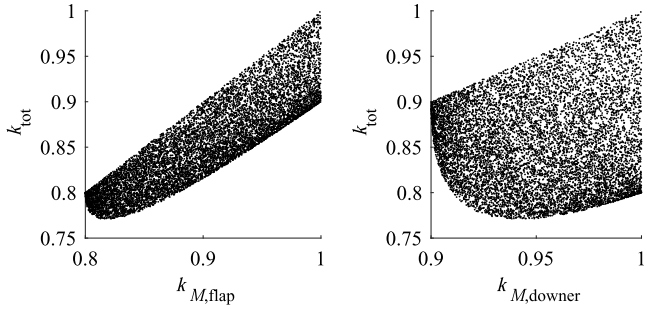

Fig. 8 Dependencies between the 2nd DOF flap and winglet downer mass impacts and their combined effect k_{tot}. These plots are the result of combining those from Fig. 7 with the fact that $k_{\text{tot}} = k_{M,\text{flap}} \cdot k_{M,\text{downer}}$.

$$\Delta M = \begin{cases} k_{M,\text{flap}} \cdot M_{\text{OE}} & \text{if only flap} \\ k_{M,\text{downer}} \cdot M_{\text{OE}} & \text{if only winglet downer} \\ k_{\text{tot}} \cdot M_{\text{OE}} & \text{if both flap and winglet downer} \end{cases} \tag{21}$$

The third technology is some engine improvement (e.g., better turbine blade cooling, to increase turbine inlet temperature) that improves the TSEC of the aircraft. It is modeled with k_{TSEC}, also shown in Table 2, and is independent from the other two technologies.

For each technology, the cost is determined through a deterministic function of its characterizing variables or a random distribution or both. These functions are crafted to give a spread in total cost, such that a tradeoff between fuel burn and cost results. The cost impacts and total cost are computed as

$$\begin{aligned} C_0 &= \mathcal{N}(\mu = 1, \sigma = 0.05), \\ \Delta_{C,\text{flap}} &= \delta_{f,\max}/100, \\ \Delta_{C,\text{downer}} &= \mathcal{N}(\mu = 0.1, \sigma = 0.02), \\ \Delta_{C,\text{engine}} &= (1 - k_{\text{TSEC}})/5 + \mathcal{N}(\mu = 0, \sigma = 0.05), \quad \text{and} \\ C &= C_0 + \Delta_{C,\text{flap}} + \Delta_{C,\text{downer}} + \Delta_{C,\text{engine}} \end{aligned}$$

Here, $\mathcal{N}$ is a Gaussian distribution. Obviously, when a certain technology is not included, its cost is not added to the total cost.

E. Verification

The mission analysis method is examined in this section. First, the response of the model for a single aircraft and design point is investigated. Second, a global sensitivity analysis is conducted to study the response of the model, in order to support the conclusions in the subsequent test cases.

1. *Mission Analysis for Single Aircraft*

An Airbus A320 aircraft is notionally represented using the available input variables to study the various responses of the mission analysis model. In Fig. 9, the results of the mission analysis for this input are depicted. It is important to observe that, during climb and descent, the flap deflection δ_f is always at its maximum value $\delta_{f,\max}$ and otherwise zero. This is a deficit in the lift–drag polar model (explained in Sec. III.D); the polars for zero and maximum deflection cross in the very small region. Therefore, C_L is either above or below it, and the one with least drag in that region is chosen. During cruise, the C_L is within the transition range from zero to maximum deflection, and hence a gradual variation results. In a more realistic setting, adverse effects of flap deflection, such as separation, have to be modeled and will influence (i.e., limit) the flap deflection.

For the rest, the mission analysis computation behaves as expected, with realistic values for all parameters (even though C_L almost reaches 5 as a result of the assumption for dV/dh). This conclusion is not enough to support the case studies, though. To make conclusions about the effect of technologies, or do probabilistic inversion, the sensitivity of this analysis method has to be measured against the variables of interest.

2. *Global Sensitivity Analysis*

A sensitivity analysis (SA) is performed to study the response of the model with respect to changes in its input. This information supports the conclusions that are drawn from PI in the various test cases in the following sections. For example, when PI modifies the distribution of one variable significantly, while another is hardly affected, there are two possible explanations. The first explanation is that the former variable simply is more effective in reaching the goal that PI is set out to achieve, even though both variables have a measurable influence on that goal. In contrast, the second explanation is that the model used to generate the samples is insensitive to the latter variable; hence, PI only changes the former because it is the only variable explaining the variance in the goal.

Commonly, sensitivity analyses are conducted around a particular point in the design space (local SA). This type of analysis is justified when investigating a specific design, rather than a wide range. For what is currently of interest, a wide range of aircraft designs, global sensitivity analysis is a more appropriate tool. It is performed as follows.

There are two quantities of interest in this study: PRE^{-1} and cost. A sensitivity analysis is only conducted on the first, as for cost there are linear equations with the technology impact variables, and as such, we already know the sensitivities are nonzero. For PRE^{-1}, the product moment correlation, rank correlation, and correlation ratio are computed. The first is a measure of the strength of a linear association of two variables, on a scale from -1 to 1. Rank correlation is similar, except that it measures how well the two variables follow a monotonic function. Finally, correlation ratio provides a sense of the extent to which a variable's variance explains another variable's variance. It is measured on a scale from 0 to 1. Therefore, unlike the other two, the correlation ratio does not give a direction of influence.

To study the sensitivity of the mission analysis method, the sensitivity of PRE^{-1} is computed with respect to the variables in Table 1, without any of the technologies. The results are shown in Fig. 10, where the sensitivities are sorted from highest to lowest, in an absolute sense. It becomes clear that C_{D_0} and TSEC are most influential. That

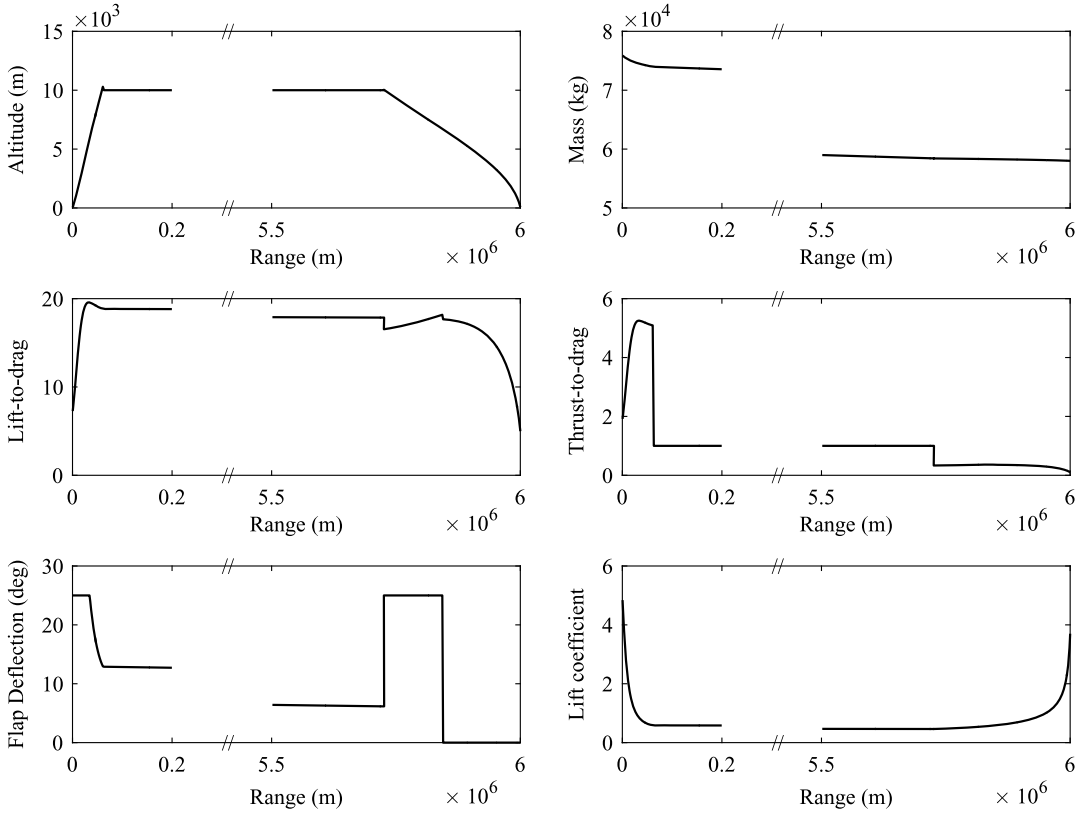

Fig. 9 Mission analysis results showing variation of altitude, mass, L/D, T/D, δ_f, and C_L with range flown.

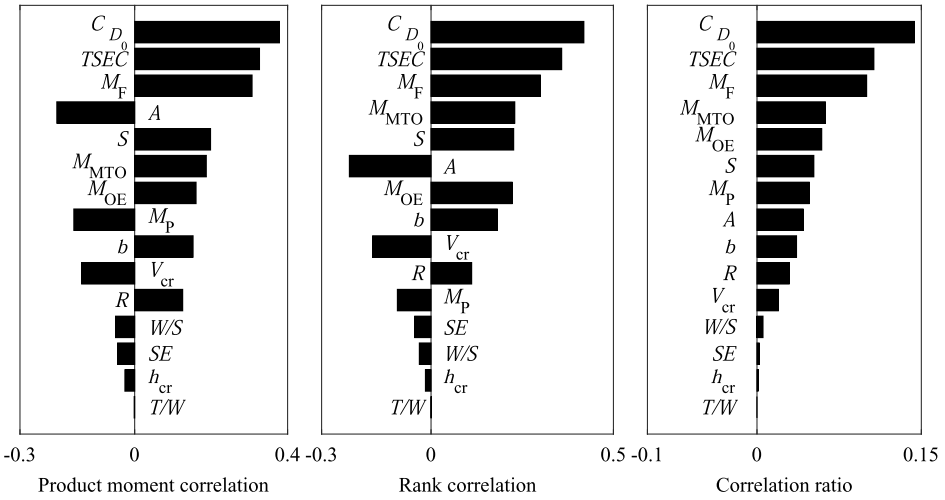

Fig. 10 Sensitivity measures of PRE^{-1} with respect to the variables in Table 1.

M_F is in third place is surprising, as it is the variable that is measured, except that it is not yet normalized with R and M_P. Therefore, we may conclude there are other variables explaining the variance in M_F other than R and M_P. Each variable's direction of influence agrees with engineering instinct. Therefore, it is concluded that the mission analysis behaves as expected and shows good sensitivity with respect to the input variables.

A similar sensitivity analysis only includes the technology impact variables from Table 2, in the dataset where all three technologies are included. This gives the sensitivities in Fig. 11, which are as expected, except for the sign of $k_{M,\mathrm{downer}}$. Thus, $k_{M,\mathrm{downer}}$ has a negative correlation, while the same sign as $k_{M,\mathrm{flap}}$ is expected because the two are effectively the same. To check correct operation of the mission analysis, the sensitivity study was repeated for when only the downer technology is included, which indeed shows the correct sign for $k_{M,\mathrm{downer}}$.

The sensitivity analyses show that the QOI PRE^{-1} is sufficiently affected by the technologies and the other input variables. Thus, PI should have little to no bias toward certain variables as a result from model insensitivities.

IV. PI for Single Technology

Even though PI shows most potential over conventional approaches in the case of multiple objectives, this section starts with a single technology and a single objective. This is because relevant design queries can be answered using this approach and it is a suitable stepping stone for the more involved multi-objective technology prioritization in the next section.

For a single technology, queries regarding what technology impacts are required for a given reduction of $x\%$ in QOI y might be of interest. Such questions are easily posed in PI, by constraining y

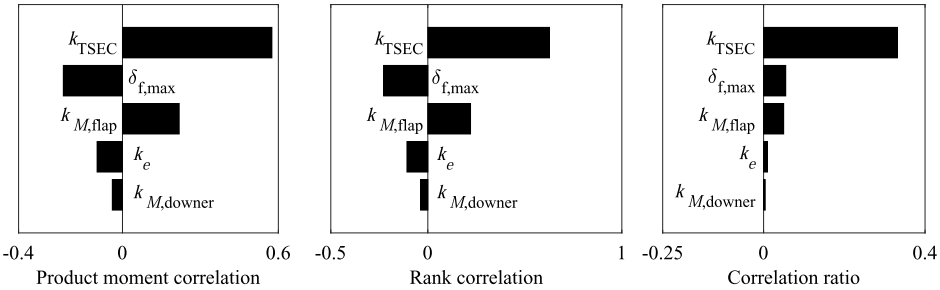

Fig. 11 Sensitivity measures of PRE^{-1} with respect to the variables in Table 2.

and observing the changes in the technology parameters. This section shows how PI enables answering such queries, specifically for the 2nd DOF flap technology, introduced in Sec. III.D. Next, the use of PI to study technology maturation is presented.

A. Inverse Query: Technology State Required for Prescribed Benefit

Studying what values the technology variables should attain for a given change in one or more QOI is done using PI. The technology variables ($\delta_{f,\max}$, $k_{M,\text{flap}}$, and k_e) are allowed to move freely, while the other input variables are held fixed. A distribution is specified on PRE^{-1} that reduces it by some amount for each percentile. The same bivariate distributions between M_{OE} and M_P and between $\delta_{f,\max}$ and $k_{M,\text{flap}}$ are constrained as well. The results in Fig. 12 follow intuition: the median PRE^{-1} has reduced by 15%, as specified. Correspondingly, the maximum flap deflection has shifted to the right, as has k_e. The deflection's median went from 17 to 23 deg: an increase of 35%. At the same time, its variance reduced by 40%. The median of k_e increased by 13%, from 0.0245 to 0.0277, with a variance reduction of 14%. As expected, the distribution for $k_{M,\text{flap}}$ shifts to the left, as a result of the increase in $\delta_{f,\max}$. Likewise, cost has increased. The fact that variance is reduced is attributed to the constraint on PRE^{-1} having lower variance than the original distribution. The corresponding PDFs are shown in Fig. 13.

While traditional approaches may be able to obtain statistical measures such as the mean and variance, PI provides a full distribution for the variables of interest, to fulfill certain goals. With pure forward propagation approaches, that is only attainable through iteratively updating the input distributions and propagating these through the models.

B. Alternative Use of PI: Investigating Technology Maturation

Technology maturation can be modeled by updating the input distributions of the technology variables [9]. However, running the analysis models again to generate thousands of samples is Time-

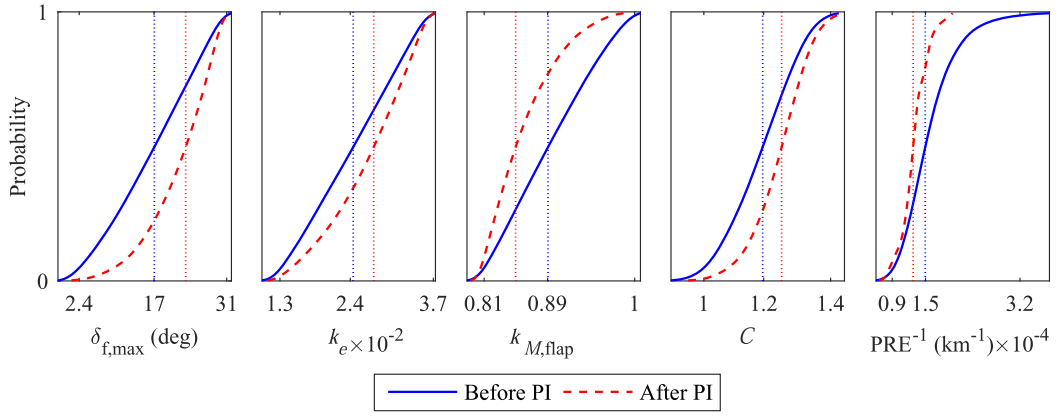

Fig. 12 Cumulative distribution functions for 2nd DOF flap technology with original (blue) distributions and for a 15% PRE^{-1} reduction (red, dashed). The vertical lines represent the medians of the two curves.

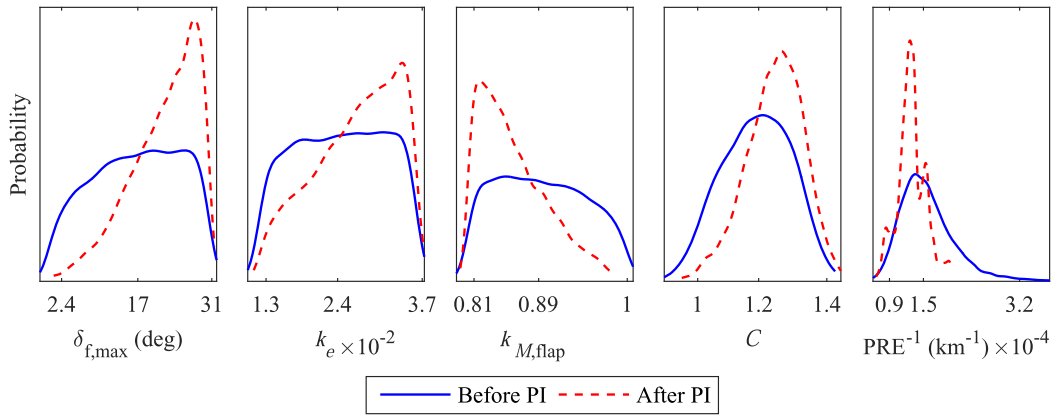

Fig. 13 Probability density functions for 2nd DOF flap technology with original (blue) distributions and for a 15% PRE^{-1} reduction (red, dashed).

Consuming. With PI, the same result can be achieved, provided the updated distributions are within the range of the original distributions. Where the traditional approach requires propagating N samples through the analysis models twice, the PI approach only requires a single forward propagation pass.

For the 2nd DOF flap, the distributions of $\delta_{f,\max}$ and k_e are adjusted. The flap deflection now follows a triangle distribution from 12 to 30 deg, with its peak at 20 deg. The Oswald efficiency impact is also provided with a triangle distribution from 0.02 to 0.036, with a peak at 0.03. In both cases, the range of values is less than the original, uniform distribution. Furthermore, more probability mass is located around the expected values: 20 deg and 0.03, respectively. Finally, because the dependency of $\delta_{f,\max}$ with $k_{M,\text{flap}}$ (recall Fig. 7, leftmost plot) has to be maintained, $k_{M,\text{flap}}$ is given an updated uniform distribution corresponding with the range of values for $\delta_{f,\max}$.

To set up the PI problem, the new distributions on $\delta_{f,\max}$, $k_{M,\text{flap}}$, and k_e are imposed as constraints. Furthermore, the joint distributions of M_{OE} and M_P and $\delta_{f,\max}$ and $k_{M,\text{flap}}$ are constrained, as are all input variables that should remain unchanged. As Fig. 14 illustrates, the results from forward uncertainty propagation and PI are virtually identical. The only noticeable difference lies in the lower percentiles of the $\delta_{f,\max}$ and k_e CDFs. This is simply a result of not constraining PI at even lower percentiles. Regardless, the updated CDFs of cost and fuel burn show the same effect. The corresponding PDFs are shown in Fig. 15.

Comparing the median (50 percentiles) of these two QOI, a 3% cost increase and an 8% fuel burn (per passenger kilometer) are measured. Furthermore, the variance of cost decreased by 51% and for fuel burn decreased by 20%. For the forward propagation, a 3% cost increase and 9% fuel burn reduction are computed. Thus, as Fig. 14 already shows, the two approaches give an almost identical result. However, the variance reduction using forward propagation is more pronounced: 59% as opposed to 30% for PI. This is attributed to the discrepancy in lower percentiles, mentioned earlier.

It should be reemphasized that the approach employed in this section only works when the matured technology distributions lie within the originally sampled distributions. That is because PI can only reweight existing samples; therefore, samples outside the original set cannot be created or inferred. However, the modified distributions should not necessarily have the same shape as the original ones.

V. PI for Technology and Portfolio Selection

In contrast to the previous test case, all three technologies are investigated here. Rather than inspecting a single technology, the differences between all portfolios resulting from multiple technologies can be studied. A portfolio is simply a set of included technologies, represented as a vector, where each entry corresponds to a technology and is 1 when it is included and 0 otherwise. Thus, there are 2^n portfolios in total, where n is the amount of technologies. Additionally, the strength of PI is showcased: dealing with multiple objectives. Rather than only having PRE^{-1} as goal, cost is a goal as well.

For the present paper, the resulting CDFs for PRE^{-1} and cost are depicted in Fig. 16 for each portfolio. It becomes clear that in terms of PRE^{-1} portfolios with technology 3 (third digit is 1) perform better

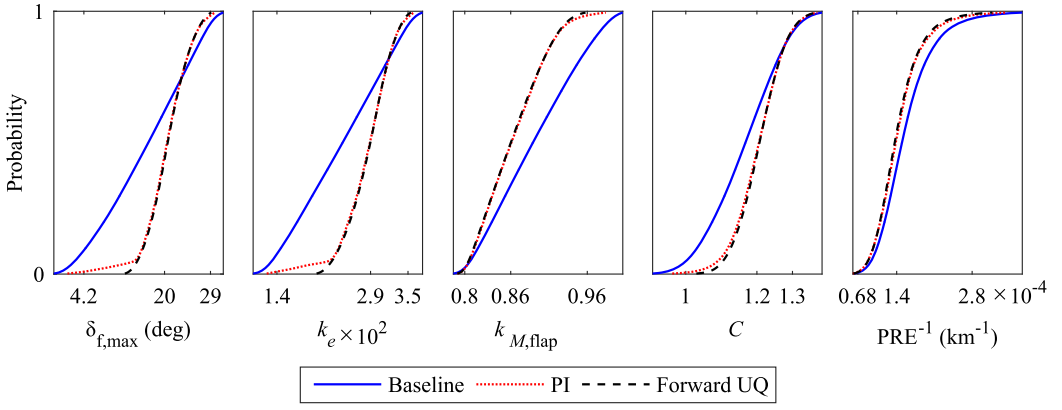

Fig. 14 Cumulative distribution functions for 2nd DOF flap technology with original (blue) distributions for $\delta_{f,\max}$ and k_e. After technology maturation, the CDFs are computed with PI (red, dotted) and with forward uncertainty propagation (black, dashed). The horizontal axis tick marks are the 5th, 50th, and 99th percentiles, in that order.

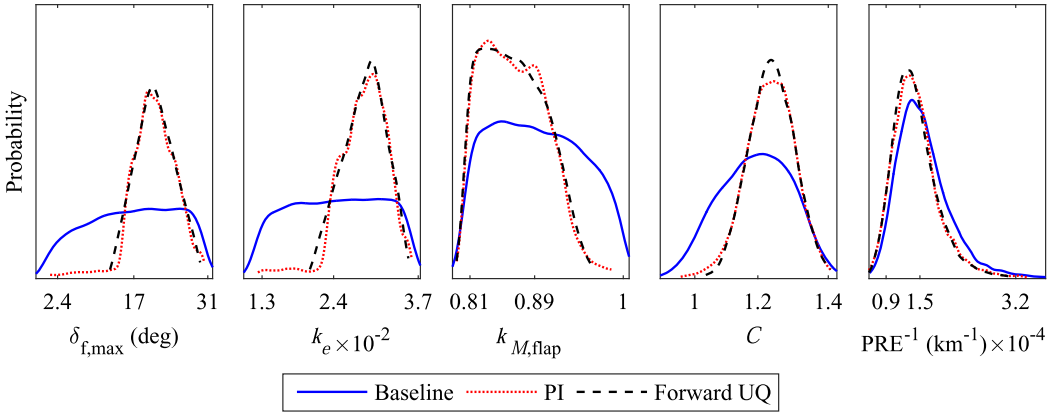

Fig. 15 Probability density functions for 2nd DOF flap technology with original (blue) distributions for $\delta_{f,\max}$ and k_e. After technology maturation, the PDFs are computed with PI (red, dotted) and with forward uncertainty propagation (black, dashed). The horizontal axis tick marks are the 5th, 50th, and 99th percentiles, in that order.

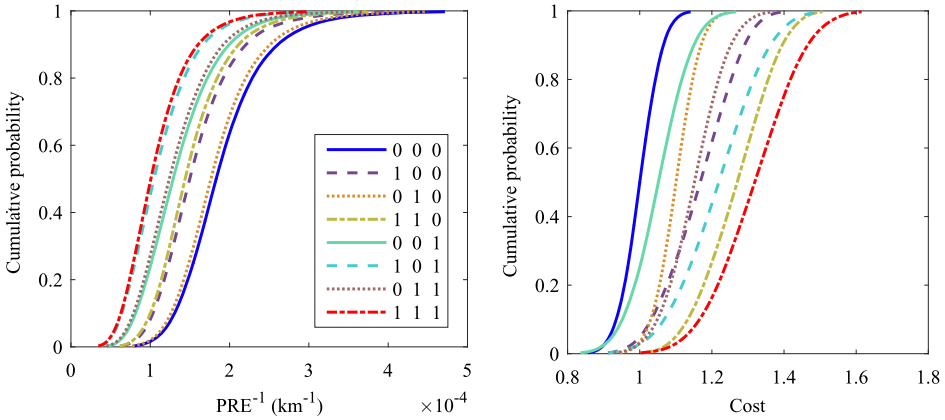

Fig. 16 CDFs of PRE^{-1} and cost for multiple portfolios.

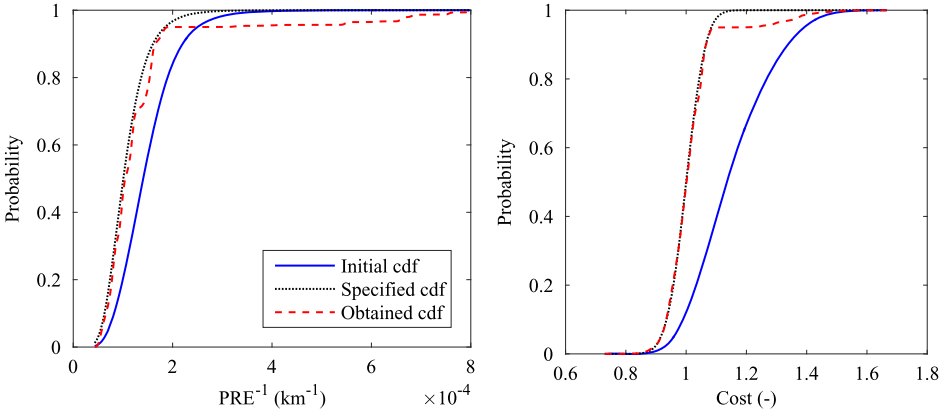

Fig. 17 Comparison of initial, specified, and obtained CDFs for PRE^{-1} and cost, using PI with a multi-objective constraint.

than the ones without it. The portfolio with all technologies included (111) performs best (it is farthest parent to the left), which can be expected because none of the technologies was defined to have negative effects on fuel burn. Second best is the portfolio 101, which combines technologies 1 and 3. The worst portfolio is the baseline (000). Thus, already from that figure, if PRE^{-1} is the only requirement, an ordering and selection of the portfolios can be made. However, with cost (see Fig. 16) there is an opposite trend: the more technologies included in a portfolio, the higher the cost. So here, portfolio 111 performs worst (farthest to the right), and the baseline performs best. There is a different spread in the CDFs for cost, so when constraints are imposed on fuel burn and cost simultaneously, there is no direct way of telling which portfolio performs best. Consequently, concluding which technologies are most promising is not directly observable anymore.

A. Technology Prioritization with Multiple Objectives

The strength of PI lies in its ability to deal with multiple, possibly conflicting, goals. However, the way in which the constraints are set up affects what query is posed to PI and provides different results. For this study, the multi-objective constraint should guide the obtained samples toward the Pareto front of fuel burn and cost. That way, we learn which distribution of portfolios is closest to the Pareto front. Thus, this distribution gives the best tradeoff between the objectives.

As pointed out by Binois et al. [44], there is an analogy between Pareto fronts and the level curves of a copula.[¶] The zero-level curve of a copula corresponds to the Pareto front of the multivariate distribution that the copula represents. To redistribute the samples closer to the Pareto front, new margins have to be specified for the objective variables. However, specifying new margins alters the copula as well because the copula is computed from the margins as explained in Sec. II.B. This is not desired because the dependency captured by this copula is a result of the model and should be treated as a physical fact. Therefore, the copula itself has to be constrained. Such a constraint can be implemented as the sum on the margins in copula space (the variables U_i in Sec. II.B). Furthermore, this sum has to be updated every iteration of IPF, because IPF updates the margins on every iteration. The quantiles of the sum are kept constant, though, as that will fix the copula.

To demonstrate this procedure, both a PRE^{-1} target and cost target are specified. The PRE^{-1} target are the percentiles of the 111 portfolio CDF. The cost target are the percentiles of the baseline CDF (000 portfolio). These are the extremes of the sample space, as shown in Fig. 16.

IPF performs exactly as intended with the constraints on the margins of PRE^{-1} and cost and a constraint on their copula. The margins are satisfied, and the copula is hardly affected, as Figs. 17 and 18 show. The obtained CDFs (see Fig. 17) match the specified CDF perfectly at the constrained percentiles. From the 95th percentile on, the CDFs do not match, which is to be expected due to the discrete nature of PI. The level curves after PI (red, dashed lines in Fig. 18) are in the same location as the initial ones (black, solid), although they wiggle around these somewhat. That is a result of the discrete constraints, in combination with only a relatively small subset of samples receiving most of the weight after PI. It can be shown that the copulas before and after PI are very similar by computing the Pearson correlation coefficients. In copula space, the correlation initially was -0.3956, while after PI, it is -0.3871, a difference of only 2%. When the constraints are less stringent, even better agreement is obtained.

The previous results may be further explained by focusing on what happens in sample space. The bivariate distribution of PRE^{-1} and

[¶]A level curve of a multivariate function f is formed by all solutions, $\{x\}$ where $f(x) = c$, in other words, where the function f has a given value c.

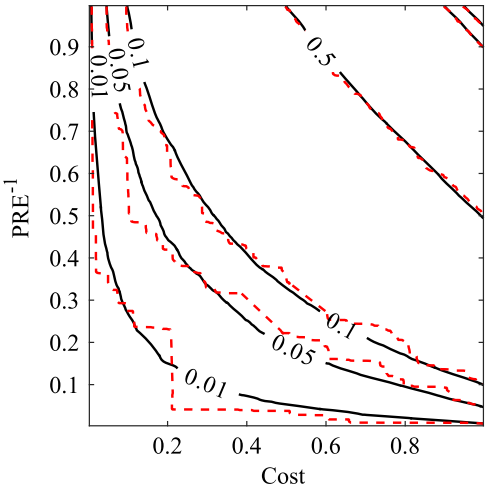

Fig. 18 Comparison of copula level curves of PRE^{-1} and cost, before and after PI with a multi-objective constraint.

cost is shown in Fig. 19 on the left, with the level curves of the two-dimensional CDF. These level curves are similar to the ones of the copula in Fig. 18. The center plot in Fig. 19 shows how these level curves shift after PI, due to the updated margins. All except the 95th percentile lines have moved toward the lower left corner; thus, reducing both PRE^{-1} and cost. Furthermore, the lines are closer together; the updated 90th percentile is on the initial 50th percentile, and the updated 50th percentile is on the initial 10th percentile. Therefore, the reweighted samples are squeezed into a region closer to the Pareto front. In the right plot in Fig. 19, the updated bivariate distribution is shown. Clearly, only few samples remain after reweighting, which explains the nonsmooth results observed in Figs. 17 and 18.

As a result of the multi-objective constraint, the portfolio and technologies frequencies in the resampled set have changed as well. This is the result we are looking for, and it is shown in Fig. 20. These frequencies are simply the amount of samples that contain a certain portfolio or technology, with respect to the total number of samples. The change with respect to the initial sample set is shown because not every portfolio or technology is equally represented in the initial set (i.e., the initial frequencies are not uniformly distributed).

Mainly portfolios 001 and 101 have received more weight, leaving technology 3 to see an increase in frequency as well. Technology 1 only suffers a small decrease, while technology 2 clearly does not satisfy the imposed constraint. When selecting a portfolio, it should be 101 as it is most featured in the resampled set, while technology 3 should receive most development resources.

B. Technology Prioritization Using Grouped Sample Reweighting

The previous results were obtained with individual sample reweighting. However, the grouped sample reweighting approach should be considered as well. The difference between the individual sample reweighting and grouped sample reweighting approaches can be considered by observing Fig. 21. On the left, the bivariate distributions for each portfolio of PRE^{-1} and cost are shown. This is the entire sample space that PI works with. When constraints are imposed and PI is performed, only a subset of these samples remains. The center plot in Fig. 21 shows this for individual sample reweighting. It becomes clear that a lot of samples receive near-zero weight to move the overall bivariate distribution to the bottom-left corner. Moreover, the shapes of the colored bivariate distributions (each corresponding

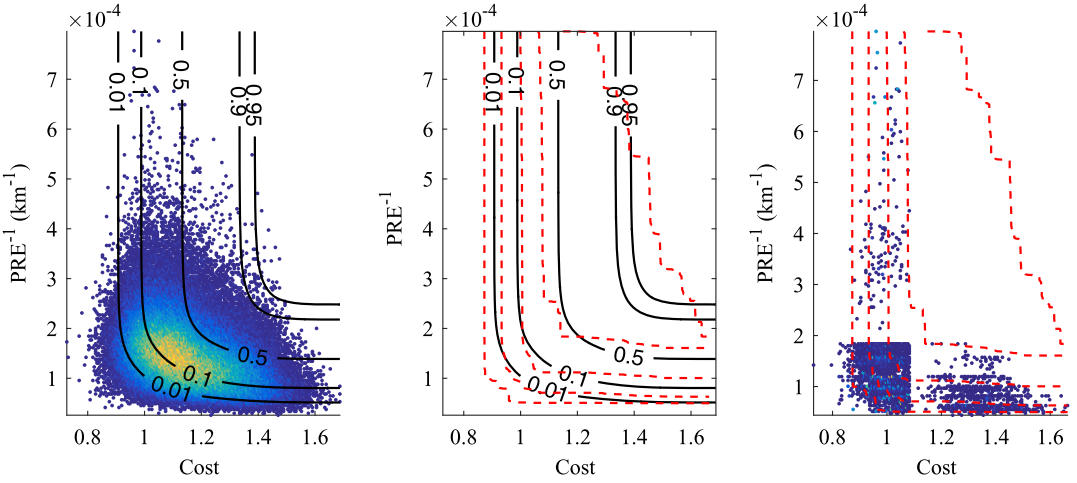

Fig. 19 Comparison of bivariate distribution of PRE^{-1} and cost, before and after PI with a multi-objective constraint.

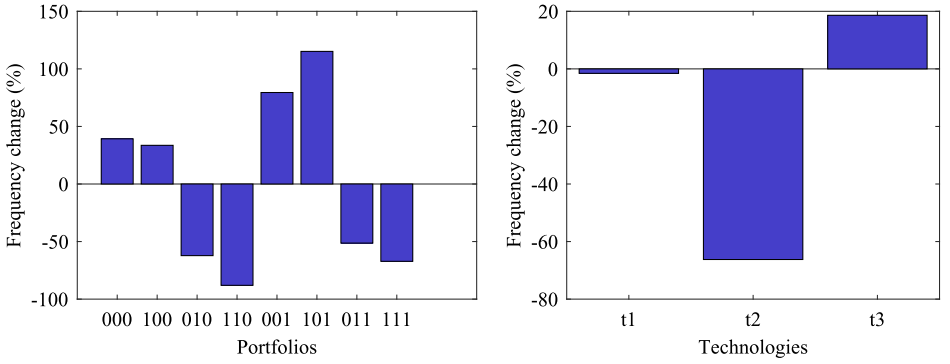

Fig. 20 Portfolio and technology frequency shifts using PI with a multi-objective constraint.

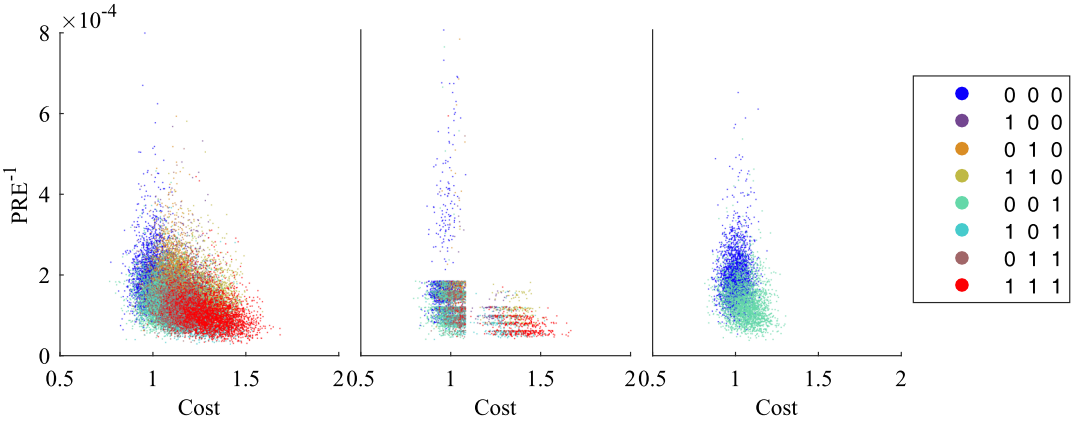

Fig. 21 Bivariate distribution of PRE^{-1} and cost: left, when no constraints are imposed, in other words, before PI; center, when the margins of cost and PRE^{-1} are constrained and with individual sample reweighting; and, right, when the margins of cost and PRE^{-1} are constrained with grouped sample reweighting.

to one portfolio) change as a result. That means the underlying distributions of the variables for each portfolio have shifted (similar to the single technology case in Sec. IV). This effect is desired when one is interested in which portfolio has the most potential to satisfy the imposed goals, when its initial distributions are not set in stone. On the contrary, when the distributions for each portfolio are well defined and should not be allowed to change, the grouped approach is needed. Its result is shown in the rightmost plot, where the bivariate distributions of two portfolios remain, while all others have zero probability mass. The shapes of these bivariates have not changed with respect to the leftmost plot.

The grouped sample reweighting has to be performed using PARFUM, because it is very likely the IPF will not converge. That is because there is too little room to play when only the weight of entire portfolios may be changed. For this reason, PARFUM will also not achieve the specified margins exactly but will get as close as possible. The distance between the initial, obtained, and specified CDFs can be computed using the metric presented by Cook and Jarrett [45]. This shows for PRE^{-1} that the initial distance to the specified constraint is $6.1 \cdot 10^{-5}$, and after PI it is $5.1 \cdot 10^{-5}$. Similarly, for cost, the initial distance is 0.24, and after PI, it is 0.055. Because the distance after PI is not zero, the constraints are not satisfied exactly. Nonetheless, the distance clearly reduces, which shows that PARFUM moves toward those constraints. It is the direction in which the portfolio weights change that we are interested in. These frequency shifts are shown in Fig. 22. Technology 3 is clearly the best choice, as it is also the only portfolio that receives increased weight, alongside the baseline portfolio. This is not too different from the individual sample reweighting result, except that technology 1 has completely dropped from the resampled set.

VI. Discussion

This paper presents a method to select from among a set of technologies based on the probabilistic assessment of several quantities of interest. Conventional approaches to this problem do not employ uncertainty. Instead, deterministic points are picked at which the technologies are evaluated and selected on. In this section, the merits of the current probabilistic approach are discussed, compared to the conventional approach. Furthermore, some drawbacks are identified as well.

Using probability distributions instead of deterministic values not only reflects uncertainty during the conceptual design phase but also allows inclusion of difficult-to-quantify QOI. Take the cost in this paper, for example, which is only reflected using engineering insight to increase with higher technology performance. A Gaussian noise is added to introduce uncertainty in that assumption. The actual value of the cost metric does not matter, as only the relative change is of importance. Similar approaches can be taken to include strategic preferences, or metrics such as reliability or aesthetics. The fact that a distribution is used rather than a deterministic, fixed point reduces the bias of assumptions made during this process.

Another argument favoring the use of a probabilistic approach is when the effect of technologies are highly nonlinear in the QOI. In the present paper, the effect of each technology is fairly linear, which shows in Fig. 16 as parallel CDFs. As soon as the CDFs cross each other more, the benefit of a technology over another becomes more ambiguous. Consequently, picking a deterministic point at which to evaluate them becomes more difficult and arbitrary.

IPF requires a set of samples, which is why sampling is used as the uncertainty propagation technique. The obvious disadvantage of sampling is the computational cost. Thousands of samples are no exception (we used 10,000 to 20,000 samples in this paper), and for each sample, the analysis method has to be run. This incurs significant computational time, especially for more advanced simulations (our simple mission analysis takes a couple of seconds to run, so for all samples, several hours are required). In future studies, therefore, strategies to reduce the amount of required samples should be researched. One idea is to not compute all portfolios, but only a

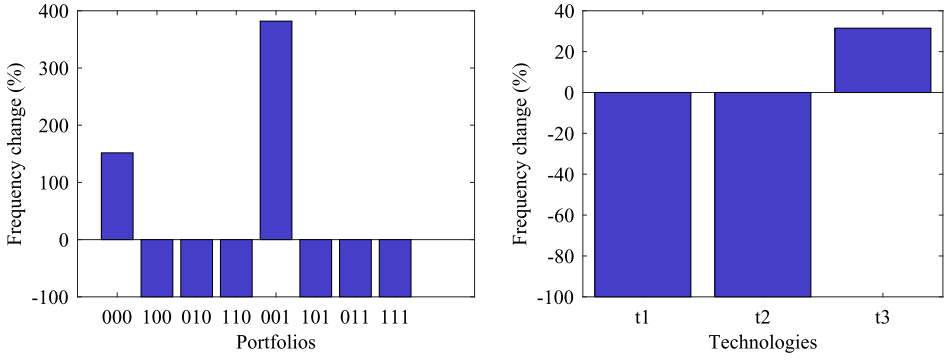

Fig. 22 Portfolio and technology frequency shifts using PI with a multi-objective constraint, grouped sample reweighting and run with PARFUM.

subset that covers the design space efficiently. Nonetheless, this issue is common to conventional sample-based approaches as well. Conversely, PI itself is very fast and incurs no significant time expenditure.

However, because PI only relies on a sample file, any analysis method can be used. Therefore, black box simulations and complex function can be employed easily. The only restriction is, again, that it is fast enough to generate thousands of samples.

Another effect of PI only relying on a sample file is that it can never exceed the boundaries defined by those samples. Thus, PI is confined to the design space covered by the initial sample set. If, however, distributions after PI would skew toward the limits of their range, that would indicate that the design space is not large enough and a more optimal solution may be found outside of the existing samples.

As the previous paragraph suggests, PI may be used to iteratively redefine the starting distributions and dependencies for the analysis. In some cases, PI incurs a dependency between variables that should not be there, requiring a statement of independence. Otherwise, PI may generate a dependency that was not modeled but should be present from a physical or model perspective. Then, this dependency can be included in subsequent runs of the analysis routine.

With respect to the accuracy of the results, two things can be said. First, PI is devised to provide an approximation of the inverse of a stochastic function. Therefore, besides the inputs and outputs being random variables, the redistribution of these variables due to PI is approximate. That is a result of the discretization of the constraints, and the resampling rather than inverting of the function. Second, after reweighting, the original samples are resampled. Whenever sampling occurs, there is an associated sampling bias, due to the random nature of a sampling process. This is expected to have very little influence on the CDFs after PI, but it could show some shift in the portfolio and technology frequency computations. Thus, it is advised to perform resampling multiple times and take an average of the results, or show an error bar. Nonetheless, the deviations due to sampling bias are expected to be small, and when a large difference in frequency is observed after PI, this likely outweighs any possible sampling bias.

Optimization approaches are an alternative to performing function inversion. When the metrics used for optimization are in the form of probability distributions, only some scalar measures of those distributions (e.g., mean and variance) are actually used as objective functions. PI uses the entire distribution directly, thus keeping more information. Furthermore, PI easily handles multiple objectives, without requiring some weighted combination of the objectives as optimization approaches do. Finally, while optimization approaches aim to find a minimum objective value, PI only targets a specified value. It therefore is less inclined to push the boundaries of the design space. However, which of the two methods is more appropriate depends on the problem at hand. For technology prioritization, however, we advocate the use of PI.

PI also differs from design exploration methods in that it defines a new design space that satisfies the specified requirements, rather than only selecting subsets of the initial design space that meet those requirements. In other words, design space exploration can only indicate the direction to move toward, while PI specifies the path to take.

VII. Conclusions

The merits of a probabilistic approach toward technology evaluation and selection are exposed in this paper. Rather than evaluating a deterministic, fixed point design (e.g., only one specific aircraft and mission), a technology is evaluated for an entire space of aircraft and missions. Using probabilistic inversion, target distributions on QOI are set, which consequently result in different distributions on the input variables. This shows how the technology variables (and possibly others) need to change in order to achieve certain requirements. Similarly, when multiple technologies are combined into portfolios, the combined effects are quantified, and probabilistic inversion shows how the frequency of the technologies should change to achieve certain goals. Multi-objective goals can be imposed on PI, without having to be converted to a single-objective function, as usually is done in design optimization. Moreover, PI can be used with any analysis method because it only relies on a set of samples obtained from the analyses. However, that limits the complexity of the model, due to computational time constraints. Thus, PI appears to be a powerful tool during conceptual design to explore the technology design space and prioritize technologies for selection.

Acknowledgments

This research is sponsored by the European Commission under the CleanSky II research program as part of project MANTA with grant agreement number 724558.

References

[1] Ozturk, B., and Saab, A., "Optimal Aircraft Design Decisions Under Uncertainty via Robust Signomial Programming," *AIAA Aviation 2019 Forum*, AIAA Paper 2019-3351, 2019. https://doi.org/10.2514/6.2019-3351

[2] Cook, L. W., and Jarrett, J. P., "Robust Airfoil Optimization and the Importance of Appropriately Representing Uncertainty," *AIAA Journal*, Vol. 55, No. 11, 2017, pp. 3925–3939. https://doi.org/10.2514/1.J055459

[3] Bianchi, D., Orra, T. H., Silva, C., and Silvestre, F. J., "Optimization Under Uncertainty: Rethinking Margins at the Conceptual Design," *2018 Multidisciplinary Analysis and Optimization Conference*, AIAA Paper 2018-3104, 2018. https://doi.org/10.2514/6.2018-3104

[4] Cooke, R. M., Zang, T. A., Mavris, D. N., and Tai, J., "Sculpting: A Fast, Interactive Method for Probabilistic Design Space Exploration and Margin Allocation," *16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA Paper 2015-3440, 2015. https://doi.org/10.2514/6.2015-3440

[5] Amadori, K., Bäckström, E., and Jouannet, C., "Future Technologies Prioritization for Aircraft Conceptual Design," *2018 AIAA Aerospace Sciences Meeting*, AIAA Paper 2018-1746, 2018. https://doi.org/10.2514/6.2018-1746

[6] Chakraborty, I., and Mavris, D. N., "Assessing Impact of Epistemic and Technological Uncertainty on Aircraft Subsystem Architectures," *Journal of Aircraft*, Vol. 54, No. 4, 2017, pp. 1388–1406. https://doi.org/10.2514/1.C034007

[7] Soban, D. S., and Mavris, D. N., "Assessing the Impact of Technology on Aircraft Systems Using Technology Impact Forecasting," *Journal of Aircraft*, Vol. 50, No. 5, 2013, pp. 1380–1393. https://doi.org/10.2514/1.C031871

[8] Huang, Y., and Soban, D. S., "Technology Impact Forecasting as a Framework for Assessment of Multi-Functional Composites," *2018 Aviation Technology, Integration, and Operations Conference*, AIAA Paper 2018-2868, 2018. https://doi.org/10.2514/6.2018-2868

[9] Gatian, K. N., and Mavris, D. N., "Facilitating Technology Development Progression Through Quantitative Uncertainty Assessments," *AIAA Aviation Forum*, AIAA Paper 2014-2170, 2014. https://doi.org/10.2514/6.2014-2170

[10] Gatian, K. N., and Mavris, D. N., "Enabling Technology Portfolio Selection Through Quantitative Uncertainty Analysis," *15th AIAA Aviation Technology, Integration and Operations Conference*, AIAA Paper 2015-2436, 2015. https://doi.org/10.2514/6.2015-2436

[11] Gatian, K. N., and Mavris, D. N., "Planning Technology Development Experimentation Through Quantitative Uncertainty Analysis," *54th AIAA Aerospace Sciences Meeting*, AIAA Paper 2016-0536, 2016. https://doi.org/10.2514/6.2016-0536

[12] Jimenez, H., and Mavris, D., "Pareto-Optimal Aircraft Technology Study for Environmental Benefits with Multi-Objective Optimization," *Journal of Aircraft*, Vol. 54, No. 5, 2017, pp. 1860–1876. https://doi.org/10.2514/1.C033688

[13] Desai, A. A., Hollingsworth, P. M., and Jinks, S., "Stochastic Aircraft Optimization and Decision Making Using a Competitive Value-Driven Design Framework," *2018 Aviation Technology, Integration, and Operations Conference*, AIAA Paper 2018-2869, 2018. https://doi.org/10.2514/6.2018-2869

[14] Zadeh, L., "Fuzzy Sets as a Basis for Possibility," *Fuzzy Sets and Systems*, Vol. 100, No. 1, 1978, pp. 9–34.

[15] Helton, J. C., Johnson, J. D., and Oberkampf, W. L., "An Exploration of Alternative Approaches to the Representation of Uncertainty in Model Predictions," *Reliability Engineering and System Safety*, Vol. 85, Nos. 1–3, 2004, pp. 39–71. https://doi.org/10.1016/j.ress.2004.03.025

[16] Liu, Z., and Huang, Y., "Technology Evaluation and Decision Making for Sustainability Enhancement of Industrial Systems Under Uncertainty," *American Institute of Chemical Engineers Journal*, Vol. 58, No. 6, 2012, pp. 1841–1852. https://doi.org/10.1002/aic.13818

[17] Bae, H. R., Grandhi, R. V., and Canfield, R. A., "Epistemic Uncertainty Quantification Techniques Including Evidence Theory for Large-Scale Structures," *Computers and Structures*, Vol. 82, Nos. 13–14, 2004, pp. 1101–1112. https://doi.org/10.1016/j.compstruc.2004.03.014

[18] Dempster, A. P., "A Generalization of Bayesian Inference," *Journal of the Royal Statistical Society: Series B (Methodological)*, Vol. 30, No. 2, 1968, pp. 205–232. https://doi.org/10.1111/j.2517-6161.1968.tb00722.x

[19] Oberkampf, W. L., Helton, J. C., and Sentz, K., "Mathematical Representation of Uncertainty," *Non-Deterministic Approaches Forum*, AIAA Paper 2001-1645, 2001. https://doi.org/10.2514/6.2001-1645

[20] Akram, F., and Mavris, D. N., "Uncertainty Propagation in Technology Valuation Based on Expert Elicitation," *58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, AIAA Paper 2012-882, 2017. https://doi.org/10.2514/6.2012-882

[21] Akram, F., Prior, M., and Mavris, D. N., "A Comparison Between Monte Carlo and Evidence Theory Approaches for Technology Portfolio Planning," *AIAA Infotech @ Aerospace*, AIAA Paper 2011-1412, 2011. https://doi.org/10.2514/6.2011-1412

[22] Akram, F., Prior, M. A., and Mavris, D. N., "An Application of Evidence Theory to Subject Matter Expert Based Technology Portfolio Analysis," *Infotech @ Aerospace*, AIAA Paper 2011-1413, 2011. https://doi.org/10.2514/6.2011-1413

[23] Akram, F., and Mavris, D. N., "Uncertainty Propagation in Technology Valuation Based on Expert Elicitation," *50th AIAA Aerospace Sciences Meeting*, AIAA Paper 2012-882, 2012. https://doi.org/10.2514/6.2012-882

[24] Akram, F., "A Methodology for Uncertainty Quantification in Quantitative Technology Valuation Based on Expert Elicitation," Ph.D. Dissertation, Georgia Inst. of Technology, Atlanta, GA, May 2012.

[25] Cooke, R., "The Anatomy of the Squizzel: The Role of Operational Definitions in Representing Uncertainty," *Reliability Engineering and System Safety*, Vol. 85, Nos. 1–3, 2004, pp. 313–319. https://doi.org/10.1016/j.ress.2004.03.019

[26] Cooke, R. M., "Messaging Climate Change Uncertainty," *Nature Climate Change*, Vol. 5, No. 1, 2015, pp. 8–10. https://doi.org/10.1038/nclimate2466

[27] Du, C., Kurowicka, D., and Cooke, R., "Techniques for Generic Probabilistic Inversion," *Computational Statistics and Data Analysis*, Vol. 50, No. 5, 2006, pp. 1164–1187. https://doi.org/10.1016/j.csda.2005.01.002

[28] Cooke, R., and Kurowicka, D., *Uncertainty Analysis with High Dimensional Dependence Modelling*, Wiley Series in Probability and Statistics, Wiley, Hoboken, NJ, 2006, Chap. 9. https://doi.org/10.1002/0470863072

[29] Kruithof, J., "Telefoonverkeersrekening," *De Ingenieur*, Vol. 52, 1937, pp. 15–25.

[30] Csiszár, I., "I-Divergence Geometry of Probability Distributions and Minimization Problems," *Annals of Probability*, Vol. 3, No. 1, 1975, pp. 146–158.

[31] Cooke, R. M., "Parameter Fitting for Uncertain Models: Modelling Uncertainty in Small Models," *Reliability Engineering and System Safety*, Vol. 44, No. 1, 1994, pp. 89–102. https://doi.org/10.1016/0951-8320(94)90110-4

[32] Kraan, B., "Probabilistic Inversion in Uncertainty Analysis: And Related Topics," Ph.D. Thesis, Delft Univ. of Technology, Delft, The Netherlands, April 2002.

[33] Kraan, B., and Bedford, T., "Probabilistic Inversion of Expert Judgments in the Quantification of Model Uncertainty," *Management Science*, Vol. 51, No. 6, 2005, pp. 995–1006. https://doi.org/10.1287/mnsc.1050.0370

[34] Elidan, G., "Copula Bayesian Networks," *Advances in Neural Information Processing Systems*, Neural Information Processing Systems Foundation, 2010, pp. 559–567.

[35] Sklar, A., "Random Variables, Joint Distribution Functions, and Copulas," *Kybernetika*, Vol. 9, No. 6, 1973, pp. 449–460.

[36] de Vries, R., Brown, M., and Vos, R., "Preliminary Sizing Method for Hybrid-Electric Distributed-Propulsion Aircraft," *Journal of Aircraft*, Vol. 56, No. 6, 2019, pp. 2172–2188. https://doi.org/10.2514/1.C035388

[37] Zaidi, T., Jimenez, H., and Mavris, D., "Copulas Theory for Probabilistic Assessment: Overview with Application to Airplane Performance Analysis," *Journal of Aircraft*, Vol. 52, No. 6, 2015, pp. 1802–1820. https://doi.org/10.2514/1.C033073

[38] White, R. J., "Improving the Airplane Efficiency by Use of Wing Maneuver Load Alleviation," *Journal of Aircraft*, Vol. 8, No. 10, 1971, pp. 769–775. https://doi.org/10.2514/3.59169

[39] Xu, J., and Kroo, I., "Aircraft Design with Maneuver and Gust Load Alleviation," *29th AIAA Applied Aerodynamics Conference*, AIAA Paper 2011-3180, 2011. https://doi.org/10.2514/6.2011-3180

[40] Mancini, A., and Vos, R., "The Effect of Maneuver Load Alleviation Strategies on Aircraft Performance Indicators," *AIAA Aviation 2019 Forum*, AIAA Paper 2019-3272, 2019. https://doi.org/10.2514/6.2019-3272

[41] Ting, E., Chaparro, D., Nguyen, N., and Fujiwara, G. E. C., "Optimization of Variable-Camber Continuous Trailing-Edge Flap Configuration for Drag Reduction," *Journal of Aircraft*, Vol. 55, No. 6, 2018, pp. 2217–2239. https://doi.org/10.2514/1.C034810

[42] Obert, E., *Aerodynamic Design of Transport Aircraft*, IOS, Amsterdam, 2009, pp. 362–363.

[43] Fonte, F., Toffol, F., and Ricci, S., "Design of a Wing Tip Device for Active Maneuver and Gust Load Alleviation," *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA Paper 2018-1442, 2018. https://doi.org/10.2514/6.2018-1442

[44] Binois, M., Rullière, D., and Roustant, O., "On the Estimation of Pareto Fronts from the Point of View of Copula Theory," *Information Sciences*, Vol. 324, Dec. 2015, pp. 270–285. https://doi.org/10.1016/j.ins.2015.06.037

[45] Cook, L. W., and Jarrett, J. P., "Horsetail Matching for Optimization Under Probabilistic, Interval and Mixed Uncertainties," *19th AIAA Non-Deterministic Approaches Conference*, AIAA Paper 2017-0590, 2017. https://doi.org/10.2514/6.2017-0590