Which uncertainty is important in multistage stochastic programs? A case from maritime transportation

Giovanni Pantuso
giopan@dtu.dk
Technical University of Denmark
Frederiksborgvej 399, 4000 Roskilde, Denmark
and Norwegian University of Science and Technology,
Afred Getz vei 3, Trondheim, 7491 Norway.

Kjetil Fagerholt
Norwegian University of Science and Technology,
Afred Getz vei 3, Trondheim, 7491 Norway.
and Norwegian Marine Technology Research Institute (MARINTEK)

Stein W. Wallace
Norwegian School of Economics,
Helleveien 30, 5045 Bergen, Norway.

January 9, 2020

Abstract

Given that the scope of stochastic programming is to suggest good decisions and not to estimate probability distributions, we demonstrate in this paper how to numerically evaluate which properties of random variables are more important to capture in a stochastic programming model. Such analysis, performed before data collection, can indicate which information should be primarily sought, and which is not critical for the final decision. We apply the analysis to a real-life instance of the maritime fleet renewal. Results show that some properties of the stochastic phenomena, such as the correlation between random variables, have very little influence on the final decision.

1 Introduction

Stochastic programming amounts to formulating and solving mathematical programming models where some of the coefficients are not known with certainty and thus represented by random variables. Besides a model of the decision problem (i.e., a the stochastic program), a model of the uncertainty is also required. The latter model is based on data which, depending on the specific case, may be available in different formats, ranging from detailed time series to simple personal beliefs about the random phenomena. As an example, data can be available in the form of experts opinions which are difficult to express numerically, such as "a patient suffering from pathology P is more likely to live an active life after one year if he receives treatment T after D days". The resulting model of uncertainty can then be more or less detailed (e.g., assigning probabilities to two or three events or drawing an extensive, possibly multi-variate, probability distribution). In this phase, understanding how important different *properties* of the uncertainty (e.g., means, supports, correlations or variances) are, for the specific decision model, can lead to better models of the uncertainty as well a to more effective data collection/analysis efforts. As an example, this can guide data collection if the available data is scarce and important properties cannot be inferred, or data analysis if the data is abundant, such that only the properties that matter are carried along in the resulting model.

In any case, to solve the resulting stochastic program, one requires a discrete representation of the uncertainty (i.e., in order to avoid solving integrals). Typically, scenario generation methods are based on sampling or on matching statistical properties of the stochastic phenomena. Particularly, the latter may be suitable for the case when the modeler is only aware of, or interested in, some of the properties of the underlying distribution. For an overview over scenario generation methods see, e.g., Dupacová *et al.* (2000) and King & Wallace (2012, Ch.4). The goal of scenario generation is to generate a set of scenarios such that the decision model behaves (almost) as if we used the original distribution. And however one models the uncertainties, there is always the question: How good is the uncertainty model? And beforehand: Which properties of the uncertain phenomena are more important to capture? We wish to give these questions an answer based on the actual decision model.

Similar questions interested Kallberg & Ziemba (1984); Broadie (1993); Chopra & Ziemba (1993) and more recently Kaut *et al.* (2007). In these studies the authors investigated the effect of errors in statistical properties (such as mean, variances or covariances) in the context of optimal portfolio selection. Our scope is that of proposing a more general analysis framework.

The purpose of this paper is, therefore, to demonstrate for multistage stochastic programs (MSPs) how to quantitatively evaluate the importance of individual properties of the stochastic phenomena and, as a result, the consequences of planning using incorrect values for a given property. Such analysis, besides addressing data collection and data analysis towards the important properties, can provide better understanding of the problem and of its drivers. The analysis, though general, is explanatorily performed on a case of the *maritime fleet renewal problem* (MFRP). The problem consists of choosing how many and which type of ships to invest in and when to do so, in order to cope with the uncertain future market situation. Most of the parameters of the problem are stochastic and stochastic optimization can be a useful tool for the MFRP. Pantuso *et al.* (2014) showed that using stochastic programming, rather than solving the mean value problem, can lower the total expected cost by up to around 10%, in some market contexts. Uncertainty is therefore an important element of the problem. In this paper we evaluate which properties of the uncertainty are more important to capture.

The remainder of this paper is organized as follows. In Section 2 we describe how we evaluate the importance of a property. In Section 3 we analyze different properties of the uncertainty in the MFRP while conclusions are drawn in Section 4.

2 Importance of a property

In this section we demonstrate how to estimate the importance of individual properties of the uncertainty in MSPs. That is, we study how optimal values of optimization problems depend on properties (such as means, variances, correlations, supports, and shapes). We do that by first defining a probability distribution across all random variables using our best estimates of all properties. Applying the scenario generation method proposed by Høyland *et al.* (2003), we generate a scenario tree for our problem using this distribution. We then repeat the process, but such that one property is changed. Now we can appreciate how optimal values are influenced by the property we changed. Below we describe how this can be done for multistage problems.

Consider a property that we wish to investigate (e.g., we are not confident about our estimates and wish to know if that matters). We define a probability distribution having a given value for that property (e.g., our best estimate) and generate a scenario tree based on it. Given the scenario tree, we solve the corresponding MSP, and store its optimal objective value. Let this represent the *expected result of using the correct property value* (ERCP). Then, in order to evaluate the expected result of using an incorrect property value, we define a probability distribution such that the value of the property under investigation (and only that) is changed. We make first-stage decisions using a scenario tree based on the new distribution. The difference between this model and the former is that this one sees a different (incorrectly described) uncertain future. The first-stage solution is then stored (this is the only part of the solution that actually would be implemented). At the beginning of the second stage, given the new information (i.e., we realize in which node of the correct tree we are) and first-stage decisions (the solution we stored), we build another model to make conditional decisions for the second stage. The farther future (stages three, four and so on) is again represented by a scenario tree matching the incorrect property value and previous decisions are fixed. Conditional second-stage solutions are stored (the only ones which would be implemented). This process continues for the subsequent stages. When we have stored decisions for all stages but the last, we observe the last-stage realizations, fix the variables at all the previous stages to the corresponding stored solutions, and solve a stochastic program for the last-stage variables. The corresponding optimal objective value represents the *expected result of using the incorrect property value* (ERIP), and can be compared to the expected value of using the correct property value calculated before.

Notice that, in this process, we never make solutions which anticipate the future. At all stages, decisions are made based on current information (including past decisions) but on incorrect descriptions of the uncertain future. This shrinking-horizon procedure is close to the one used by Escudero *et al.* (2007) to calculate the "dynamic value of the stochastic solution" for MSPs. When calculating the expected return of using the mean value problem, at each node in the scenario tree (except those the last stage), they make decisions based on past decisions, available information, and on approximating future stages by their expected values. Our procedure is identical, except for the fact that in our case future stages are still uncertain, but incorrectly described. A similar procedure is used by Fleten *et al.* (2002). They test the performance of two different approaches to portfolio management on a shrinking-horizon scheme. They make first-stage decisions for both approaches based on the same scenario tree. Then they generate a number of simulation scenarios, and for each of them they generate conditional solutions for the following stages taking into account the first-stage solution, again for both approaches. Each scenario gives a measure of the performance of each approach.

Assume the possible values for the property under investigation can be limited to a set Π. The set Π can be based, for example, the decision makers knowledge about the problem or on the observation of the available data. Designate $\pi \in \Pi$ and $\pi' \in \Pi$ as the correct and incorrect value for the property, respectively.

Definition 2.1 The Pairwise Property Error – denoted by $PPE(\pi, \pi')$ – is the loss incurred when using π' rather than the correct property value π', and is calculated as

$$PPE(\pi, \pi') = |ERCP - ERIP|$$

By evaluating the $PPE(\pi, \pi')$ for different pairs of property values in the given set, one can come to an understanding of how the specific property can influence optimal objective values. It gives a measure of how wrong one can be if the value of that property is incorrect, and consequently the value of obtaining the right property value.

The procedure to calculate the $PPE(\pi, \pi')$ is described in Fig. 1 for an example three-stage problem. Assume that the scenario tree in Fig. 1a represents the one matching by the correct property value, say π (note that its nodes are drawn with solid lines). When planning with the incorrect property value, say π', at the root node one would make decision x_0 based on future outcomes mistakenly described by a scenario tree matching incorrect property value π'. The incorrect part of the scenario tree is drawn in dashed lines in Fig. 1b. Solution x_0 is saved. At stage $t = 1$, one can end up in either node 1 or 2 of the correct tree in Fig. 1a, as shown in Fig. 1c. Past decisions (i.e., those belonging to the root node) are fixed at the value x_0 previously saved (indicated by a crossed node). Decisions for nodes 1 and 2 are made based on an incorrect description of the farther future (i.e., based on π' for $t = 2$), see Fig. 1c. Finally, at $t = 2$ all the uncertainty is disclosed. Past decisions are fixed and one needs to solve for the last-stage variables, see Fig. 1d. The corresponding objective value represents the ERIP, i.e., the expected result of using the incorrect property value, and can be compared to the ERCP obtained by solving the stochastic program over the scenario tree in Fig. 1a, in order to calculate the $PPE(\pi, \pi')$.

3 Application to the Maritime Fleet Renewal Problem

In this section we estimate the importance of different properties of the uncertainty affecting the Maritime Fleet Renewal Problem (MFRP). The problem consists of deciding how to modify the available fleet of ships in order to efficiently meet future market requirements. The fleet can be expanded by ordering new ships or buying second-hand ones, and shrunk by selling or demolishing available ships. For short-term needs, ships can also be chartered in and out for periods of time (e.g., days, weeks or months). When deciding upon fleet renewal plans shipping companies have to take ship operations into account in order to make a correct estimation of the tonnage requirement. In liner shipping, as an example, ships sail on *trades*, which consist of fixed routes between two geographic areas (e.g., Europe to North America). Every time a ship sails on a trade it transports cargo from the origin to the destination. Therefore, in a given period, a number of

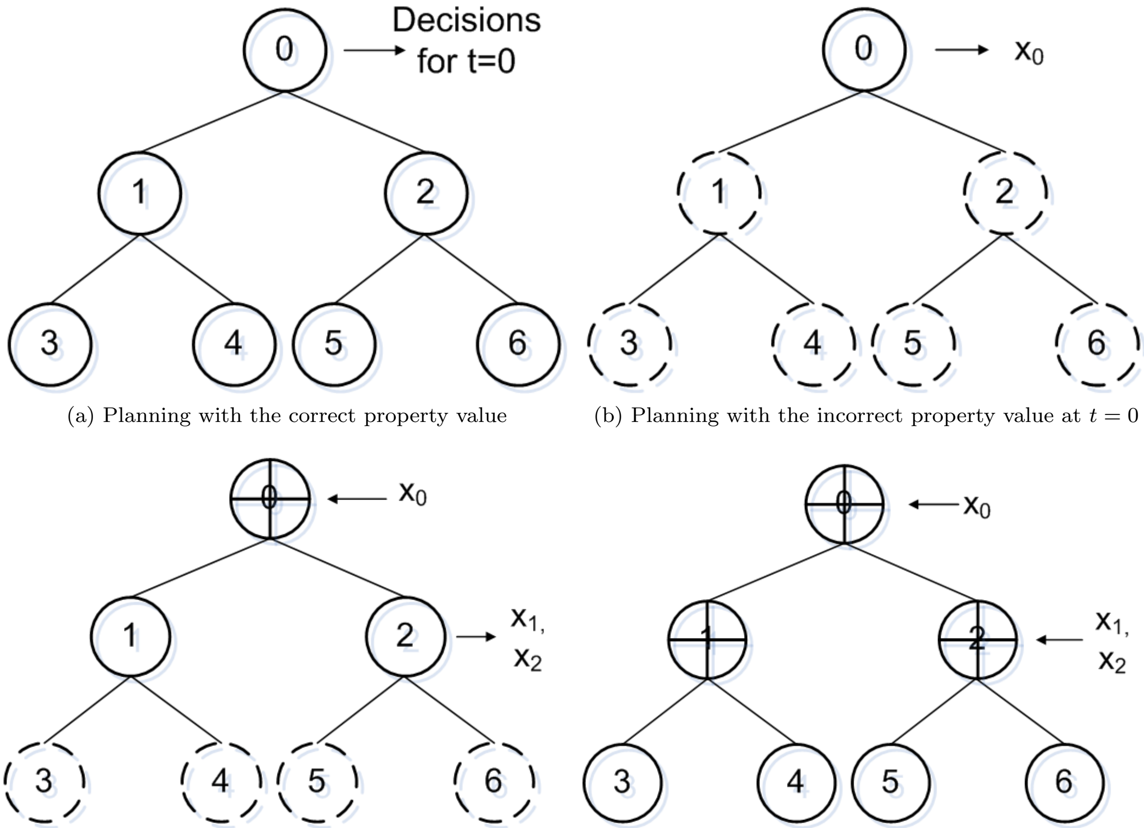

(a) Planning with the correct property value

(b) Planning with the incorrect property value at $t = 0$

(c) Planning with the incorrect property value at $t = 1$

(d) Last-stage decisions

Figure 1: *Procedure to calculate the* $PPE(\pi, \pi')$*. Nodes drawn with solid lines represent the correct representation of uncertainty, while those in dashed lines represent an incorrect one. Crossed nodes represent nodes whose variables are fixed.*

voyages on each trade are in general performed in order to satisfy the corresponding transportation demand. Different voyages may be performed by different ships, and each ship can alternate voyages on different trades. Ocean-going ships come usually in a big variety of types differing at least in size, speed, cost and technological characteristics.

The problem is highly affected by uncertainty. Ship prices, charter rates, fuel prices, demolition rates and demand are in fact very volatile, and attempts to make accurate forecasts are usually hopeless in practice. Furthermore, given the long lifetime of ships, a long planning horizon must be considered. In addition, fleet renewal decisions are typically made periodically, e.g., every year, in light of new information about the market status (and the technological development) obtained. This makes the problem stochastic and multistage. A more thorough description of the problem can be found in Pantuso *et al.* (2014). The authors also propose a multistage stochastic model for the problem which will be used in the tests performed in what follows.

We assess the effect of planning with incorrect values of given properties of the uncertainty on a real-life stochastic problem, based on the $PPE(\pi, \pi')$ defined in Section 2. Pantuso *et al.* (2014) showed that, when the charter market does not offer many ships for chartering – which is typical in markets with very specialized ships types – using stochastic programming can noticeably lower the total expected cost. Uncertainty is therefore a very important element of the problem. Here we touch upon what statistical properties are more important to capture in the model of uncertainty. In Section 3.1 we introduce the instances we tested. In Section 3.2 we discuss the property "correlations", in Section 3.3 we discuss the properties "mean" and "support", and in Section 3.4 we discuss the property "shape of the distribution". Finally, in Section 3.5 we evaluate the effect of combined mis-specification of property values.

3.1 Instances

We use instances based on the real case of Wallenius Wilhelmses Logistics (WWL), a major liner shipping company engaged in transportation of rolling equipment. They transport three main types of cargo: cars, high and heavy vehicles (e.g., trucks or tractors), and breakbulk cargo (i.e., high-weight/volume items such as train coaches or big engines). Instances vary in the number of ship types, trades and scenarios considered. Table 1 reports the instances considered. Letters S, M and L stand for small medium and large, respectively, with respect to the number of ship types and trades, while the following number represents the number of stages (e.g., M3 indicates a three-stage medium-size instance).

Scenarios are generated by means of a variant of the heuristic proposed by Høyland *et al.* (2003), which creates scenarios matching the first four marginal moments of the distribution plus the given correlation matrix. Acceptable *in sample stability* (see Kaut & Wallace (2007)) is achieved with 15 scenarios in the second-stage and 10 the third. However, the reliability of results we propose in the following sections was confirmed by using up to 200 scenarios for two-stage programs, and 1600 (40 times 40) for the three-stage programs. Since we decided to work only with problems which we can solve to optimality, in order for the results not to be biased by the quality of heuristic solutions, we only consider up to three stages. The length of the planning horizon considered is six year, which represents the length of the forecasts made by WWL.

Table 1: Instances of the MFRP

Instance	S2	M2	L2	S3	M3	L3
# Ship types	6	8	10	6	8	10
# Trades	5	8	12	5	8	12
# Stages	2	2	2	3	3	3
# Scenarios	15	15	15	150	150	150

We consider a total of six random variables for each stage t, $\xi_1^t, \ldots, \xi_6^t$, each determining the realization of associated parameters of the MFRP at stage t. Each parameter associated to a given random variable is therefore itself a random variable (*random parameter* in what follows). Random parameters are connected to the underlying random variable by a relationship of type: $RP_t = E[RP_t](1 + \delta\xi_i^t)$, where RP_t is the realization of the random parameter at stage t, $E(RP_t)$ is its expected value, $\delta \in [0, 1]$ is a constant which determines the support width, and ξ_i^t, with $i \in \{1, \ldots, 6\}$ is the underlying random variable. We will assume that the company forecast represent expected values.

Variable ξ_1^t determines the realization of fuel prices at stage t. Variable ξ_2^t determines the realization of second-hand ship purchase and selling prices, newbuilding prices and charter (in and out) rates. Variable ξ_3^t determines the realization of the steel price (and consequently the scrapping rates). Finally, variables ξ_4^t, ξ_5^t and ξ_6^t determine the realizations of the demand of cars, high and heavy vehicles and breakbulk cargo, respectively. All random parameters associated to the same random variable are assumed perfectly correlated (e.g., ship prices and charter rates are therefore assumed perfectly correlated). Since for each random parameter the forecasted value is assumed as expected value, random variables represent errors in the forecast. For this reason we assume random variables to be uncorrelated along time, i.e., ξ_i^t is uncorrelated with $\xi_i^{t'}$ for each t, t'.

In all the tests proposed in what follows, except for where otherwise indicated, all the random variables have uniform marginal distributions over the support $[-1, +1]$. The actual support, for each associated random parameter is defined by a constant δ. Values of δ are 0.6 for demands and second-hand prices, 0.5 for charter rates, 0.15 for newbuilding costs, 0.4 for the fuel price and 0.2 for the scrapping rates. As an example $\delta = 0.5$ corresponds to allowing the random parameter to vary by $\pm 50\%$ from its expected value. These values are chosen in order to reflect the fact that demand, ship prices and charter rates are in general more volatile than newbuilding prices and scrapping rates (which besides the status of the market also depends on the cost of labor and of materials and are somewhat more predictable). The importance of the support and of the distribution are tested in Section 3.3 and Section 3.4, respectively.

Table 2: *Correlation matrices at all stages*

(a) Matrix A

	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6
ξ_1	1	0.65	0.7	0.85	0.75	0.75
ξ_2		1	0.9	0.85	0.85	0.85
ξ_3			1	0.8	0.8	0.8
ξ_4				1	0.8	0.7
ξ_5					1	0.7
ξ_6						1

(b) Matrix B

	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6
ξ_1	1	0.65	0.7	0.85	0.8	0.75
ξ_2		1	0.9	0.4	0.4	0.4
ξ_3			1	0.4	0.4	0.4
ξ_4				1	0.8	0.7
ξ_5					1	0.7
ξ_6						1

(c) Matrix C

	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6
ξ_1	1	0.65	0.7	0.4	0.4	0.4
ξ_2		1	0.9	0.4	0.4	0.4
ξ_3			1	0.4	0.4	0.4
ξ_4				1	0.3	0.3
ξ_5					1	0.3
ξ_6						1

(d) Matrix D

	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6
ξ_1	1	0.65	0.7	0.85	0.3	0.3
ξ_2		1	0.9	0.85	0.3	0.3
ξ_3			1	0.8	0.3	0.3
ξ_4				1	0.3	0.3
ξ_5					1	0.3
ξ_6						1

Problems have been modeled and solved using the Java callable libraries of ILOG Cplex 12.2 on a 2x2.4GHz AMD Opteron 2431, 6 core machine, with 24Gb ram.

3.2 Importance of the Correlations

In this section we evaluate the property "correlations". Consider the four correlation matrices in Table 2. Matrix A (Table 2a) represents a situation in which the three demands are strongly positively correlated between each other, and to all the costs and prices. All the components of the problem tend to move in the same direction. Matrix B (Table 2b) represents the case in which the demands are still strongly correlated between each other, but are little correlated to ship prices, charters and scrapping rates. Matrix C (Table 2c) represents the case in which the demands are weakly correlated between each other and to all prices and costs. Finally, Matrix D (Table 2d) describes the case in which only the demand of cars is correlated to the costs, prices, charters and scrapping rates of RoRo ships, and to fuel prices, while the other demands are weakly correlated with both the car demand and all the costs and prices.

Assume we do not know which one correctly describes the market configuration, but we can limit the choice to the four matrixes. We want to find out what happens if we use the incorrect correlation matrix. We understand that not all correlation matrices may represent real market configurations, but we choose to work with significantly different matrices, in order to test relatively extreme cases. When a correlation matrix is chosen, it applies at all stages, therefore we simplify the notation by dropping the t index in Table 2.

Table 3 reports, for each pair of correlation matrices the average (Avg) and max (Max) $PPE(\pi, \pi')$ over all the instances. For each instance the $PPE(\pi, \pi')$ is calculated ten times, each time with new scenario trees (one matching the correct and the other the incorrect property value). The $PPE(\pi, \pi')$ is expressed as a percentage of the expected cost obtained working with the correct property value, $ERCP$, in order to preserve confidentiality relative to the objective value.

The increase in the expected cost for working with the incorrect correlation matrix is relatively small in all cases. This might seem surprising at first sight. One may expect that, when demands are strongly correlated, the initial investment is higher because the company is more likely to be facing simultaneous peaks in the demand of all the three products. Indeed, when demands are strongly positively correlated, the stochastic program suggests a slightly bigger investment/smaller disinvestment. However, the solution obtained with weakly correlated demands, though suggesting a smaller investment, is still near optimal. Some of the reasons are: 1) the objective function is relatively flat, presenting several near optima, 2) in both cases the most flexible ship types are chosen, providing the ability to adapt to different market configurations, and 3) the availability of different types of recourse actions such as, second-hand purchases/sales, time and voyage

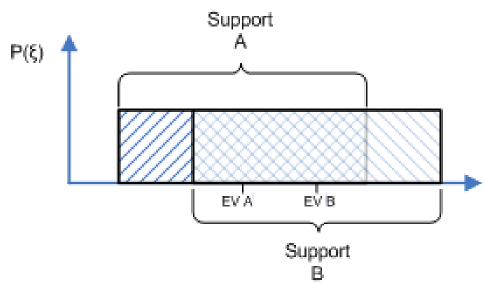

Figure 2: Supports for uniformly distributed random variables

charters. Related to that, a reasonable question might be whether planning with incorrect correlations is more dangerous when the recourse actions become more expensive. Indeed, we experienced an increase in the $PPE(\pi, \pi')$ when increasing the voyage charter costs. However, the impact of the errors in the correlations was still limited even when the voyage charter cost was doubled. It is always possible to obtain different numerical results by arbitrarily playing with the parameters of the problem. However, using realistic values for voyage and time charter costs, showed that the impact of mistakes in the correlations is limited for our case study.

Table 3: *Average (Avg) and maximum (Max)* $PPE(\pi, \pi')$ *for every pair of correlation matrices. The correct property value* π *is indicated by the row, while the incorrect property value,* π'*, is indicated by the column.*

	A		*B*		*C*		*D*	
	Avg [%]	Max [%]	Avg [%]	Max [%]	Avg [%]	Max [%]	Avg [%]	Max [%]
A	-	-	0.05	0.21	0.06	0.20	0.05	0.37
B	0.05	0.23	-	-	0.05	0.29	0.06	0.26
C	0.06	0.23	0.05	0.41	-	-	0.05	0.22
D	0.06	0.21	0.05	0.20	0.06	0.30	-	-

3.3 Importance of Mean and Support

Now consider these four supports for the random variables, $[-1, +1]$, $[-0.9, +1.1]$, $[-0.8, +1.2]$ and $[-0.7, +1.3]$ with mean values 0.0, 0.1, 0.2 and 0.3, respectively. Each random variable has uniform marginal distribution, therefore a shift in the support determines a shift in the mean. Fig. 2 shows two different uniform distributions where the supports have the same width, but different positions and expected values. Notice however that the width of the support of each random parameter is influenced by the corresponding value of the constant δ (see Section 3.1). What happens if we set an incorrect mean in the model of uncertainty?

Fig. 3 reports the average and the maximum $PPE(\pi, \pi')$ for each pair of mean values. The position of the support, and consequently the mean value, is a more critical property. In order to understand the reasons behind this we tested the $PPE(\pi, \pi')$ for the support of each individual random parameter (i.e. we tested the effect of mistaking individual supports while all the others are correct). Planning with higher expected demand or expected charter rate, leads to tonnage surpluses. In the former case the model will prepare to carry more cargo, in the latter to use less charters, thus a bigger fleet. Similarly, planning with a lower expected demand or expected charter rate, leads to tonnage deficits. The model is in fact setting up less capacity because it expects less cargo, in the former case, or to use more charters, in the latter. Conversely, planning with incorrect expected variable operating costs has very little impact on the solution, due to the fact that in any case, the model suggests investing in more efficient ship types.

As far as the width of the support is concerned, consider the supports $[-0.6, +0.6]$ and $[-0.1, +0.1]$. The support width is generated by adjusting the constant δ introduced in Section 3.1, while the support of the random variables remains $[-1, +1]$. In the former case the random parameter can vary between $\pm 60\%$ from the expected value, in the latter case between $\pm 10\%$. Fig. 4 reports the $PPE(\pi, \pi')$ for the property "support width" for charter rates, variable operating costs, second-hand prices and demands. We can notice

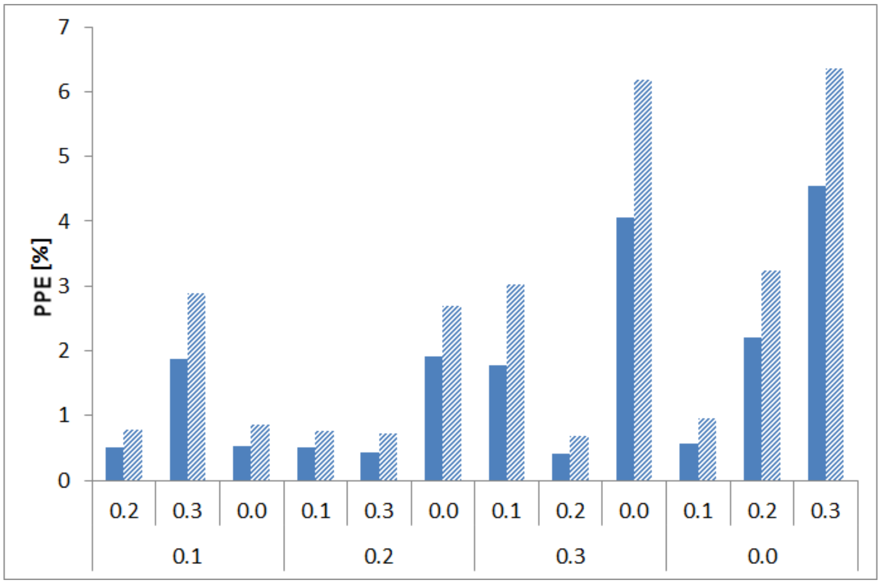

Figure 3: $PPE(\pi, \pi')$ for the property "Mean value". The horizontal axis reports the correct mean value, π, in the lower row, and the incorrect mean value, π', in the upper row. Plain bars indicate average values, while patterned bars indicate maximum values.

that the problem is more sensitive to mistakes in the width of the support of the demand, while the loss generated by incorrect supports for second-hand prices, charter rates and variable operating cost is lower. Based on the discussion above, if the mean and support for the demands and the mean of the charter rates are correct, the model has sufficient information to determine the correct tonnage requirement. However, setting an incorrect width for the support of the demand has an impact. This indicates that, if one is not sure about the support, just assuming it wider might not always be a safe way out.

This can be precious information from a shipping company perspective. WWL engage in transportation contracts which commit the company to carry a given percentage of the customer's production. The production is however uncertain and the customer might sign transportation contracts with different transport providers. If the company is able to negotiate upper and lower bounds on the amount to carry, this would be highly beneficial when making fleet renewal plans, as shown in Fig. 3 and Fig. 4.

3.4 Importance of the Shape of the Distribution

Now, we assume the modeler is fairly confident about the supports (i.e., $[-1, +1]$), correlations (i.e., matrix A in Table 2a), and mean (i.e., 0), but is not certain about how the probability is distributed. Consider the following marginal probability distributions uniform ($U[-1, 1]$), triangular ($T[-1, 1, 0]$), and "four-parameters" beta with $\alpha = \beta = 2$ ($B[2, 2, -1, 1]$). The uniform distribution is relatively different from the other two as it does not have a peak. The triangular and beta distributions are instead similar as they are both symmetrical around the mean.

In Table 4 we report the average and max $PPE(\pi, \pi')$ due to errors in the shape of the marginal distributions. Rows report the correct shape of the distribution while columns report the incorrect one. Although ignoring the peak seems to have a slightly higher impact on the fleet renewal plan (i.e., the average $PPE(\pi, \pi')$ is higher), the loss is relatively low. This suggests that the fourth moment of the distribution is not an important parameter to capture, or that at least that the first moment (mean) and the variance (at least for the demand) have a higher impact.

3.5 Effect of combined errors

Finally, we evaluate the effect of combined errors, that is, the effect of planning with incorrect values in two properties. In order to keep the notation readable, we continue referring to combined errors as $PPE(\pi, \pi')$

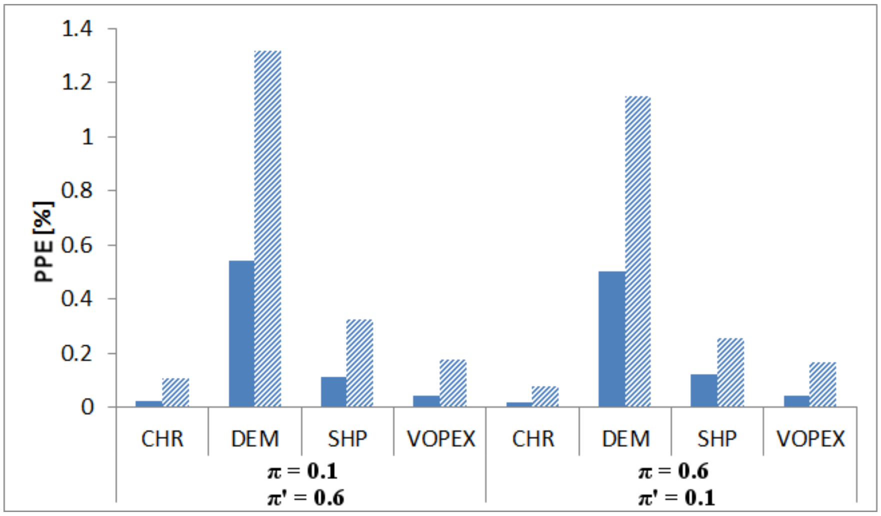

Figure 4: $PPE(\pi, \pi')$ relative to the support width for variable operating costs (VOPEX), second-hand prices (SHP), charter rates (CC) and demands (DEM). Plain bars indicate average values, while patterned bars indicate maximum values.

Table 4: $PPE(\pi, \pi')$ relative to the property "distribution". Rows report the correct distribution while columns report the incorrect one.

	T[-1,1,0]		U[-1,1]		B[2,2,-1,1]	
	Avg [%]	Max [%]	Avg [%]	Max[%]	Avg [%]	Max [%]
T[-1,1,0]	-	-	0.20	1.17	0.06	1.45
U[-1,1]	0.08	0.28	-	-	0.21	0.84
B[2,2,-1,1]	0.03	0.33	0.20	0.53	-	-

but in this case π and π' represent couples of property values. As an example, π might represent couple (shape=triangular, mean=0.2)".

We start by evaluating the $PPE(\pi, \pi')$ when an incorrect shape of the distribution is used together with an incorrect value of the mean. We assume the support to be $[-1, +1]$ and the correlations to be described by matrix A in Table 2a. Consider the two "four-parameters" beta distributions $B[\alpha, \beta, -1, 1] = B[3, 1.5, -1, 1]$ and $B[1.5, 3, -1, 1]$. The first beta distribution is left skewed while the second one is right skewed. The mean value is 0.3 for $B[3, 1.5, -1, 1]$ and -0.3 for $B[1.5, 3, -1, 1]$. If one plans, for example, with the distribution $U[-1, 1]$, but $B[1.5, 3, -1, 1]$ out to be the correct distribution, one is planning with both incorrect shape and mean value of the distribution. Therefore, we will in turn consider distributions $B[1.5, 3, -1, 1]$, $B[3, 1.5, -1, 1]$, $T[-1, 1, 0]$, and $U[-1, 1]$ to be the correct distributions. This allows us to evaluate the $PPE(\pi, \pi')$ on a combined error on shape and mean of the distribution.

Table 5: $PPE(\pi, \pi')$ on combined errors in means and shape of the marginal distributions.

	T[-1,1,0]		U[-1,1]		B[3,1.5,-1,1]		B[1.5,3,-1,1]	
	Avg [%]	Max [%]	Avg [%]	Max[%]	Avg [%]	Max [%]	Avg [%]	Max [%]
T[-1,1,0]	-	-	-	-	10.51	28.07	6.94	22.05
U[-1,1]	-	-	-	-	8.23	11.88	8.22	14.30
B[3,1.5,-1,1]	8.91	12.95	6.96	10.23	-	-	26.18	39.4
B[1.5,3,-1,1]	14.35	20.93	17.48	25.50	61.36	86.72	-	-

Table 5 reports the results of the tests. Clearly, the magnitude of the $PPE(\pi, \pi')$ increases noticeably with respect to the values obtained in the previous sections. The main driver is, however, again the mean value. Notice, in fact, that between $B[1.5, 3, -1, 1]$ and $B[3, 1.5, -1, 1]$ there is a 0.6 difference in mean value, higher than what tested in Section 3.3. However, the simultaneous error in the shape of the distribution increases the relevance of errors in the mean. Notice in fact that PPE($B[3, 1.5, -1, 1]$,$U[-1, 1]$) is caused by a 0.3 difference in the mean as in Fig. 3, but the simultaneous error in the shape of the distribution determines an increase of $PPE(\pi, \pi')$ of a few percentage points. Notice also that PPE($B[1.5, 3, -1, 1]$,$B[3, 1.5, -1, 1]$) is much higher than PPE($B[3, 1.5, -1, 1]$, $B[1.5, 3, -1, 1]$) showing that, in our case, overestimating the mean is far more dangerous than underestimating it. Overestimating the mean leads in fact to extremely high tonnage oversupply and, consequently unjustified investments in new ships.

Let us continue by evaluating the effect of simultaneous errors in the shape of the marginal distributions and in the correlations. Assume the support to be $[-1, +1]$ and the mean to fall in 0. In Table 6 we report the $PPE(\pi, \pi')$ obtained by assuming as the correct combination of property values, in turn, distribution $U[-1, 1]$ with correlations matrix B (see Table 2b), and distribution $T[-1, 1, 0]$ with correlations matrix C (see Table 2c). Notice that these combination have mean value and support in common. Matrices B and C have been chosen as they gave the highest $PPE(\pi, \pi')$ in Table 3. The two combinations are indicated as U-B and T-C, respectively, in Table 6. Notice that the error in the correlations determined only a slight increase in $PPE(\pi, \pi')$ with respect to Table 4, where only the distribution is incorrect.

Table 6: $PPE(\pi, \pi')$ on combined errors in correlations and shape of the marginal distributions.

	U-B		T-C	
	Avg [%]	Max [%]	Avg [%]	Max [%]
U-B	-	-	0.61	1.46
T-C	0.77	2.67	-	-

Finally, we evaluate the $PPE(\pi, \pi')$ caused by simultaneous errors in mean and correlations. In order to do that we focus on uniform marginal distributions, and evaluate the effect of shifting the mean while changing the correlation matrix. We consider two configurations, namely uniform distribution over the support $[-1, 1]$ with correlation matrix B (see Table 2b), say $B - U[-1, 1]$, and uniform distribution over the support $[-0.7, 1.3]$ with correlation matrix C (see Table 2c), say $C - U[-0.7, 1.3]$. Notice that the two configurations have different mean and correlation matrix. Table 7 reports the $PPE(\pi, \pi')$ relative to simultaneous errors in mean and correlations. Also in this case, incorrect correlations slightly amplified the magnitude of errors in the means. In fact, the results in Table 7 show a slight increase in $PPE(\pi, \pi')$ with respect to the results in Fig. 3, which were obtained assuming perfect knowledge of the correlation matrix.

Table 7: $PPE(\pi, \pi')$ on combined errors in correlations and means.

	$B - U[-1, 1]$		$C - U[-0.7, 1.3]$	
	Avg [%]	Max [%]	Avg [%]	Max [%]
$B - U[-1, 1]$	-	-	6.39	10.08
$C - U[-0.7, 1.3]$	6.10	9.58	-	-

The analysis of combined errors, on our case study, confirms the absolute importance of having correct mean values. It also supports the statement that the shape of the marginal distributions and the correlations play a secondary role. In fact, the results shown in Table 6 and Table 7 are consistent with the results shown in Table 4 and Fig. 3 where the shape of the marginal distributions and the means, respectively, where analyzed individually. However, this section demonstrates that properties which individually have little impact on decisions, may become misleading if incorrectly estimated when the decision maker has little knowledge about other relevant properties. In this case, having correct values for the correlations can lower the loss in case other properties are incorrectly estimated.

4 Conclusions

The scope of modeling uncertainty is to represent, in an abstract way, stochastic phenomena, capturing the key features. Understanding the relative importance of properties of the phenomena, based on the specific decision problem, might require much more than the modeler's skills and intuition. We presented an analysis framework and measures which can help understanding which properties of uncertainty are more important to capture in a given model for decisions under uncertainty. The analysis is suitable for general stochastic programs, and especially for inherently multistage ones, where using incorrect models of uncertainty can lead to repeating poor decisions.

The analysis has been performed on a case of the maritime fleet renewal problem, which is modeled as a multistage stochastic program. The tests reported show that some properties have very little influence on the final decisions (e.g., the correlations between the random variables) while others (e.g., the mean values) can lead to noticeable increases in the expected cost if incorrectly estimated. This analysis can be of help for shipping companies when performing market analysis, data analysis and contract negotiation, in order to work out the correct value for the critical properties. However, these results are not meant to suggest neglecting properties such as the correlations. The results are in fact only valid for this specific case. The analysis framework is instead general and can support decision makers any time models of the uncertainty have to be created.

Acknowledgements

We thank Dr. Michal Kaut for providing us code and support for the utilization of the scenario generating heuristic. This work is part of the MARFLIX project partly financed by the Norwegian Council of Research. We also thank the two anonymous reviewers for their insightful comments which helped us improve the paper.

References

ALVAREZ, J. F., TSILINGIRIS, P., ENGEBRETHSEN, E. S. & KAKALIS, N. M. P. (2011) Robust fleet sizing and deployment for industrial and independent bulk ocean shipping companies, *INFOR:Information Systems and Operational Research,* **49**, 93–107.

BROADIE, M. (1993) Computing efficient frontiers using estimated parameters, *Annals of Operations Research,* **45**, 21–58.

CHOPRA, V. K. & ZIEMBA, W. T. (1993) The effect of errors in means, variances, and covariances on optimal portfolio choice, *The journal of portfolio management,* **19**, 6–11.

CHRISTIANSEN, M., FAGERHOLT, K., RONEN, D. & NYGREEN, B. (2007) Maritime transportation. In: C. BARNHART & G. LAPORTE, ed. *Handbook in Operations Research and Management Science,* Elsevier, 189–284.

DUPACOVÁ, J., CONSIGLI, G. & WALLACE, S. W. (2000) Scenarios for Multistage Stochastic Programs, *Annals of Operations Research,* **100**, 25–53.

ESCUDERO, L.F., GARÍN, A., MERINO, M. & PÉREZ G. (2007) The value of the stochastic solution in multistage problems, *TOP,* **15**, 377–398.

FLETEN,S.-E., HØYLAND,K. & WALLACE, S. W. (2002) The performance of stochastic dynamic and fixed mix portfolio models, *European Journal of Operational Research,* **140**, 37–49.

HØYLAND,K., KAUT, M. & WALLACE, S. W. (2003) A Heuristic for moment-matching scenario generation, *Computational Optimization and Applications,* **24**, 169–185.

KALLBERG, J. G. AND ZIEMBA, W. T. (1984) Mis-Specifications in Portfolio Selection Problems In: G. BAMBERG & K. SPREMANN, ed. *Risk and Capital,* Springer Berlin Heidelberg, 74–87.

KAUT, M., VLADIMIROU, H., WALLACE, S. W. & ZENIOS, S. A. (2007) Stability analysis of portfolio management with conditional value-at-risk, *Quantitative Finance,* **7**, 397–409.

KAUT, M. & WALLACE, S. W. (2007) Evaluation of Scenario-Generation Methods for Stochastic Programming, *Pacific Journal of Optimization,* **3**, 257–271.

KING, A.J. & WALLACE, S.W. (2012) *Modeling with Stochastic Programming,* New York: Springer.

PANTUSO,G., FAGERHOLT, K. & WALLACE,S. W. (2014) Uncertainty in fleet renewal: a case from maritime transportation, *Forthcoming in Transportation Science.*