

Article

An Open-Data-Driven Enhanced Bayesian Decision Network for System-Level UAV Accident-Severity Analysis and Response Simulation

Ruimin Hao , Anning Ni *, Jingbo Yin , Linjie Gao , Yutong Zhu, Xi Wang, Yizhou Wang and Xiaoning Zhang

State Key Laboratory of Ocean Engineering, School of Ocean and Civil Engineering, Shanghai Jiao Tong University, Shanghai 200240, China; haoruimin@sjtu.edu.cn (R.H.); jingboyin@sjtu.edu.cn (J.Y.); zhuja2024@sjtu.edu.cn (Y.Z.); wx0107@sjtu.edu.cn (X.W.); wang_yizhou@sjtu.edu.cn (Y.W.); zhangxiaoning@sjtu.edu.cn (X.Z.)

* Correspondence: nianning@sjtu.edu.cn

Highlights

Please indicate how your work links to systems science via your contributions to systems practice, theory, and/or methodology.

- The work treats UAV accident severity as an outcome of a coupled human–machine–environment–management system, converting heterogeneous accident narratives and matched weather data into a traceable, four-layer risk structure.
- It links systems methodology to practice by combining Leaky Noisy-OR parameter reduction, Bayesian diagnosis, and influence-path analysis with assumption-explicit simulation of alternative operational responses.

What are the main findings and/or the implications of the main findings?

- Across 633 recorded accidents, human and environmental dimensions showed the strongest model-based associations with severe consequences; airspace exposure, wind, and operation error ranked highest under severe-consequence conditioning.
- Under baseline safety-priority assumptions, the reroute strategy achieved the highest expected utility and reduced the modelled high-risk probability from 81% to 45%; the adjust strategy ranked first under alternative preference weights, making the preferred response context-dependent.

Abstract

Unmanned aerial vehicle (UAV) accidents pose growing challenges to public safety and airspace management. This study develops an open-data-driven hierarchical Bayesian network (BN) with a Leaky Noisy-OR mechanism to analyse factors associated with consequence severity among recorded UAV accidents, and extends it to a Bayesian decision network for response simulation. Using 633 public accident records and matched meteorological data, 14 binary risk-factor nodes spanning human, machine, environmental, and management dimensions were constructed. Stratified five-fold cross-validation yielded a mean validation F1 score of 0.922 and an AUC of 0.784. Backward inference ranked airspace exposure, wind, and operation error highest under severe-consequence conditioning, whereas sensitivity analysis identified wind, bad weather history, and operation error as the most influential root-node parameters. Under the assumed directed acyclic graph (DAG), the bad weather history→weather→environment→risk state path had the highest average edge-influence score (0.853). Under the baseline safety-priority assumptions, the reroute strategy was preferred, yielding the highest expected utility (31.967) and reducing



Academic Editors: Fábio Duarte and Leticia Peret Antunes Hardt

Received: 8 July 2026

Revised: 14 August 2026

Accepted: 15 August 2026

Published: 17 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

the model-estimated post-decision high-risk probability from 81% to 45%. Alternative preference settings ranked the adjust strategy first. The framework integrates open-data severity analysis with assumption-explicit response simulation.

Keywords: UAV accident severity; Bayesian network; Leaky Noisy-OR; influence-path analysis; decision simulation

1. Introduction

As a strategic emerging industry, the low-altitude economy is experiencing rapid growth. Unmanned aerial vehicles (UAVs), which serve as the core vehicles in this sector, have accelerated their transition from military applications to a wide range of civil and commercial fields [1]. According to data released by the Civil Aviation Administration of China, the number of registered UAVs in China has continued to grow rapidly, reaching approximately 3.29 million by the end of 2025, representing a year-on-year increase of approximately 51% [2]. Meanwhile, the Federal Aviation Administration (FAA) listed 837,513 registered drones in the United States on its public statistics webpage in June 2026 [3]. However, alongside the surge in operational scale and flight frequency, UAV safety incidents are also on the rise. The FAA receives over 100 reports each month of unauthorised UAV sightings near airports [4]. Global UAV accident records also indicate a growing safety-management challenge. Consequently, understanding UAV accident causation has become essential for improving flight safety. Meanwhile, the growing availability of public accident records and environmental data provides new opportunities for data-driven accident analysis.

Several methods have been developed for accident causation analysis, among which statistical analysis, fault tree analysis (FTA), and Bayesian networks (BNs) are the most widely adopted. Statistical approaches primarily rely on descriptive statistics or regression analysis to identify accident patterns and high-frequency risk factors. For example, Grindley et al. conducted a systematic human-factor analysis of 77 UAV accidents over a 12-year period and identified key causal factors [5]. However, such approaches are limited in representing the causal relationships and interactions among multiple risk factors. FTA, which combines qualitative and quantitative analysis, employs logical gates to describe the relationships among events, enabling systematic tracing of the root causes of accident events [6,7]. Nevertheless, conventional FTA assumes deterministic logical relationships and therefore cannot represent uncertainty or incomplete information in complex systems. By contrast, Khakzad et al. systematically compared FTA with BNs. Bayesian networks are capable of causal inference and uncertainty quantification with limited sample sizes, thereby addressing both the lack of causal modelling in statistical methods and the inability of FTA to handle uncertainty [8]. Consequently, Bayesian networks have been widely applied in risk analysis [9,10].

BNs have also been applied to UAV safety risk modelling across various scenarios. Wang et al. proposed a systematic BN for UAV operations, establishing causal pathways from underlying driving factors to key performance indicators such as safety, mission success, and third-party risks [11]. Sun et al. adopted the tree-augmented naive Bayes approach to learn from UAV accident data, uncovering the dependencies among risk factors and their combined effects on accident severity [12]. Ye et al. employed a common cause failure BN to identify critical risk factors of UAV accidents, finding that flights in no-fly zones, obstacles and operational errors are the primary causes [13]. Collectively, these studies demonstrate that BNs are suitable for representing the multi-factor, uncertain, and

interdependent nature of UAV accident causation. However, existing studies still rely heavily on expert judgement, limited case samples, or predefined causal assumptions, while systematic modelling based on public accident records and external environmental information remains insufficient. Moreover, although UAV accident narratives contain rich causal information, the transformation of unstructured accident descriptions into well-defined and reproducible risk nodes has not been fully addressed. From the modelling perspective, another challenge lies in the rapid growth of CPT parameters when multiple parent nodes point to the same child node, which makes parameter estimation difficult under limited accident samples [14]. Although parameter-reduction mechanisms such as Leaky Noisy-OR have been used in other fields to alleviate CPT complexity [15], their application to UAV accident risk modelling, particularly in open-data-based analysis, remains relatively limited.

Furthermore, accident-causation analysis alone is insufficient for reducing UAV operational risk. Effective safety management requires a closed-loop framework that connects risk identification with response-strategy evaluation. Bayesian decision networks (BDNs), which extend BNs by incorporating decision and utility nodes, provide a suitable tool for assessing the expected effects of alternative risk-management strategies. Xue et al. constructed a BN to analyse ship accident risk and evaluated risk improvement under different decisions using a utility matrix [16], offering a useful reference for decision-oriented safety analysis. In the UAV domain, Herashchenko et al. proposed a BDN-based decision-support method for UAV flight control [17], while other studies have evaluated UAV risk-mitigation measures within BN frameworks [18]. Nevertheless, studies that systematically integrate accident-causation analysis, risk propagation, and decision simulation in UAV safety remain limited.

These limitations indicate the need for an integrated framework that can transform publicly available UAV accident records and external environmental information into interpretable risk variables, reduce the parameter burden of BN modelling, and further support the quantitative evaluation of risk-response strategies. Against this background, the present study addresses three closely connected research questions: how public accident narratives and matched meteorological data can be structured for probabilistic accident-severity analysis; which factors and network paths show the strongest model-based associations with severe consequences; and how these results can support the comparison of response strategies under explicit decision assumptions. Accordingly, this study proposes an open-data-driven enhanced Bayesian decision network for systemic UAV accident risk analysis and response simulation. The overall research framework is illustrated in Figure 1, and the main contents are as follows:

- (1) Open-data-based risk-factor identification and node construction. Public UAV accident records and corresponding historical meteorological data are organised into a structured dataset. Key risk factors are then identified from the constructed dataset and transformed into BN risk nodes covering human, equipment, environmental, and management dimensions.
- (2) BN parameter reduction and training. The Leaky Noisy-OR mechanism is introduced to simplify the CPT under multi-parent-node conditions, thereby reducing the complexity of BN parameter estimation.
- (3) Accident-severity association and network-influence analysis. Based on the trained BNs, backward inference, sensitivity analysis, and influence-path analysis are conducted to identify the factors with the strongest conditional associations, the most locally sensitive parameters, and high-influence paths under the assumed network structure.

- (4) Decision simulation for risk-response strategies. Decision nodes and utility nodes are further incorporated into the BN to develop an enhanced Bayesian decision-network architecture. Different risk-response strategies are then simulated and compared to evaluate their impacts on accident-risk variations and expected utility, thereby providing decision-making references for UAV accident risk control.




Figure 1. The flowchart of this study.

The remainder of this paper is organised as follows. Section 2 introduces the data used in this study and the mathematical models related to the BNs. Section 3 elaborates on the construction, training, analysis, and decision simulation of the BNs. Section 4 concludes the study.

2. Data and Methods

2.1. Data Collection and Preprocessing

The UAV accident data used in this study were obtained from the Aviation Safety Network (ASN) database (<https://aviation-safety.net/database/databases.php> (accessed on 1 March 2026)). Operated and maintained by the Flight Safety Foundation (FSF), ASN compiles publicly available aviation accident records from multiple countries and regions and has been used as an open data source in aviation safety research [19]. The database of UAVs features broad coverage, a large number of cases, and rich field-level information, making it well suited for UAV accident risk analysis. Examples of the extracted raw data are presented in Table 1.

To evaluate the influence of meteorological factors on UAV accidents, the weather data corresponding to the time and location of each accident were retrieved through Open-Meteo (<https://open-meteo.com/en/docs/historical-weather-api> (accessed on 15 March 2026)). The meteorological factors acquired in this study include wind speed, precipitation, visibility, and temperature.

Following data retrieval and duplicate checking, 1258 candidate UAV accident records were obtained. Of these, 35 records were excluded because their accident narratives contained fewer than 50 characters and therefore provided insufficient information for causal-factor identification. A further 590 records were excluded because fewer than five predefined risk dimensions could be reliably identified from their narratives, preventing

systematic risk-node coding. The final dataset comprised 633 records with complete labels for the 14 risk-factor nodes and the target risk state. The accidents occurred between 26 June 2003 and 19 December 2025 and covered 25 countries and regions, including the United Kingdom, Australia, South Africa, the United States, and China.

Table 1. Sample of UAV accidents from the ASN database.

Field	Description	Example
Date	Date of accident	15 March 2022
Type	UAV model	MQ-9A Reaper
Location	Location of accident	Niamey Airport, Niger
Phase of flight	Flight phase	En route/Landing/Take-off
Nature of flight	Nature of mission	Military/Surveillance/Training
Damage	Level of aircraft damage	Destroyed/Substantial/Minor/None
Fatalities	Number of casualties	0
Narrative	Detailed description of the accident	The UAV experienced a loss of link during a routine surveillance mission. The operator attempted to regain control, but the aircraft entered an uncontrolled descent and impacted terrain.

Based on information explicitly reported in the accident records, the identifiable UAV types included multirotor, fixed-wing, single-rotor, unmanned helicopter, and hybrid or vertical take-off and landing configurations. The operational purposes included surveying and investigation, test flights, private operations, training, aerial patrol, agricultural operations, cargo transport, and demonstration or display flights. The accidents occurred during different flight phases, including en route flight, operational manoeuvring, take-off, initial climb, approach, and landing. The reported accident outcomes included substantial damage, complete destruction, minor damage, and aircraft loss. These characteristics illustrate the diversity of aircraft configurations, operational scenarios, flight phases, and accident consequences represented in the dataset.

2.2. Bayesian Network Construction

2.2.1. Fundamentals of the BN

The BN represents conditional dependencies and causal relationships among variables using a directed acyclic graph (DAG) [20]. Its core principle is to encode domain knowledge through directed causal links and quantify these relationships using conditional probabilities [21]. In this study, the BN is employed as an interpretable probabilistic framework in which, given a predefined network structure, the model parameters can be learned from limited observational data [22]. The DAG was specified a priori from the literature, accident narratives, and domain judgement, whereas its numerical parameters were learned from observational ASN accident records. Accordingly, the directed arcs represent assumed directional dependencies rather than causal effects identified from the data. Compared with purely data-driven machine learning models, this structure-guided parameter learning process offers stronger small-sample effectiveness and clearer causal interpretability, which is particularly valuable for accident analysis where data availability is often limited and model transparency is required [21,23]. Moreover, the BN supports both forward prediction and backward diagnostic reasoning, making it suitable for analysing accident scenarios involving missing information or uncertain causal mechanisms [8].

Let the BN contain a set of variables $X = \{X_1, X_2, \dots, X_n\}$. According to the chain rule and the conditional-independence assumption, the joint probability distribution can be decomposed into the product of the conditional probabilities of all nodes:

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^n P(X_i \mid Pa(X_i)) \quad (1)$$

where $Pa(X_i)$ denotes the set of parent nodes of X_i in the DAG. For a root node without any parents, $P(X_i \mid Pa(X_i))$ degenerates into the prior probability $P(X_i)$. This decomposition exploits the conditional-independence structure, simplifying the representation of a high-dimensional joint distribution into a combination of local conditional probabilities, thereby substantially reducing the number of model parameters and computational complexity. Once the states of certain variables are observed, the posterior probabilities of any unobserved variables in the network can be updated via Bayes' theorem:

$$P(X_i \mid E) = \frac{P(E \mid X_i) P(X_i)}{P(E)} \quad (2)$$

Here, $P(X_i)$ and $P(E)$ denote the prior probabilities of the variable X_i and the evidence E , $P(X_i \mid E)$ represents the posterior probability of X_i given the occurrence of E , and $P(E \mid X_i)$ represents the likelihood of E given the occurrence of X_i .

Based on the above formula, the BN integrates prior knowledge with observational evidence to dynamically update and infer the occurrence probabilities of variables in the system, thereby providing a quantitative probabilistic-reasoning foundation for subsequent accident causation analysis and risk assessment.

2.2.2. Node Selection for the BN

The selection of BN nodes is primarily based on a review of the literature on UAV safety and the causes of accidents. Existing research indicates that UAV accidents are rarely caused by a single factor; rather, they typically involve interactions between human and operational factors, technical or equipment failures, environmental conditions, and organisational or management-related factors [12,24,25]. Building on this, this study utilised these dimensions as a preliminary analytical framework for node identification and further examined the dataset of UAV accident investigation reports.

At the network structural level, drawing on the modelling paradigm of hierarchical BNs and based on the causal relationships between factors, nodes are categorised into three levels: risk factors, intermediate nodes and target nodes [26]. Specifically, the risk factors represent directly observable basic risk factors, the intermediate nodes characterise risk dimensions arising from the coupled effects of multiple basic factors, and the target node represents accident-consequence severity among ASN-recorded accidents. The introduction of intermediate nodes is intended to enhance the interpretability of the causal structure and to facilitate subsequent causal-chain analysis [27]. With respect to node-state definition, the risk levels of all nodes are quantified based on UAV accident data. In this study, binary discrete variables are employed to represent node risk levels, namely the high-risk state and the low-risk state. This binarisation effectively reduces data requirements [28].

2.2.3. Implementation of the Leaky Noisy-OR Model

In a BN, when multiple parent nodes point to the same child node, the number of parameters required by the CPT of the child node grows exponentially with the number of parents, owing to the decomposition property of the Bayesian joint probability. The Noisy-OR model is suitable for handling situations in which multiple causes lead to the same effect [29], and it is well suited to the binary-variable nodes employed in this study. Its core assumption is that each parent node independently activates the high-risk state of the child node with a certain probability, and the child node is activated if and only if at least one parent node successfully exerts an activating effect. On this basis, a leak

probability is further introduced to capture the influence of latent causes not explicitly incorporated into the model, thereby enhancing the robustness.

Let the child node Y be a binary variable with the state set $\{0,1\}$, denoting the inactive and active states, respectively, and let its parent nodes $X_1, X_2, \dots, X_n$ likewise be binary variables. In the Leaky Noisy-OR model, the activating effects of the parent nodes on Y are assumed to be mutually independent. For each parent node X_i , the event $Y_{X_i} = 1$ is defined as the activation of Y caused solely by X_i whose probability is given by:

$$P_i = P(Y_{X_i} = 1 \mid X_i = 1) \quad (3)$$

In addition, a leak probability P_L is introduced to characterise the probability that the child node Y is still activated by unknown causes when none of the known parent nodes is activated, namely:

$$P_L = P(Y = 1 \mid X_1 = 0, X_2 = 0, \dots, X_n = 0) \quad (4)$$

Consequently, the conditional probability that Y is activated, given the states of its parent nodes, is expressed as:

$$P(Y = 1 \mid X_1 = x_1, \dots, X_n = x_n) = 1 - (1 - P_L) \times \prod_{i=1}^n (1 - P_i)^{x_i} \quad (5)$$

Through this mechanism, the data input required for parameter fitting of intermediate nodes in the BNs can be substantially simplified.

In Equations (3)–(5), the parent-specific activation probabilities and the leak probability are free parameters of the Leaky Noisy-OR model rather than externally fixed expert estimates. In the final model, these parameters were randomly initialised in GeNIe Modeler and subsequently estimated jointly from the training data using the Expectation–Maximisation (EM) algorithm. The random values served only as numerical starting points for optimisation and were replaced by the learned parameter estimates after convergence.

2.2.4. Performance Evaluation of the BN

To reduce dependence on a single random split, the 633 labelled samples were divided into five stratified folds using a fixed random seed of 42. In each iteration, four folds (approximately 80%) were used for BN parameter learning, and the remaining fold (approximately 20%) was used for validation. This process was repeated five times so that each fold served once as the validation set. With the network structure predefined, the network parameters were estimated using the Expectation–Maximisation (EM) algorithm in GeNIe Modeler. The 14 basic risk-factor nodes and the target node were directly observed in the labelled accident records, whereas the intermediate nodes were treated as latent variables. The prior distributions of the fully observed root nodes were estimated from their empirical frequencies in the training data. For all non-root Leaky Noisy-OR nodes, the parent-specific activation probabilities and leak probabilities were randomly initialised in GeNIe Modeler and subsequently estimated jointly from the training data using EM. The random values served only as numerical starting points and were iteratively updated until convergence; no formally elicited expert probabilities or manually fixed leak probabilities were retained in the final model.

Considering that the accident risk is categorised into two classes, the model performance can be assessed from the confusion matrix. *Accuracy*, *Precision*, *Recall*, *F1 score*,

true positive rate (*TPR*), false positive rate (*FPR*) and the area under the receiver operating characteristic curve (*AUC*) are selected as evaluation metrics:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (6)$$

$$Precision = \frac{TP}{TP + FP} \quad (7)$$

$$TPR = Recall = \frac{TP}{TP + FN} \quad (8)$$

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (9)$$

$$FPR = \frac{FP}{FP + TN} \quad (10)$$

where, *TP*, *TN*, *FP* and *FN* denote the numbers of true-positive, true-negative, false-positive, and false-negative samples, respectively. *Accuracy* is used to measure the overall proportion of correct classifications made by the model. *Precision* reflects the proportion of samples classified as high-risk that are actually high-risk. *Recall* reflects the proportion of actual high-risk samples correctly identified by the model. *F1* balances *Precision* and *Recall*, providing a comprehensive measure of the model's performance in identifying the positive class. *TPR* is equivalent to *Recall*. *FPR* denotes the probability that low-risk samples are erroneously classified as high-risk. *AUC* measures the model's ability to distinguish between cases with severe and less severe consequences across different classification thresholds.

2.3. Bayesian Network Analysis Methods

2.3.1. Backward Inference Method

Backward inference is employed to infer, in reverse, the posterior probabilities that individual risk factors are in the high-risk state, given that the target risk event has occurred or is set to the high-risk state. Accordingly, this method can be used to identify the key risk factors that contribute to an elevated UAV accident risk.

In this study, the probability that the UAV accident is in the high-risk state is set to 100%, and the principal causal factors leading to this high-risk outcome are then identified through the results of the backward inference.

Backward inference is used to update the posterior probability that each risk factor is in its high-risk state after the target node is conditioned on the severe-consequence state. In this study, $P(R = 1)$ was set to 1, and the resulting posterior probabilities were compared with their baseline values. A higher posterior probability indicates a stronger conditional association with the severe-consequence state within the assumed BN structure; it does not identify a causal effect or establish that the factor caused the recorded accident.

2.3.2. Node Sensitivity Analysis Method

Sensitivity analysis is employed to evaluate the extent to which probability variations in factor nodes affect the target node within a BN. For certain factors, even slight changes in their probabilities can induce substantial variations in the probability of the target node. The sensitivity coefficient is obtained by taking the partial derivative of the posterior probability of the target node with respect to each numerical parameter in the network, and is formulated as:

Sensitivity analysis was performed in GeNIe with the severe-consequence state of the target node *R* as the target outcome and with no additional evidence. For each binary root risk-factor node *X_i*, the varied parameter was $p_i = P(X_i = 1)$. The complementary

probability was adjusted as $P(X_i = 0) = 1 - p_i$, while all other root-node priors and CPT parameters were held at their learned baseline values. The local sensitivity coefficient was evaluated at the learned baseline value of p_i and is defined as:

$$S_i = \left| \frac{\partial P(R = 1)}{\partial p_i} \right|_{p_i = \hat{p}_i}, \quad p_i = P(X_i = 1), \quad P(X_i = 0) = 1 - p_i \quad (11)$$

where S_i measures the local rate of change in the model-estimated probability $P(R = 1)$ with respect to p_i , and $\hat{p}_i$ is the learned baseline value. A larger S_i indicates greater local parameter sensitivity within the fitted BN.

2.3.3. Influence-Path Analysis Based on Edge-Strength Method

The Strength of Influence function in GeNIe was used in Static mode with the Maximum option and Hellinger distance. For a directed edge $X \rightarrow Y$, let $Z = Pa(Y)$ denote the other parent nodes of Y , and let $\mathcal{Z}$ denote the set of all possible state configurations of Z . During each comparison, the other parent nodes are held at the same configuration z . The edge-influence strength is the maximum Hellinger distance across these configurations:

$$I_{max}(X \rightarrow Y) = \max_{z \in \mathcal{Z}} \left\{ \frac{1}{\sqrt{2}} \left[\sum_{y \in \{0,1\}} \left(\sqrt{P(Y = y | X = 1, Z = z)} - \sqrt{P(Y = y | X = 0, Z = z)} \right)^2 \right]^{\frac{1}{2}} \right\} \quad (12)$$

The Maximum option retains the largest local change in the conditional distribution of the child node rather than diluting a strong connection through averaging. Hellinger distance is symmetric and bounded between 0 and 1. For a complete directed path $\pi = (e_1, e_2, \dots, e_m)$, the path-level average edge-influence score is calculated as:

$$S_{path}(\pi) = \frac{1}{m} \sum_{k=1}^m I_{max}(e_k) \quad (13)$$

The arithmetic mean removes the direct effect of different path lengths and retains the same 0–1 scale as the individual edge values. The resulting score is a descriptive path-comparison measure under the assumed DAG, rather than a joint propagation probability or an identified total causal effect.

2.3.4. Decision Simulation Analysis Method

In UAV risk decision-making, the introduction of a utility node into the BN enables the optimal decision to be selected based on the utility values associated with different decision actions.

The two states of UAV accident risk are denoted as $R = \{R_1, R_2\}$. Let there be n decision options $F = \{F_1, F_2, \dots, F_n\}$. For each combination of a risk state and a decision, the corresponding utility value is denoted as $U(R_i, F_j)$, and these values together form the decision utility matrix:

$$U = \begin{bmatrix} U(R = 0, F_1) & U(R = 0, F_2) & \dots & U(R = 0, F_n) \\ U(R = 1, F_1) & U(R = 1, F_2) & \dots & U(R = 1, F_n) \end{bmatrix} \quad (14)$$

Once an evidence E is observed, the BN updates the posterior probabilities of the risk states accordingly. For the j decision, the expected utility can be expressed as:

$$EU(F_j|E) = \sum_{i=1}^2 P(R_i|E, F_j) U(R_i, F_j) \quad (15)$$

where $EU(F_j|E)$ denotes the expected utility of adopting the decision option F_j given the evidence E . $P(R_i|E, F_j)$ represents the posterior probability that the risk state is R_i when the evidence E is observed and the decision option F_j is taken. $U(R_i, F_j)$ is the utility value corresponding to the combination of the given risk state and decision option. The optimal decision is the one that yields the maximum expected utility.

In summary, the UAV risk decision system can be formulated as $D = \{R, F, U\}$. Through Bayesian inference and expected-utility calculation, the system enables a quantitative comparison among alternative risk-response strategies under risk uncertainty, thereby providing quantitative support for UAV accident prevention and safety-management decision-making.

Because the attribute scores, attribute weights, risk-impact factors, and post-decision transition probabilities are decision-scenario parameters rather than quantities estimated entirely from the 633 accident records, a compact parameter-sensitivity analysis was conducted. Five weight schemes were first examined: the baseline safety-priority scheme, equal weights, and cost-, time-, and mission-priority schemes. A one-at-a-time (OAT) analysis was then performed. The baseline setting itself was not included in these counts. In every OAT case, only one parameter was varied whilst all remaining parameters were held at their baseline values, and the expected utilities and rankings of the five strategies were recalculated.

3. Results and Discussion

3.1. Node Selection and BN Construction

Based on the narrative texts of valid UAV accident records in the ASN database, 14 nodes were ultimately determined as the basic risk factors of the BNs through item-by-item review and coding, covering four dimensions: human factors, machine factors, environmental factors, and organisational/management factors. In addition, to provide the conditions required for supervised parameter learning of the BNs, the final risk level of each accident was also annotated. The nodes and the definition of the high-risk states are summarised in Table 2. All 633 records were first coded according to the binary criteria in Table 2 and were then independently re-coded by a second researcher using the same criteria. A node was labelled as high risk only when explicit evidence satisfying the corresponding criterion was identified in the accident narrative or the matched external data; otherwise, including when the available description was ambiguous or insufficient, it was labelled as low risk. Consistent with the predefined binary structure of the BN, no separate unknown state was introduced. Accordingly, the low-risk state denotes that no explicit high-risk evidence was identified in the available record, rather than that the underlying factor was confirmed to be absent.

As the accident records mainly consist of unstructured narratives, many risk factors can only be reliably identified as present or absent, rather than consistently classified into multiple risk levels. Binary coding reduces subjective labelling errors and improves consistency across heterogeneous accident records.

The ASN narratives and the time- and location-matched Open-Meteo observations were used as complementary evidence for the same binary nodes. The wind node was restricted to wind speed, gusts, and turbulence, whereas bad weather history represented non-wind adverse conditions, including extreme temperature, precipitation, snowfall, icing, and low visibility. The term history refers to archived meteorological observations matched to the accident time and location rather than to long-term climatic history. When the same meteorological phenomenon was identified in both sources, it was recorded as a single binary activation and was not counted twice. The two nodes were simultaneously labelled as high risk only when distinct wind and non-wind adverse conditions were present.

Table 2. Node names and the definition of the high-risk state.

Node	Definition of the High-Risk States
special ops	The mission is non-routine in nature, such as training, test flight, verification, demonstration, or experimental operation.
operation error	A definite procedural deviation, checklist omission, or critical step error by the pilot, operator, or ground control station is identified and causally linked to the accident.
airspace exposure	The operational environment involves airports, runways, approach areas, controlled airspace, manned aircraft routes, public roads, or populated areas.
coordination failure	Communication breakdown, handover error, coordination failure, or loss of monitoring occurred among the pilot, visual observer, ground station, or command personnel.
power loss	A malfunction or performance degradation of the powerplant, motor, propulsion system, or power supply chain resulted in degraded flight control or crash.
battery issue	A battery fault, voltage anomaly, or insufficient charge triggered an automatic return-to-home or forced landing, and is associated with the accident.
power software issue	A firmware defect, power-management software anomaly, battery management error, motor/ESC control parameter misconfiguration, or power-protection logic failure is related to the accident.
sensor issue	Anomalous signals, magnetic interference, or positioning deviation of GNSS, compass, inertial measurement unit, or other sensors is associated with the accident.
link issue	The command-and-control link, data transmission, telemetry, or radio communication was disrupted by electromagnetic interference or signal loss, resulting in loss of contact or controllability.
wind	Wind speed, gusts, or turbulence adversely affected the aircraft attitude, airspeed, return-to-home, or landing, or the measured wind reached the operational limit of the aircraft type.
bad weather history	Adverse meteorological conditions such as extreme temperature, precipitation, snowfall, icing, or low visibility were present during the accident period.
obstacles terrain	Wires, trees, buildings, towers, cranes, or complex terrain were present in the operational space and directly related to collision risk or operational constraints.
planning process issue	Evident deficiencies or omissions existed in risk assessment, pre-flight inspection, maintenance support, procedural execution, parameter review, or mission planning.
authorisation issue	The operation involved special permits, waivers, or operational authorisations, or constituted an unauthorised operation in violation of regulations.
risk state	The aircraft damage is destroyed, aircraft missing, or written off. And if fatalities/other fatalities > 0, or if the narrative clearly indicates serious injury, amputation, hospitalisation, major injury, fire with severe consequences, etc.

Following the modelling paradigm of hierarchical BNs [30], two layers of intermediate nodes are introduced between the 14 risk factors and the target node, forming a four-layer BN for UAV accident risk (Figure 2). This hierarchical structure clarifies causal relationships and supports the subsequent causal-chain analysis. The first intermediate layer comprises five nodes: operator risk, operation execution, power failure, control failure and weather conditions. The second intermediate layer consolidates subsystem risks into four risk dimensions: the human dimension, comprising operator risk and operation errors, the machine dimension, comprising UAV power and control failures, the environmental dimension, comprising weather conditions and obstacles terrain, and the management dimension, comprising planning processes and authorisation issues. These dimensional nodes jointly point to the target node risk state. This hierarchical design reduces the complexity of Leaky Noisy-OR parameter estimation and makes the relationships among risk factors easier to trace, which is useful for the subsequent causal-chain analysis.




Figure 2. Architecture of the BNs.

In this study, GeNIe Modeler (5.0.4830) is employed as the modelling and inference platform for the BNs, which is a mature software for probabilistic graphical models, supporting visual node construction, structural editing, parameter learning, exact inference, and sensitivity analysis [31]. It has been widely applied to risk-assessment studies in fields such as aviation safety and reliability engineering [32,33].

3.2. Data Labelling and Network Training Performance

Using the criteria in Table 2, two researchers independently coded all 633 retained records. Before reconciliation, percent agreement across the 14 risk-factor nodes ranged from 88.0% to 98.7% (mean 93.5%). Agreement for the target risk-state node was 97.0%. Each disagreement was re-examined against the original ASN narrative and the matched meteorological record. The final label was assigned by applying the predefined evidence rule: a high-risk state required explicit supporting evidence, whereas unresolved or ambiguous cases were assigned to the low-risk state. The resulting reconciled labels were used for all parameter learning and subsequent analyses. The empirical high-state frequency of every observed node is reported in Table 3.

As shown in Table 3, 79% of the 633 accidents were classified as severe-consequence cases. The risk factors exhibit a clearly stratified distribution. Airspace exposure (35%), wind (28%), operation error (26%), special ops (25%), and obstacles terrain (25%) form a

high-frequency cluster, each labelled as high-risk in at least one quarter of the accident records, which indicates their prevalence in UAV accidents. Power software issue (16%) and power loss (16%) constitute a mid-frequency group, showing that equipment-related failures, although not the most common feature, still occur at a non-trivial rate. Coordination failure (1%) has the lowest prior probability, indicating that multi-operator coordination failure is a low-probability event in the current dataset.

Table 3. Frequency of the high-risk state for each node.

Node	Prior Probability (%)	Node	Prior Probability (%)
risk state (target)	79	sensor issue	13
airspace exposure	35	bad weather history	12
wind	28	planning process issue	10
operation error	26	battery issue	9
special ops	25	authorisation issue	7
obstacles terrain	25	link issue	3
power software issue	16	coordination failure	1
power loss	16		

After completion of the stratified five-fold cross-validation, the final BN was retrained on all 633 labelled records using the same predefined network structure and EM learning procedure. For this final model, the prior distributions of the 14 fully observed root risk-factor nodes were estimated from their empirical frequencies in the full dataset, whereas the parent-specific activation probabilities and leak probabilities of all non-root Leaky Noisy-OR nodes were randomly initialised and subsequently estimated jointly using the EM algorithm in GeNIe Modeler. The random initial values were used only to initiate optimisation, and no expert-specified CPT entries or fixed leak probabilities were retained. The parameters reported in Table 4 and all subsequent inference, sensitivity, influence-path, and decision analyses are based on this final full-sample model. Its unrounded target-node marginal probability was approximately 0.806, displayed as 81% in Figure 3 and used in the decision simulation.




Figure 3. The marginal probability distribution of every node after parameter learning.

Table 4. EM-learned Leaky Noisy-OR parameters of the final Bayesian network.

Child Node	EM-Learned Parent-Specific Activation Probabilities	EM-Learned Leak Probability
Operator risk (N1)	Special ops: 0.83, Airspace exposure: 0.84	0.05
Operation execution (N2)	Operation error: 0.80, Coordination failure: 0.59	0.03
Power failure (N3)	Power loss: 0.18, Battery issue: 0.19, Power software issue: 0.81	0.08
Control failure (N4)	Sensor issue: 0.22, Link issue: 0.60	0.08
Weather (N5)	Wind: 0.86, Bad weather history: 0.97	0.04
Human (Y1)	Operator risk: 0.91, Operation execution: 0.93	0.01
Machine (Y2)	Power failure: 0.80, Control failure: 0.59	0.07
Environment (Y3)	Weather: 0.77, Obstacles terrain: 0.96	0.01
Management (Y4)	Planning process issue: 0.25, Authorisation issue: 0.42	0.03
Risk state (R)	Human: 0.98, Machine: 0.79, Environment: 0.95, Management: 0.24	0.01

It is important to note that out of 633 cases, 500 were labelled as severe consequences (approximately 79%). This empirical proportion differs from the marginal probability of the fitted target node in the final Bayesian network. Due to the inclusion of potential intermediate nodes and the constrained Leaky Noisy-OR relationship estimated by EM, its fitted target node marginal probability is approximately 81%. This slight difference reflects model fit, rather than variations in sample composition or target labels.

As a brief robustness check, all EM-learned leak probabilities in the final model were jointly scaled to 0.5 and 1.5 times their baseline values while the root-node priors and parent-specific activation probabilities were held constant. Airspace exposure, wind, and operation error remained the three leading posterior factors, and bad weather history → weather → environment → risk state remained the highest-ranked influence path; its average edge-influence score ranged from 0.833 to 0.878.

The 633 labelled samples were evaluated using stratified five-fold cross-validation. The mean training and validation results across the five folds are summarised in Table 5.

Table 5. Mean performance of the BN under stratified five-fold cross-validation.

Dataset	Accuracy	Precision	Recall	F1	TPR	FPR	AUC
Training folds (mean)	0.8893	0.9084	0.9557	0.9315	0.9557	0.3565	0.8382
Validation folds (mean)	0.8709	0.8863	0.9599	0.9216	0.9599	0.4464	0.7838

Across the five validation folds, the model achieved mean values of 0.8709 for accuracy, 0.8863 for precision, 0.9599 for recall and TPR, 0.9216 for F1 score, 0.4464 for FPR, and 0.7838 for AUC. The high recall indicates a strong ability to identify severe-consequence cases, whereas the relatively high FPR suggests a tendency to classify some less-severe-consequence cases as severe. Compared with the original single random split, these results provide a more reliable estimate of the model's generalisation performance within the current dataset. The marginal probability distribution of every node after parameter learning is shown in Figure 3. The full dataset contained 500 severe-consequence and 133 less-severe-consequence records. An always-high-risk majority-class baseline therefore achieved accuracy and precision of 0.7899, recall and TPR of 1.0000, F1 of 0.8826, FPR of 1.0000, and AUC of 0.5000. The BN improved accuracy, precision, F1, and AUC while reducing FPR to 0.4464.

At the dimension level, human attains the highest high-risk probability (56%), followed by environment (48%), machine (28%), and management (8%), indicating that human and environmental factors are the dimensions with the highest inferred high-state probabilities within the recorded-accident model. At the subsystem level, operator risk (49%) exceeds operation execution (24%), suggesting that the mission context and airspace conditions

surrounding the operator contribute more to human-factor risk than operation errors themselves. Weather (35%) surpasses the direct contribution of obstacles terrain (25%), confirming that meteorological conditions are the primary source of environmental risk. Within the machine dimension, power failure (23%) is higher than control failure (11%), reflecting that propulsion- and power-supply-related faults occur more frequently than flight-control and communication faults.

3.3. Backward Inference Analysis

When backward inference is performed, the target node risk state is fixed in its high-risk state ($P = 100\%$) as the observed evidence, and Bayes' theorem is applied to update the posterior probability of each risk factor being in its high-risk state. A higher posterior probability indicates a stronger conditional association with the severe-consequence state within the fitted BN; it does not identify a causal effect [34]. The backward inference results are shown in Figure 4 and Table 6.




Figure 4. Results of backward inference.

Table 6. Comparison between prior and posterior probabilities.

Node	Prior (%)	Posterior (%)	Absolute Increase (%)	Relative Increase (%)
airspace exposure	35	40	5	14.3
wind	28	33	5	17.9
operation error	26	30	4	15.4
special ops	25	29	4	16
obstacles terrain	25	29	4	16
power software issue	16	18	2	12.5
power loss	16	16	0	0
bad weather history	12	14	2	16.7
sensor issue	13	13	0	0
battery issue	9	10	1	11.1
planning process issue	10	10	0	0
authorisation issue	7	7	0	0
link issue	3	4	1	33.3
coordination failure	1	1	0	0

As shown in Table 6, airspace exposure has the highest posterior probability when the target node is conditioned on the severe-consequence state, reaching 40% and increasing by 5 percentage points from its prior of 35%. It is therefore the highest-ranked factor by posterior probability under severe-consequence conditioning among the ASN-recorded cases, rather than an identified cause of accident occurrence or severity. The elevated posterior probability indicates that airspace exposure near airports, in controlled airspace, or over populated areas is conditionally associated with severe consequences within the fitted BN. Wind and operation error rank second and third, with posterior probabilities of 33% and 30%, respectively, indicating similarly strong, although lower, conditional associations with the severe-consequence state.

From the perspective of probability increase, airspace exposure and wind both exhibit an absolute rise of 5%, reflecting the strongest activation under accident conditions. Operation error, special ops, and obstacles terrain each rise by 4% and form the second tier. Notably, link issue shows a relative increase of 33.3%, indicating that although communication-link disruption is a low-probability event, its occurrence is the most strongly conditioned on accident outcomes. Bad weather history exhibits a relative increase of 16.7%, likewise demonstrating that severe weather, while infrequent in routine operations, is markedly more associated with accidents when it does occur.

At the dimension level, the human dimension shows the largest aggregate increase in posterior probability of 12%, followed closely by the environment dimension with an aggregate increase of 10%. This confirms that human and environmental factors are the dominant sources of UAV accident risk. By contrast, the machine dimension exhibits posterior probabilities that remain close to its priors, indicating a limited marginal contribution from equipment-related failures under accident conditions. This does not imply that machine factors are unimportant. Rather, it reflects that machine failures tend to occur as relatively independent stochastic events with weak conditional dependence on the target node.

Overall, the backward inference results indicate that UAV accidents arise from the coupled action of multiple factors. Airspace exposure, wind, and operation error constitute the three factors most strongly associated with the severe-consequence state and should receive priority in subsequent risk-mitigation strategies. Special ops and obstacles terrain act as additional factors showing notable conditional associations, and their management warrants comparable attention.

3.4. Node Sensitivity Analysis

Sensitivity analysis evaluates the marginal effect of changes in a factor's risk level on the model-estimated probability of severe consequences, providing more actionable guidance for prioritising risk interventions [35]. The sensitivity results of the risk factor nodes are shown in Figure 5 and Table 7.

Table 7. Sensitivity coefficients of the risk factors with respect to the target node.

Node	Sensitivity	Node	Sensitivity
wind	0.184	power software issue	0.110
bad weather history	0.183	link issue	0.050
operation error	0.182	battery issue	0.017
obstacles terrain	0.177	power loss	0.017
airspace exposure	0.177	authorisation issue	0.016
special ops	0.150	planning process issue	0.011
coordination failure	0.123	sensor issue	0.009

As shown in Table 7, wind has the highest sensitivity coefficient (0.184), followed by bad weather history (0.183) and operation error (0.182), indicating that changes in

meteorological and human operational factors exert the strongest marginal influence on model-estimated severe-consequence probability. Obstacles terrain (0.177) and airspace exposure (0.177) are tied for fourth place and likewise belong to the high-sensitivity group. The lowest sensitivity coefficients are observed for sensor issue (0.009), planning process issue (0.011), and authorisation issue (0.016), suggesting that within the current network structure these factors exert only a limited marginal effect on the target node. Low sensitivity, however, does not imply that such factors are unimportant. This pattern more likely reflects the attenuation of low-frequency signals along relatively long hierarchical propagation paths.




Figure 5. Sensitivity analysis of the risk factors.

It is worth noting that the sensitivity ranking does not fully coincide with the posterior-probability ranking obtained from backward inference, and this discrepancy carries meaningful implications for risk management. Coordination failure has the lowest posterior probability in backward inference yet attains a sensitivity coefficient of 0.123, and bad weather history has a prior probability of only 12% but a sensitivity coefficient as high as 0.183. These two factors exhibit low base rates in the current accident data, but once they switch from a low-risk to a high-risk state, they exert a substantial marginal effect on the system-level accident probability. Such factors are easily overlooked in routine operations, yet they can rapidly shift the overall risk state once they emerge, which gives them importance for assessing potential consequence severity [36]. Conversely, airspace exposure ranks first in backward inference but falls into the joint fourth-ranked group in the sensitivity analysis, with a sensitivity coefficient of 0.177. This indicates that, although the factor exhibits the highest posterior probability when conditioned on the severe-consequence state, its already high baseline level (35%) causes the marginal effect of further changes to approach saturation.

3.5. Influence-Path Analysis Based on Edge Strength

To characterise the local probabilistic connections along the assumed DAG, the maximum Hellinger-based influence strength of every directed edge was calculated as defined in Section 2.3.3. The results are visualised in Figure 6, where arrow thickness represents the corresponding local edge-influence strength.




Figure 6. The influence strengths of all directed edges in the BN.

For every complete directed path from a risk-factor node to the target node, the arithmetic mean of its edge-influence strengths was calculated as the path-level average edge-influence score. This length-normalised measure allows paths containing different numbers of edges to be compared on the same scale. The paths are ranked in descending order in Table 8.

Table 8. Path-level average edge-influence score of each directed path.

Path	Average Edge-Influence Score
bad weather history → weather → environment → risk state	0.853
wind → weather → environment → risk state	0.791
operation error → operation execution → human → risk state	0.789
airspace exposure → operator risk → human → risk state	0.764
special ops → operator risk → human → risk state	0.762
obstacles terrain → environment → risk state	0.761
coordination failure → operation execution → human → risk state	0.759
power software issue → power failure → machine → risk state	0.645
link issue → control failure → machine → risk state	0.535
power loss → power failure → machine → risk state	0.483
battery issue → power failure → machine → risk state	0.482
sensor issue → control failure → machine → risk state	0.410
authorisation issue → management → risk state	0.326
planning process issue → management → risk state	0.283

The arrow (→) indicates the direction of the path.

The seven highest-ranked paths belong to the environmental or human dimensions and have average edge-influence scores above 0.75. The path bad weather history → weather → environment → risk state has the highest score (0.853), indicating that it contains the strongest maximum local edge influences on average among the evaluated paths. Within the machine dimension, the power software issue path ranks highest (0.645), while the two management paths have the lowest scores (0.326 and 0.283). These differences reflect the learned CPTs and the assumed DAG and should not be interpreted as total causal effects or physical propagation probabilities.

The mean score of the seven environmental and human paths is 0.783, compared with 0.452 for the machine and management paths. Under the current structure and

learned parameters, weather and operation execution may therefore be treated as candidate intermediate nodes for closer monitoring or scenario-based assessment of potential safety barriers. This priority is structure- and model-dependent and does not demonstrate that interventions at these nodes will causally reduce consequence severity.

3.6. Decision Simulation Analysis

After the causal analysis of UAV accidents, this study introduces decision and utility nodes to construct a BDN [37], which translates the risk-assessment results into actionable response strategies and provides an interpretable decision-support basis for UAV safety management. The purpose of the decision simulation is not to estimate the absolute economic value of each response action, but to compare the relative suitability of alternative strategies under a unified probabilistic and utility-based framework. Drawing on real UAV operational scenarios and prior studies [38,39], five typical decision alternatives are defined:

- (1) Continue, in which the current flight mission is maintained without additional measures.
- (2) Adjust, in which flight speed is reduced and flight altitude is modified to adapt to the current risk condition.
- (3) Reroute, in which the planned trajectory is altered to avoid the identified high-risk areas.
- (4) Hover, in which the UAV remains stationary in the air and the mission is suspended until the risk condition is alleviated.
- (5) Return, in which the UAV returns to base or performs an emergency landing, terminating the current mission to maximise vehicle safety.

These five alternatives span the full response spectrum from maintaining the original mission to terminating the flight. They therefore represent different levels of intervention intensity and can be matched to UAV operations under different risk states.

To quantify the overall performance of each alternative, this study constructs a normalised four-dimensional utility evaluation scheme. Each attribute is scored within the interval $[-100, 100]$, where positive values denote relative benefits and negative values denote relative costs or losses. The four attributes are defined as follows:

- (1) Economic cost, a negative-oriented indicator capturing the additional expenditure of a decision, including battery consumption, equipment depreciation, and operational expenses.
- (2) Time impact, a negative-oriented indicator measuring the time loss caused by a decision, including mission delay, waiting, and detour duration.
- (3) Safety, a positive-oriented indicator evaluating the contribution of a decision to reducing the risk of crash, collision, loss of control, and personal injury.
- (4) Mission impact, a bidirectional indicator reflecting the overall effect of a decision on mission completion and continuity.

To reduce arbitrariness in the utility specification, the utility scores are assigned according to ordinal consistency rules derived from UAV operational logic. Specifically, strategies that terminate or suspend the mission, such as return and hover, are assigned higher safety benefits but larger penalties in economic cost, time impact, and mission continuity. Strategies that preserve mission continuity, such as continue and adjust, impose lower operational costs but provide weaker safety improvement. Reroute represents an intermediate intervention, as it avoids high-risk areas while maintaining partial mission continuity. Therefore, the utility matrix reflects structured trade-offs among safety, mission completion, and operational burden rather than isolated subjective preferences. To distinguish empirical learning from scenario specification, the probabilistic parameters of the original BN were estimated from the 633 recorded accidents. By contrast, the attribute

scores, baseline attribute weights, risk-impact factors, and post-decision CPT used in the decision layer were manually specified as baseline scenario assumptions. The accident records do not contain counterfactual outcomes under alternative response actions and therefore cannot identify the intervention effects of reroute, adjust, hover, or return. Accordingly, the decision results are interpreted as scenario comparisons under a consistent set of assumptions, rather than as causal intervention effects validated by the accident data. The resulting baseline attribute scores are summarised in Table 9.

Table 9. Utility scores of the five decision alternatives across the four attributes.

Decision	Economic Cost	Time Impact	Safety	Mission Impact
continue	−10	0	−70	60
adjust	−25	−15	55	55
reroute	−45	−35	80	45
hover	−55	−60	65	−70
return	−75	−70	90	−95

Following the multi-attribute utility aggregation method in management science proposed by Pidd [40], and considering the safety-priority principle in UAV operation, the four attributes are weighted as follows. The baseline weights were specified to represent a safety-priority decision preference and were normalised to sum to 1. These values were treated as baseline modelling assumptions rather than universally applicable weights, and their influence on the strategy ranking was examined through the parameter-sensitivity analysis. Safety receives the largest weight of 0.40, reflecting its central role in UAV operational decisions. Mission impact is weighted at 0.25, capturing the importance of mission completion and continuity. Economic cost is weighted at 0.20, and time impact is weighted at 0.15.


The general utility of each alternative is then computed as the weighted sum of its attribute scores. The resulting values are −15.00 for continue, 28.50 for adjust, 29.00 for reroute, −11.50 for hover, and −13.25 for return.

Because the practical effect of a given alternative depends on the current risk state, two risk impact factors corresponding to the high-risk and low-risk states are introduced to correct the general utility [41]. The risk impact factor is a contextual correction coefficient that characterises the extent to which a decision's general utility is amplified or attenuated under a specific risk state. Under high-risk conditions, active risk-reduction strategies become more valuable because their safety benefits are more pronounced, whereas continue is penalised because maintaining the original flight plan may further increase exposure to hazardous conditions. Under low-risk conditions, highly disruptive strategies such as hover and return become less attractive because their additional safety benefit may not compensate for the losses in time, cost, and mission continuity. For example, the negative utility of continue is amplified in the high-risk state because maintaining the mission under elevated risk substantially increases potential losses, whereas the negative utility of return is attenuated in the high-risk state because the safety benefit of returning becomes more pronounced and its economic and time costs become relatively acceptable. The risk impact factors and corrected utilities of the five alternatives are listed in Table 10.

A decision node, a post-decision risk state node, and a utility node were added to the original BN to form a BDN with decision-simulation capability, as illustrated in Figure 7. The post-decision risk state is jointly influenced by the original risk state and the chosen decision, and the utility node is jointly determined by the decision and the post-decision risk state.

Table 10. Risk impact factors and corrected utilities of the five decision alternatives under high- and low-risk states.

Decision	General Utility	High-Risk Impact Factor	Low-Risk Impact Factor	High-Risk Utility	Low-Risk Utility
continue	−15	1.6	0.8	−24	−12
adjust	28.5	1.15	0.95	32.775	27.075
reroute	29	1.35	0.9	39.15	26.1
hover	−11.5	0.7	1.25	−8.05	−14.375
return	−13.25	0.6	1.4	−7.95	−18.55

**Figure 7.** BDN for UAV risk decision simulation.

To make the decision-simulation block reproducible, Table 11 reports the complete conditional probability table of the post-decision risk-state node for every combination of the original risk state and decision alternative. These probabilities were manually specified as baseline intervention-scenario assumptions rather than learned from the 633 accident records, because the observational data do not contain counterfactual outcomes under alternative response actions. The CPT satisfies two ordinal constraints. First, for every decision, the probability of a high post-decision state is greater when the original state is high than when it is low. Second, under both original states, the residual high-state probability decreases in the order continue, adjust, reroute, hover, and return.

Table 11. Conditional probability table of the post-decision risk state.

Original Risk State	Decision	P(Post High)	P(Post Low)
High	Continue	0.95	0.05
High	Adjust	0.70	0.30
High	Reroute	0.50	0.50
High	Hover	0.42	0.58
High	Return	0.25	0.75
Low	Continue	0.58	0.42
Low	Adjust	0.44	0.56
Low	Reroute	0.24	0.76
Low	Hover	0.21	0.79
Low	Return	0.09	0.91

The corrected utilities in Table 10 and the conditional probabilities in Table 11 were entered into GeNIe Modeler to complete the BDN. Using the built-in expected-utility algo-

rithm, the expected utility of each alternative was calculated. The calculations used the unrounded marginal probability of the original high-risk state (approximately 0.806, reported as 81% after rounding). For each decision, the marginal post-decision high-risk probability was obtained by combining the two original-state-specific probabilities in Table 11 with this marginal probability. The resulting expected utilities, marginal post-decision high-risk probabilities, and changes relative to the original probability are summarised in Table 12 and Figure 8. These post-decision probabilities should be interpreted as scenario outcomes conditional on the specified intervention assumptions rather than as causal risk reductions independently validated by the accident data.

Table 12. Expected utilities and changes in post-decision high-risk probability for the five alternatives.

Decision	Expected Utility	Post High Risk	Risk Change
reroute	31.967	45%	−36%
adjust	30.778	65%	−16%
hover	−11.976	38%	−43%
return	−16.229	22%	−59%
continue	−22.539	88%	7%




Figure 8. Decision simulation results of UAV risk.

In terms of expected utility, reroute ranks first with a value of 31.967 and reduces the high-risk probability from 81% to 45%, indicating that it achieves the most favourable balance among safety improvement, mission preservation, and operational cost under the adopted utility assumptions. Adjust ranks second with an expected utility of 30.778, providing moderate reduction in severe-consequence probability with limited disruption to mission continuity, and is therefore suitable for moderate-risk situations. Hover and return both reduce risk more substantially, lowering the high-risk probability to 38% and 22%, respectively; however, their negative expected utilities indicate that mission interruption, time loss, and operational cost offset their safety benefits. Return is thus more appropriate for emergency scenarios in which risk control is prioritised over mission completion. Continue performs worst, with an expected utility of −22.539 and an increased high-risk probability of 88%, suggesting that taking no action may further deteriorate safety conditions.

To examine whether the decision ranking depends on the relative preference assigned to the four utility attributes, five weight schemes were evaluated. In the order of economic

cost, time impact, safety, and mission impact, the weight vectors were: baseline safety-priority, (0.20, 0.15, 0.40, 0.25); equal weights, (0.25, 0.25, 0.25, 0.25); cost-priority, (0.40, 0.20, 0.20, 0.20); time-priority, (0.20, 0.40, 0.20, 0.20); and mission-priority, (0.20, 0.20, 0.20, 0.40). The corresponding expected utilities are reported in Table 13.

Table 13. Expected utilities under alternative attribute-weight scenarios.

Decision	Baseline Safety Priority	Equal Weights	Cost Priority	Time Priority	Mission Priority
Continue	−22.539	−7.513	−9.016	−6.010	12.021
Adjust	30.778	18.899	9.719	11.879	26.998
Reroute	31.967	12.401	0.000	2.205	19.842
Hover	−11.976	−31.242	−36.449	−37.490	−39.573
Return	−16.229	−45.931	−55.117	−53.892	−60.016

A local one-at-a-time (OAT) sensitivity analysis was also conducted for the remaining decision-layer parameters. In each case, only one parameter was changed whilst all other parameters were held at their baseline values. All 20 attribute scores in Table 9 were evaluated under −10% and +10% changes, producing 40 cases. The time-impact score of continue is zero, and its two directional calculations therefore remained unchanged; these were retained as two baseline-equivalent cases to preserve the same two-direction procedure for every score. The 10 risk-impact factors in Table 10 were each varied by $\pm 10\%$, producing 20 cases, and the 10 independent P(post high) values in Table 11 were each varied by ± 0.10 , producing another 20 cases. The baseline setting itself was not included in these counts. The results are summarised in Table 14.

Table 14. Local parameter sensitivity of the decision ranking.

Parameter Block	Perturbation	OAT Cases	Reroute First	Adjust First	Range of EU(Reroute)-EU(Adjust)
Attribute scores	All 20 scores at −10% and +10%	40	36 (90.0%)	4 (10.0%)	−2.338 to 4.717
Risk-impact factors	Each of 10 factors at $\pm 10\%$	20	17 (85.0%)	3 (15.0%)	−0.940 to 3.318
Post-decision CPT	Each of 10 P(post high) at ± 0.10	20	20 (100%)	0	0.137 to 2.241

Table 13 shows that the first-ranked strategy depends on the attribute-weight assumptions. Reroute ranks first only under the baseline safety-priority scheme, whereas adjust ranks first under the equal-weight, cost-priority, time-priority, and mission-priority schemes. Thus, the small baseline advantage of reroute over adjust is not preserved under all decision preferences. As shown in Table 14, reroute ranks first in 36 of the 40 attribute-score cases and in 17 of the 20 risk-impact-factor cases, whilst adjust ranks first in the remaining four and three cases, respectively. By contrast, reroute remains first-ranked in all 20 post-decision CPT perturbation cases, with its expected-utility advantage over adjust ranging from 0.137 to 2.241.

Overall, reroute is locally robust to the tested post-decision CPT variations and remains first-ranked under most attribute-score and risk-impact-factor perturbations, but its ranking is sensitive to the attribute weights. It should therefore be interpreted as the preferred strategy under the baseline safety-priority scenario rather than as a universally optimal response. Return remains an important emergency option when safety is assigned an overriding priority.

4. Concluding Remarks

To address the multifactor interactions among risk factors and the difficulty of using unstructured accident narratives directly for quantitative modelling, this study develops a BN framework for UAV accident-severity analysis and decision simulation. Drawing on real UAV accident records from the ASN database and integrating historical meteorological data, the framework extracts risk factors from accident narratives and constructs a hierarchical BN covering the human, machine, environmental, and management dimensions. Backward inference, sensitivity analysis, influence-path analysis, and decision simulation are then used to characterise severity-associated factors, local parameter sensitivity, structure-dependent network paths, and alternative response scenarios. The main work and findings are summarised as follows.

- (1) Based on a large body of real accident data, a set of UAV accident risk factors was screened, and a BN with strong risk-identification capability was constructed. Fourteen risk factor nodes were extracted from the accident narratives and historical meteorological data, and a hierarchical BN comprising risk factor node, intermediate, and target nodes was built. Introducing the Leaky Noisy-OR mechanism effectively reduced the parameter size of the conditional probability tables under multi-parent conditions and improved the feasibility of network parameter estimation under limited-sample conditions. Under stratified five-fold cross-validation, the model achieved mean validation values of 0.871 for accuracy, 0.960 for recall, 0.922 for F1 score, and 0.784 for AUC, supporting its ability to distinguish cases with severe and less severe consequences within the recorded accident dataset and to support the subsequent causation analysis and decision simulation.
- (2) Backward inference and sensitivity analysis showed that the environmental and human dimensions have the strongest model-based associations with severe consequences among the recorded accidents. Airspace exposure, wind, and operation error have the highest posterior probabilities under severe-consequence conditioning, while wind, bad weather history, and operation error have the highest local sensitivity coefficients. Under the assumed DAG, the path bad weather history $\rightarrow$ weather $\rightarrow$ environment $\rightarrow$ risk state has the highest path-level average edge-influence score (0.853), followed by wind $\rightarrow$ weather $\rightarrow$ environment $\rightarrow$ risk state and operation error $\rightarrow$ operation execution $\rightarrow$ human $\rightarrow$ risk state. These rankings describe structure-dependent posterior and local influence patterns.
- (3) The decision-simulation results show that the alternatives differ in both changes in post-decision high-risk probability and expected utility, and that their ranking depends on the trade-off among safety, mission continuity, time, and cost. Under the baseline safety-priority weights and the specified intervention-scenario assumptions, reroute achieves the highest expected utility of 31.967 and lowers the post-decision high-risk probability to 45%, providing the most favourable balance among safety improvement, mission preservation, and operational burden. However, adjust ranks first under the equal-weight, cost-priority, time-priority, and mission-priority schemes, and some local changes in the attribute scores and risk-impact factors also reverse the ranking of reroute and adjust. By contrast, reroute remains first-ranked under all 20 local ± 0.10 perturbations of the post-decision CPT probabilities. Reroute should therefore be regarded as the preferred strategy under the baseline safety-priority scenario rather than as a universally optimal intervention. The decision-simulation results represent scenario comparisons conditional on the specified utility preferences and intervention-effect assumptions.

A notable strength of this study is that risk-factor variables were reconstructed from unstructured ASN accident narratives and matched open-access meteorological data, enabling

heterogeneous accident records to be analysed within a unified probabilistic modelling framework. The dataset spans multiple countries, time periods, UAV types, and operational scenarios, while stratified five-fold cross-validation provides a more robust estimate of model performance within the current dataset. These strengths should nevertheless be interpreted in light of several limitations. First, ASN is a public accident database rather than a complete census of UAV operations. Military or larger platforms, serious accidents, and well-documented cases may be disproportionately represented, whereas minor or poorly reported incidents may be underrepresented. Consequently, the findings are most directly applicable to accidents recorded in ASN and cannot represent the full population of UAV operations. Second, although the integration of accident narratives and meteorological data facilitates consistent probabilistic modelling, the derived variables remain dependent on the completeness and consistency of the source reports, and their binary representation may not capture differences in risk intensity. Finally, stratified five-fold cross-validation assesses internal model performance but does not establish external transferability, while the decision results remain conditional on the specified utility parameters rather than on values calibrated from observed operational outcomes.

These limitations point to corresponding directions for future research. First, the external validity of the framework should be examined by integrating national investigation databases, regulatory reports, operator incident records, and insurance claims, and by testing the model on independent samples from different countries, periods, and UAV types. Second, access to more detailed operational data would allow the current binary nodes to be extended to ordinal or continuous representations and would support evaluation of the framework across a broader range of low-altitude aerial vehicles. Third, the utility parameters in the decision layer should be empirically calibrated using observed operating costs, mission benefits, and safety-loss data to improve the practical validity of utility estimates and strategy comparisons. Together, these extensions would strengthen the generalisability, analytical resolution, and operational relevance of the framework while preserving the advantages of its open-data-driven probabilistic modelling approach.

In summary, the proposed BN framework for UAV accident-severity analysis and decision simulation provides a systematic analytical foundation for model-based risk identification. By integrating unstructured accident narratives, meteorological data, and Bayesian inference, the framework identifies the factors most strongly associated with severe consequences and the highest-influence paths under the assumed network structure, and it compares alternative response strategies under explicit scenario assumptions. The findings can inform UAV operation management, mission planning, and safety oversight.

Author Contributions: Conceptualisation, R.H. and A.N.; methodology, R.H., J.Y. and Y.W.; software, J.Y., X.W. and Y.W.; validation, X.W. and X.Z.; formal analysis, R.H.; investigation, L.G., Y.Z. and X.Z.; data curation, L.G. and Y.Z.; writing—original draft preparation, R.H.; writing—review and editing, R.H. and A.N.; visualisation, R.H.; supervision, A.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by the 2025 Shanghai Jiao Tong University Research Program on Low-Altitude Economy and Industrial Technology, grant number 2025DKYA05.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The complete labelled dataset used in this study is publicly available at <https://github.com/hdd-2003/uav-accident-bayesian-network-data> (accessed on 4 August 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

UAV	Unmanned Aerial Vehicle
BN	Bayesian Network
BDN	Bayesian Decision Network
FTA	Fault Tree Analysis
DAG	Directed Acyclic Graph
CPT	Conditional Probability Table
EM	Expectation–Maximisation
ASN	Aviation Safety Network
FSF	Flight Safety Foundation
FAA	Federal Aviation Administration
CAAC	Civil Aviation Administration of China
AUC	Area Under Curve
TPR	True Positive Rate
FPR	False Positive Rate
OAT	One-at-a-time
GNSS	Global Navigation Satellite System
ESC	Electronic Speed Controller
RQ	Research Question

References

1. Asghari, O.; Ivaki, N.; Madeira, H. UAV Operations Safety Assessment: A Systematic Literature Review. *ACM Comput. Surv.* **2025**, *57*, 266. [\[CrossRef\]](#)
2. CAAC. Releases Statistical Bulletin of Civil Aviation Industry Development in 2025. Available online: https://www.caac.gov.cn/English/News/202604/t20260422_230642.html (accessed on 6 July 2026).
3. Aircraft Registration | Federal Aviation Administration. Available online: https://www.faa.gov/licenses_certificates/aircraft_certification/aircraft_registry (accessed on 6 July 2026).
4. Drone Sightings Near Airports | Federal Aviation Administration. Available online: https://www.faa.gov/uas/resources/public_records/uas_sightings_report (accessed on 6 July 2026).
5. Grindley, B.; Phillips, K.; Parnell, K.J.; Cherrett, T.; Scanlan, J.; Plant, K.L. Over a Decade of UAV Incidents: A Human Factors Analysis of Causal Factors. *Appl. Ergon.* **2024**, *121*, 104355. [\[CrossRef\]](#) [\[PubMed\]](#)
6. Gao, G.; Zhu, E.; Chen, A.; Zhang, W.; Wang, X. Analysis of Accident Factors of UAV Logistics Distribution Using Fault Tree Analysis and Analytic Hierarchy Process. In *Proceedings of the Second International Conference on Rail Transit and Transportation (ICRTT 2025)*; Li, F., Sun, Y., Eds.; SPIE: Nanjing, China, 2026; p. 6.
7. Abdallah, R.; Kouta, R.; Sarraf, C.; Gaber, J.; Wack, M. Fault Tree Analysis for the Communication of a Fleet Formation Flight of UAVs. In *Proceedings of the 2017 2nd International Conference on System Reliability and Safety (ICSRS)*; IEEE: Milan, Italy, 2017; pp. 202–206.
8. Khakzad, N.; Khan, F.; Amyotte, P. Safety Analysis in Process Facilities: Comparison of Fault Tree and Bayesian Network Approaches. *Reliab. Eng. Syst. Saf.* **2011**, *96*, 925–932. [\[CrossRef\]](#)
9. Kabir, S.; Papadopoulos, Y. Applications of Bayesian Networks and Petri Nets in Safety, Reliability, and Risk Assessments: A Review. *Saf. Sci.* **2019**, *115*, 154–175. [\[CrossRef\]](#)
10. Yang, C.; Mott, J.H. HFACS Analysis of U.S. General Aviation Accidents Using Bayesian Network. *Proc. Hum. Factors Ergon. Soc. Annu. Meet.* **2020**, *64*, 1655–1659. [\[CrossRef\]](#)
11. Wang, L.; Zhu, M.; Li, N. A Bayesian Network Approach for Systemic Risk Analysis in Unmanned Aerial Vehicle (UAV) Operations. *Sci. Rep.* **2026**, *16*, 13546. [\[CrossRef\]](#) [\[PubMed\]](#)
12. Sun, X.; Hu, Y.; Qin, Y.; Zhang, Y. Risk Assessment of Unmanned Aerial Vehicle Accidents Based on Data-Driven Bayesian Networks. *Reliab. Eng. Syst. Saf.* **2024**, *248*, 110185. [\[CrossRef\]](#)
13. Ye, B.; Zhu, S.; Wang, H. Identification of Key Risk Factors for Unmanned Aerial Vehicle Accidents Based on Common Cause Failure Bayesian Network. *Reliab. Eng. Syst. Saf.* **2026**, *272*, 112591. [\[CrossRef\]](#)
14. Das, B. Generating Conditional Probabilities for Bayesian Networks: Easing the Knowledge Acquisition Problem. *arXiv* **2004**. [\[CrossRef\]](#)
15. Anand, V.; Downs, S.M. Probabilistic Asthma Case Finding: A Noisy OR Reformulation. *AMIA Annu. Symp. Proc.* **2008**, *2008*, 6–10. [\[PubMed\]](#)
16. Xue, J.; Yang, P.; Wang, Z.; Li, Q.; Hu, H. Cause Analysis and Management Strategies for Ship Accidents: A Bayesian Decision Support Model. *Ocean Eng.* **2025**, *320*, 120291. [\[CrossRef\]](#)

17. Herashchenko, M.; Bashinsky, K.; Kamak, Y.; Nesterenko, S.; Rudnichenko, S.; Isachenko, O. Decision Support with Bayesian Influence Network During UAV Flight Control. In *Mathematical Modeling and Simulation of Systems*; Shkarlet, S., Morozov, A., Palagin, A., Vinnikov, D., Stoianov, N., Zhelezniak, M., Kazymyr, V., Eds.; Lecture Notes in Networks and Systems; Springer International Publishing: Cham, Switzerland, 2022; Volume 344, pp. 451–463.
18. Lin, J.; Luo, Y.; Li, C.; Zheng, Y. Effectiveness Evaluation of UAV Risk Mitigation Measures Based on Bayesian Networks. In *Proceedings of the 2nd Aerospace Frontiers Conference (AFC 2025)*; Springer Nature: Singapore, 2026; pp. 335–344. [\[CrossRef\]](#)
19. Ranter, H. Aviation Safety Netwok: An Exclusive Serice Flight Safety Foundation. Available online: <https://aviation-safety.net/database/> (accessed on 6 May 2026).
20. He, L.; Ma, X.; Qiao, W.; Yang, L. A Methodology to Assess the Causation Relationship of Seafarers' Unsafe Acts for Ship Grounding Accidents Based on Bayesian SEM. *Ocean Coast. Manag.* **2022**, *225*, 106189. [\[CrossRef\]](#)
21. Pearl, J. *Causality: Models, Reasoning, and Inference*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2009.
22. Fenton, N.; Neil, M. *Risk Assessment and Decision Analysis with Bayesian Networks*, 2nd ed.; Chapman and Hall/CRC: Boca Raton, FL, USA, 2019.
23. Mihaljević, B.; Bielza, C.; Larrañaga, P. Bayesian Networks for Interpretable Machine Learning and Optimization. *Neurocomputing* **2021**, *456*, 648–665. [\[CrossRef\]](#)
24. Ghasri, M.; Maghrebi, M. Factors Affecting Unmanned Aerial Vehicles' Safety: A Post-Occurrence Exploratory Data Analysis of Drones' Accidents and Incidents in Australia. *Saf. Sci.* **2021**, *139*, 105273. [\[CrossRef\]](#)
25. Xiao, Q.; Li, Y.; Luo, F.; Liu, H. Analysis and Assessment of Risks to Public Safety from Unmanned Aerial Vehicles Using Fault Tree Analysis and Bayesian Network. *Technol. Soc.* **2023**, *73*, 102229. [\[CrossRef\]](#)
26. Skibniewski, M. Bayesian-Network-Based Safety Risk Analysis in Construction Projects. *Reliab. Eng. Syst. Saf.* **2014**, *131*, 29–39. [\[CrossRef\]](#)
27. Wang, J.; Fan, H.; Chang, Z.; Lyu, J. Unleashing Data Power: Driving Maritime Risk Analysis with Bayesian Networks. *Reliab. Eng. Syst. Saf.* **2025**, *264*, 111310. [\[CrossRef\]](#)
28. Marcot, B.G.; Penman, T.D. Advances in Bayesian Network Modelling: Integration of Modelling Technologies. *Environ. Model. Softw.* **2019**, *111*, 386–393. [\[CrossRef\]](#)
29. Oniško, A.; Druzdzel, M.J.; Wasyluk, H. Learning Bayesian Network Parameters from Small Data Sets: Application of Noisy-OR Gates. *Int. J. Approx. Reason.* **2001**, *27*, 165–182. [\[CrossRef\]](#)
30. Ma, L.; Ma, X.; Chen, L. A Data-Driven Bayesian Network Model for Pattern Recognition of Maritime Accidents: A Case Study of Liaoning Sea Area. *Process Saf. Environ. Prot.* **2024**, *189*, 115–133. [\[CrossRef\]](#)
31. BayesFusion, LLC. Genie Modeler. User Manual. Available online: <https://support.bayesfusion.com/docs/> (accessed on 21 October 2019).
32. Druzdzel, M.J. SMILE: Structural Modeling, Inference, and Learning Engine and GeNIe: A Development Environment for Graphical Decision-Theoretic Models. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, Orlando, FL, USA, 18–22 July 1999; pp. 902–903.
33. Cai, B.; Liu, Y.; Fan, Q.; Zhang, Y.; Liu, Z.; Yu, S.; Ji, R. Multi-Source Information Fusion Based Fault Diagnosis of Ground-Source Heat Pump Using Bayesian Network. *Appl. Energy* **2014**, *114*, 1–9. [\[CrossRef\]](#)
34. Trucco, P.; Cagno, E.; Ruggeri, F.; Grande, O. A Bayesian Belief Network Modelling of Organisational Factors in Risk Analysis: A Case Study in Maritime Transportation. *Reliab. Eng. Syst. Saf.* **2008**, *93*, 845–856. [\[CrossRef\]](#)
35. Castillo, E.; Gutiérrez, J.M.; Hadi, A.S. Sensitivity Analysis in Discrete Bayesian Networks. *IEEE Trans. Syst. Man Cybern.-Part A Syst. Hum.* **1997**, *27*, 412–423. [\[CrossRef\]](#)
36. Fan, S.; Blanco-Davis, E.; Yang, Z.; Zhang, J.; Yan, X. Incorporation of Human Factors into Maritime Accident Analysis Using a Data-Driven Bayesian Network. *Reliab. Eng. Syst. Saf.* **2020**, *203*, 107070. [\[CrossRef\]](#)
37. Howard, R.A.; Matheson, J.E. Influence Diagrams. *Decis. Anal.* **2005**, *2*, 127–143. [\[CrossRef\]](#)
38. Quinones-Grueiro, M.; Biswas, G.; Ahmed, I.; Darrah, T.; Kulkarni, C. Online Decision Making and Path Planning Framework for Safe Operation of Unmanned Aerial Vehicles in Urban Scenarios. *Int. J. Progn. Health Manag.* **2021**, *12*. [\[CrossRef\]](#)
39. Neogi, N.; Bhattacharyya, S.; Griessler, D.; Kiran, H.; Carvalho, M. Assuring Intelligent Systems: Contingency Management for UAS. *IEEE Trans. Intell. Transp. Syst.* **2021**, *22*, 6028–6038. [\[CrossRef\]](#)
40. Pidd, M. Tools for Thinking—Modelling in Management Science. *J. Oper. Res. Soc.* **1997**, *48*, 1150. [\[CrossRef\]](#)
41. Hanea, A.M.; Christophersen, A.; Alday, S. Bayesian Networks for Risk Analysis and Decision Support. *Risk Anal.* **2022**, *42*, 1149–1154. [\[CrossRef\]](#) [\[PubMed\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.