Journal of Environmental Management 384 (2025) 125373

Contents lists available at ScienceDirect

Journal of Environmental Management

journal homepage: www.elsevier.com/locate/jenvman

Research article

Participatory Bayesian Networks for uncovering reflexive unknowns in strategic environmental risk management

Annukka Lehikoinen[a,b,*], Tapio Reinekoski[c], Nina Janasik[b], Marko Ahvenainen[d,e], Janne I. Hukkinen[f]

[a] Kotka Maritime Research Centre, Salakuljettajantie 4, 48100, Kotka, Finland
[b] University of Helsinki, Viikinkaari 1 (PB 65), 00014, Helsinki, Finland
[c] Tampere University, Faculty of Social Sciences (SOC), 33014, Tampere, Finland
[d] National Emergency Supply Agency, Aleksanterinkatu 48 A, 00100, Helsinki, Finland
[e] Finland Futures Research Centre, University of Turku, Turku School of Economics, 20014, Turku, Finland
[f] Environmental Policy Research Group, Helsinki Institute of Sustainability Science, Faculty of Social Sciences, Unioninkatu 37 (PB 54), 00014, Helsinki, Finland

ARTICLE INFO

Keywords:
Participatory modelling
Bayesian Networks
Socio-environmental disruptions
Leverage points
Strategic risks
Systems analysis
Future pathways

ABSTRACT

Strategic environmental risk management and planning must account for uncertainty and complexity, necessitating methods that facilitate scenario development under incomplete knowledge. This paper introduces a participatory modelling (PM) -based knowledge co-production and strategic planning approach utilizing one type of AI tool - Bayesian Networks (BN) - for systemic scenario development, analysis and resilience-building. The developed method integrates diverse perspectives and expertise of participants through a structured BN model, enabling co-imagination and -construction of causal pathways, translating them into probabilistic dependencies, and diagnostically identifying potential leverage points for strategic resilience-increasing actions. We illustrate and test this approach using a case study of a chemical transportation accident in an urban environment, documenting the participatory process and the algorithm to translate the participants' thinking to a computational BN. Through content analysis of transcribed audio recordings, we demonstrate how the exercise helped uncover "reflexive unknowns" – previously unrecognized threats that became apparent and thinkable only through the collaborative modelling process. An example of such a reflexive unknown in our case exercise is the prospect of toxic rainfall following the accident and its short- and long-term implications for the built and natural environment. This was a blind spot in the thinking of the participants, and it appeared and became a scenario to be acted upon only as a result of the process of collective cross-sectoral causal thought represented with a BN model. The paper provides a detailed description of the developed participatory BN approach and methodology, enabling their applicability in various contexts. Through a qualitative analysis of the exercise's implementation, the article also demonstrates how the approach fostered collective, iterative reflection, generating new insights to socio-environmental resilience.

1. Introduction

Strategic decisions and management measures for environmental regulation and security are implemented in the face of an uncertain and complex future, requiring consideration of combinations of events that have not yet materialized (Ludwig et al., 1993; Schindler and Hilborn, 2015). This creates a need for approaches that enable imagination and elaboration of scenarios in the light of prevailing knowledge and uncertainties, allowing identification of policies that are robust to uncertainty, thus likely leading to resilient management strategies. We present a participatory Bayesian network (BN) approach that enhances long-term strategic resilience by uncovering *reflexive unknowns* in the collectively crafted risk scenarios, i.e., unknowns that are systemic in nature and can be articulated only with the help of a BN model and the collective work put into it. In scenario planning, instead of accurate predictions about the future, the core idea is to creatively think about a variety of possible futures (Peterson et al., 2003; Hichert et al., 2021) and how these different, contingent futures compel different kinds of

* Corresponding author. Kotka Maritime Research Centre, Salakuljettajantie 4, 48100, Kotka, Finland. University of Helsinki, Viikinkaari 1 (PB 65), 00014, Helsinki, Finland.

E-mail addresses: annukka.lehikoinen@merikotka.fi, annukka.lehikoinen@helsinki.fi (A. Lehikoinen).

https://doi.org/10.1016/j.jenvman.2025.125373
Received 30 September 2024; Received in revised form 6 February 2025; Accepted 12 April 2025
Available online 5 May 2025

management actions (cf. Esposito, 2024; Reinekoski, 2025). Incorporating diverse perspectives into the scenario planning process is recommendable, as cross-sectoral and -disciplinary scenarios may help to identify social-ecological-technological interlinkages, enabling the exploration of the dynamics and consideration of holistically and collectively good (i.e. sustainable) pathways and management measures (Düspohl et al., 2012 and the references therein). Participatory scenario development provides a way to integrate the knowledge and realities of practitioners, citizens, and scientists, filling information gaps, increasing the relevance of the scenarios, as well as the acceptance and legitimacy of scenario findings and recommendations (Hichert et al., 2021 and the references therein).

As Bayesian Networks (BN) – including their decision analytical extension, Influence Diagrams – have characteristics particularly suitable for analyzing management problems where decision-making must account for complex socio-environmental interactions, systemic feedback, and uncertainties arising from diverse origins, the approach has been applied in numerous environmental management studies over time (see e.g. the overviews by Barton et al., 2012; Sperotto et al., 2017; Parviainen et al., 2021; Kaikkonen et al., 2021). By structuring knowledge and stakeholder perspectives into probabilistic models, the approach supports the exploration of alternative futures and enhances the capacity to identify management strategies that remain effective across a range of plausible scenarios. This makes BN modelling a valuable tool also for scenario planning, facilitating the integration of diverse expertise and fostering resilient, adaptive environmental governance.

Technically BNs are classified as one type of Artificial Intelligence (AI) tools, representing structural and conditional aspects of a specific phenomenon or problem domain, and allowing systemic inquiries in the presence of uncertainty (Korb and Nicholson, 2010). A BN consists of two key elements: a graphical component allows for organizing and illustrating causal information and conceptual frameworks (Carriger et al., 2018; Carriger and Parker, 2021; Luoma et al., 2021, 2024), while the numerical, probabilistic aspect – as combined with the causal structure – permits the predictive and diagnostic examination of how the system operates (e.g. Kuikka et al., 1999; Uusitalo et al., 2011; Lehikoinen et al., 2013, 2014, 2015; Helle et al., 2015; Penman et al., 2020).

BN applications have become increasingly popular in recent decades, particularly in the realm of various forward-looking analyses like environmental risk assessments (Kaikkonen et al., 2021). As these analyses typically deal with yet unmaterialized scenarios, it is often justifiable to rely on logical reasoning and expert judgments to consider the possibility and plausibility of diverse events and their consequences (O'Hagan et al., 2006; Düspohl et al., 2012; Pihlajamäki et al., 2020; Janasik, 2021). BNs have been applied as participatory modelling (PM) tools in numerous environmental studies covering different levels and aspects of participation (e.g. Hukkinen, 1993; Bromley et al., 2005; Castelletti and Soncini-Sessa, 2007; Carmona et al., 2011; Mäntyniemi et al., 2013; Andriatsitohaina et al., 2020; Singto et al., 2020; Mayfield et al., 2023; see also the overview by Düspohl et al., 2012). As visually accessible models, they have been suggested to function as a type of boundary objects (Star and Griesemer, 1989) for diverse groups to share their knowledge and perspectives (Moglia et al., 2018; Müller et al., 2024; but see also Parviainen et al., 2022).

In this article, we present a participatory BN-based modelling approach developed in conjunction with a Finnish case study but specifically designed for broad applicability in enhancing the strategic long-term resilience of municipalities against complex socio-environmental disruptions. A recent analysis of 58 expert interviews in Finnish cities and other public organizations clearly shows a need for such approaches and solutions (Salomaa et al., 2025). A disconnect was found between the explorative approach required by strategic resilience management and the institutionally established practices of crisis planning and operations. This disjointedness permits socio-environmental crises to creep in, because the tried-and-true methods of crisis planning and operations constrain the ability to account for plausible but poorly understood long-term threats.

Aiming to tackle this issue, acknowledging challenges recognized with another tested approach (Järvensivu et al., 2021 – more in section 2.1), we designed and tested a participatory knowledge co-production and strategic joint-planning exercise that utilises BN as the platform and tool for systemic scenario and strategy development. Our approach consists of the processes of and associated techniques for (1) representing the participants' joint view of the scenario system as a causal structure of the BN, (2) translating the participants' thinking about plausible causal dependencies to a probabilistic format, to populate the conditional probability tables of the BN, and (3) the computational use of the BN to help the participants in identifying the key *leverage points* (Fischer and Riechers, 2019) of their scenario system to start discussing about rational management actions for improving resilience.

Here we describe the developed approach and analyse its functionality based on data and observations from a use case on an urban chemical transportation accident and its impacts on the long-term social, environmental, and economic objectives of a city. We end up arguing that a *prognostically* (causally: from causes to effects) co-produced Bayesian network and its *diagnostic* (inference from effects to causes) use can make actionable something we call "**reflexive unknowns**". Ekengren (2024) has classified surprises into "known unknowns" (such as strategic surprises and expected surprises), "unknown unknowns" (unexpected surprises) and "known-corporeally unknowns" (surprises that you became aware of only after direct corporeal contact). To this list we add *a fourth type*: the **reflexively known unknowns**. The concept means observations that could be articulated only thanks to the emergent systemic mechanisms that would have remained unknown without the BN model and the collective work put into constructing it, pinpointing shortcomings in the scenario system or/and its management, such as inadequate expertise or infrastructures for dealing with the implications of plausible but poorly understood environmental threats.

While our test case focuses on an urban chemical accident scenario involving a city organization, the approach is designed to be applicable to any group and adaptable to a wide range of scenarios that pose risks to environmental and social objectives relevant to the participants. The contribution of this article is twofold: (1) we provide a detailed description of the exercise methodology, ensuring its applicability – if not direct replicability – in various contexts, whether concerning specific socio-environmental hazards and risks or the strategic planning processes of different organizations; and (2) through a qualitative analysis of the exercise's implementation, we demonstrate how the modelling procedure fosters collective, iterative reflection, generating new insights and observations.

Notably, the key purpose of our participatory BN-modelling approach is not to provide accurate predictions on the actual state of the target variables, but to *understand the systemic mechanisms through which they might be at risk*: the key factors, their direct and indirect dependencies and non-dependencies, the mutual plausibility of different combinations, and finally the parts of the system whose manipulation or control would be most useful – from a systemic viewpoint – for effective risk management strategies under uncertainty. Our approach does not question the validity of the participants' mental model but accepts it as given, encouraging groups to co-imagine, then focusing on structuring and analyzing their thinking, and finally using the outcomes to support the co-diagnosis of where in the system actions would best improve the resilience – meaning the city organization's capacity to protect their values (i.e. strategy objectives) based on the proposed mental model. Thus, in parallel with the more traditional, prognostic, first order modelling, our diagnostic "thinking about thinking" -type of modelling exercise contributes to second order research, as defined and recalled by Fazey et al. (2018). This is also somewhat corresponding to what Voinov et al. (2018) and Zellner and Campbell (2015) call single vs. double loop learning.

To our knowledge, this is the first published participatory BN application that focuses on the diagnostic (from unwanted consequences to their potential causes) analysis of participants' thought scenarios, in order to identify the key systemic leverage points for management. We are also not aware of any other studies that have examined and documented, through qualitative content analysis of a participatory BN modelling exercise, *how* the process actually generates new insights and knowledge and facilitates collaborative reflection and learning among the participants.

The present approach and case study have earlier been described in Finnish by Ahvenainen et al. (2021). This article is the first presentation of the concept in English, and the first presentation of the translation algorithm from experts' thinking about causalities to a probabilistic BN. In addition, this article contains a totally new analysis of the functionality and generalizability of the approach. The data and analysis of test users' experiences and reactions as well as the following conclusions are entirely new.

The structure of the article is as follows: In Section 2, we provide the background of the study, introduce the research questions, and describe the phases of the exercise in a way that makes them replicable in other contexts. We also describe the types of data associated with our test exercise. The theory and principles of BNs, as well as the algorithm we developed for translating the participants' descriptions of causalities into conditional probability distributions, are presented in Appendix 1. In Section 3, we first present the types of management actions that, based on the outcomes of the test exercise, can be identified using the presented method. We then elaborate on the key findings of the qualitative content analysis conducted on the test exercise, focusing on how the tested PM methodology facilitated collective learning and systemic understanding, helping to identify factors and actions significant for the city's resilience against environmental crises similar to the exercise scenario. In Section 4, we discuss these findings from diverse perspectives.

2. Material and methods

2.1. Background and objectives

The original aim behind the study was to design so called Policy Operations Room exercises (POR) to strengthen long-term resilience of Finnish cities against diverse socio-environmental disruptions (Hukkinen et al., 2022). A key idea behind the POR concept is to demonstrate the potential long-term consequences of diverse acute crisis management choices made under time pressure and remarkable uncertainties. The hypothesis behind PORs is that acknowledgement of such future pathway mechanisms leads to more resilient policies and decision-making. In an earlier POR exercise (Järvensivu et al., 2021), a set-up where participants were presented with fixed narrative situational pictures and forced to make fast choices between given pre-defined strategic decision options with long-term disruptive consequences, some focal design issues were identified, which we now wanted to tackle. These were: a) the need for more engaging, socio-politically more relevant and complex, scenarios; b) the too narrow and unrealistic strategy (decision) options among which to choose; c) the need for more comprehensive involvement of the participants in the planning and development of the exercise – which could have also solved the issues a) and b). In addition, in the preceding POR, participants couldn't turn back and reconsider their choices, seeing where the other pathways would have taken them. While choosing they fully closed the possibility to consider the other, alternative future pathways. Although this may be the case in real life, we questioned whether such exercise optimally develops the participants' capacity to make resilience-increasing strategic choices under pressure.

Against this background we decided to experiment with PM and try designing an exercise for co-producing a model that would allow the participating actors themselves to identify and elaborate the relevant future pathways and recognize resilience-creating factors and actions, amidst uncertainty that typically arises from a multitude of (manageable and non-manageable) sources. After preliminary discussions with a representative of our case city, the goal was refined to help municipal sectoral managers identify factors and actions critical to long-term strategic crisis resilience of the city by modelling their co-imagination about long-term implications of a locally relevant acute disruption scenario: a chemical transportation accident. We wanted to make the managers consider path-dependencies between the acute scenario and its potential long-term consequences for the city organization, not forgetting the associated uncertainties affecting plausibility of diverse scenarios.

Given our goals, we needed a modelling approach allowing systemic representation and dynamic analysis of (presumably complex and strategically wicked) scenario pathways, and the consideration of uncertainties. The method also had to be accessible to participants without modelling expertise. Therefore, we selected BN as the modelling tool to be tested. We hypothesized that if we succeeded in representing as a computational BN model the participants' cross-sectoral collective thinking about the pathways connecting the acute disruption scenario and its relevant long-term consequences, we might be able to provide them with a new order of observation and help them identify such leverage points and associated interventions that are critical to the city's strategic resilience against corresponding incidents.

Thus, our research questions were (1) what kind of participatory BN-modelling process could support our goals, and (2) to what extent such an approach could be incorporated into strategic risk management more broadly. These questions are elaborated through the data and experiences from a true use case example. For security reasons, not all details of the case study can be presented publicly; instead, the description is kept at a level that allows the examination of the above-mentioned research questions.

2.2. The participatory modelling process

Amidst the COVID-19 pandemic, in January–March 2021, the exercise was conducted online (Teams video calls) over three sessions with experts from various sectoral branches of a Finnish municipality. In the selection of participants, we sought to include key actors within the city organization who could provide diverse perspectives and had some degree of power (Voinov and Bousquet, 2010) in the planning and decision-making processes of the city. To foster active participation and meaningful discussions, also considering the online format, we prioritized the "depth" (over "breadth") of participation (Voinov et al., 2018), aiming to keep the number of participants relatively small. The six participants were department heads and planners from three different sectoral departments of the city. The selection process was also influenced by practical factors, such as schedule compatibility and the participants' ability to attend all three sessions.

Initially, participants were placed in an imaginary scenario in the near future, where the long-term consequences could jeopardize the strategic objectives of the city. The starting acute crisis scenario was a hypothetical but plausible event discussed in background interviews and regional risk assessments: a major industrial chemical transportation accident posing risks to the environment and public health, set in a real urban environment.

Bayesian network (BN) was applied as the central method of systems analysis and cognitive PM. Key constraints that influenced the design of the modelling exercise included the limited time resources and engagement of the participating experts: the exercise had to be meaningful and feasible without taking too much of their time. We did not want to restrict the richness (i.e. complexity) of their thinking, nor use their time and mental capacity for daunting elicitation of probabilities. Instead, we decided to follow the group's discussion and do our best to represent their thinking as a BN that could offer them additional insights into the system they had collectively described.

The modelling process can be divided into three phases, as follows:

2.2.1. *First stage: systemic prognosis (representing the participants' joint scenario space as a causal structure)*

In the beginning of the first meeting, we anchored the initial scenario in the actual physical environment through an audiovisual crisis narrative, a *crisis dashboard* (see Hukkinen et al., 2024 preprint), to enhance the participants' immersion in the fictional situation and motivate them to take the exercise seriously. A video simulating a newscast and a BN model representing the situational picture right after the hazard were presented to the participants (N6), who were managers and experts of different operative spheres in the city organization. In the test scenario, a truck carrying toxic chemicals derails from a bridge onto a chemical train passing under, causing the release of a chemical cocktail to the ground, the sea, and the air. The accident closes a central passageway to the city center that is located on an island. A major summer festival is underway in the city, and there may be tens of thousands of people in the area. The acute BN model of the possible immediate impact mechanisms of such an accident, developed by the authors (AL and MA), was based on the material found from the literature, and served as the starting point for the PM (Fig. 1).

Continuing from the acute situation model, the city's experts considered the impact chains in the longer-term future from the perspective of three variables aggregating acute harm effects ("long-term drivers" in Fig. 1): (1) level of harm to the ecosystem caused by toxic chemical leaks, (2) secondary human health issues through the polluted environment, and (3) human health issues due to acute exposure to a gasifying toxin. The task was, through discussion and drawing facilitated by the researchers, to formulate causal paths connecting the acute harm variables to the variables representing the indicator metrics of the city's long-term strategic goals.

Utilizing harm-specific four-field illustrations shown in Fig. 2, the participants were first directed by MA to discuss such cause-and-effect paths. The goals and their indicator metrics (see Fig. 2) were extracted from the city's strategy document for the year 2025.

Each of the three acute harm perspectives were treated one-by-one and their connections to the strategic goals structured by means of causal drawing. During the discussion, AL outlined a graphical causal representation of the mentioned influential factors and the causal connections between them. Immediately after each discussion, the drawn causal map was reviewed together with the participants and supplemented as necessary. Methodologically, this part of the exercise represents collective construction of cognitive maps (LaMere et al., 2020).

After the first meeting, common variables connecting the so far separate harm-specific networks were identified, combining them to one complex network of pathways connecting the acute harm variables to the long-term strategic goals (see Fig. A1 in Appendix 1). In the second meeting, the participants reviewed and approved the combined causal diagram which had been sent to them beforehand. After that, they were asked to refine their thinking about the causal relationships by discussing and describing the quality (positive vs. negative dependency) and strength (weak vs. moderate vs. strong effect) of each cause-and-effect link. In addition, the uncertainty associated with assessing the quality of each link was defined either as low-moderate or particularly high.

2.2.2. *Second stage: numerical translation (representing the participants' thinking of causalities probabilistically, to populate conditional probability tables of the BN)*

After the second meeting, AL transformed the mental model defined by the participants into a numerical Bayesian Network that enables a novel analysis of the causal scenario system described. The core idea in the developed probabilistic "translation" approach is to assume linear dependency and increase uncertainty around the expected values applying simple triangular distributions. The translation algorithm is described in Appendix 1.

While developing the translation algorithm, it was noticed that the

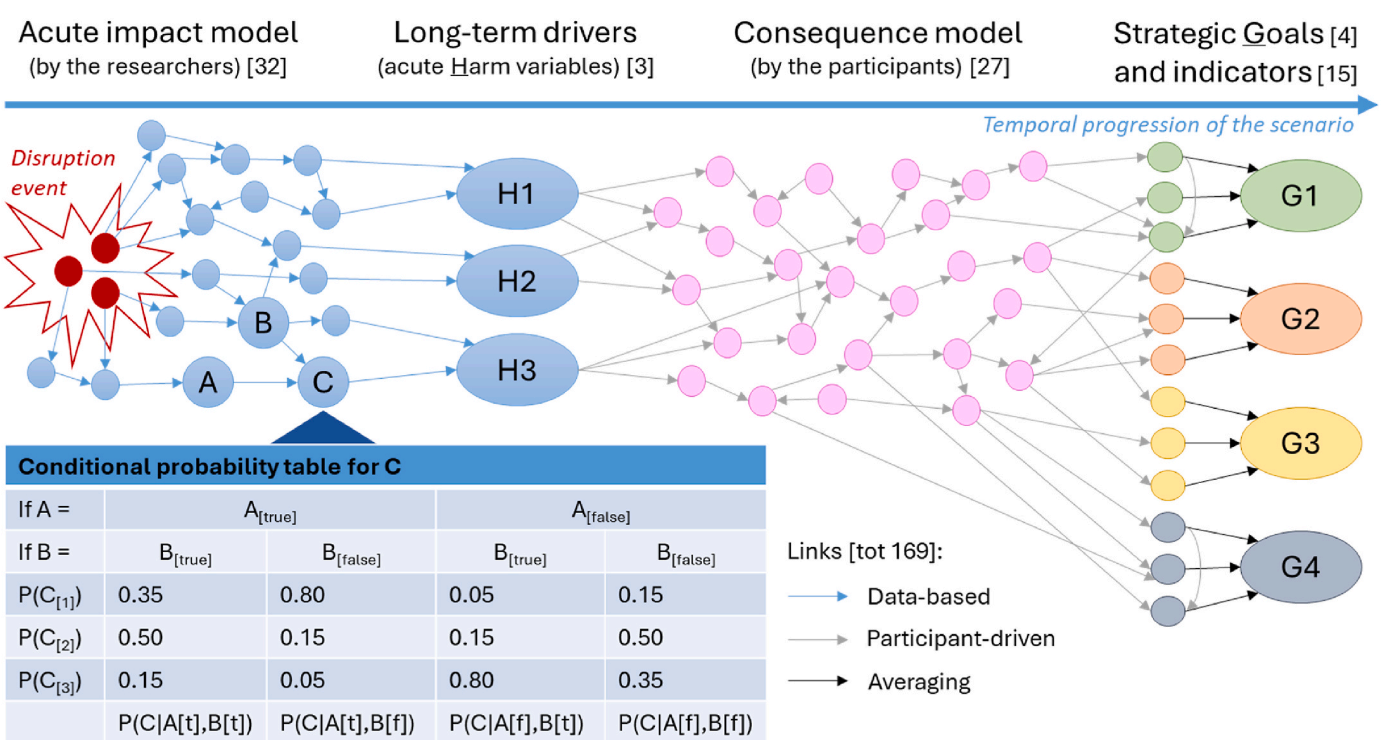

Conditional probability table for C

If A =	$A_{[true]}$		$A_{[false]}$	
If B =	$B_{[true]}$	$B_{[false]}$	$B_{[true]}$	$B_{[false]}$
$P(C_{[1]})$	0.35	0.80	0.05	0.15
$P(C_{[2]})$	0.50	0.15	0.15	0.50
$P(C_{[3]})$	0.15	0.05	0.80	0.35
	P(C|A[t],B[t])	P(C|A[t],B[f])	P(C|A[f],B[t])	P(C|A[f],B[f])

Fig. 1. Overview of the exercise. Based on data from the literature, the authors created an acute impact model from which the participants proceeded to build a long-term consequence model linking driver (acute harm) variables to city strategy target indicators. The brackets denote the numbers of nodes or links in the BN created as part of our case study, but the illustrated model structure is exemplary. The probabilistic representation of dependencies in a BN's conditional probability table is shown for a hypothetical three-state variable C, with parents A and B being two-state (true/false) variables. Each column represents a conditional probability distribution, where the probabilities of the states of C sum to one. The operating principle of Bayesian networks is explained in Appendix 1.

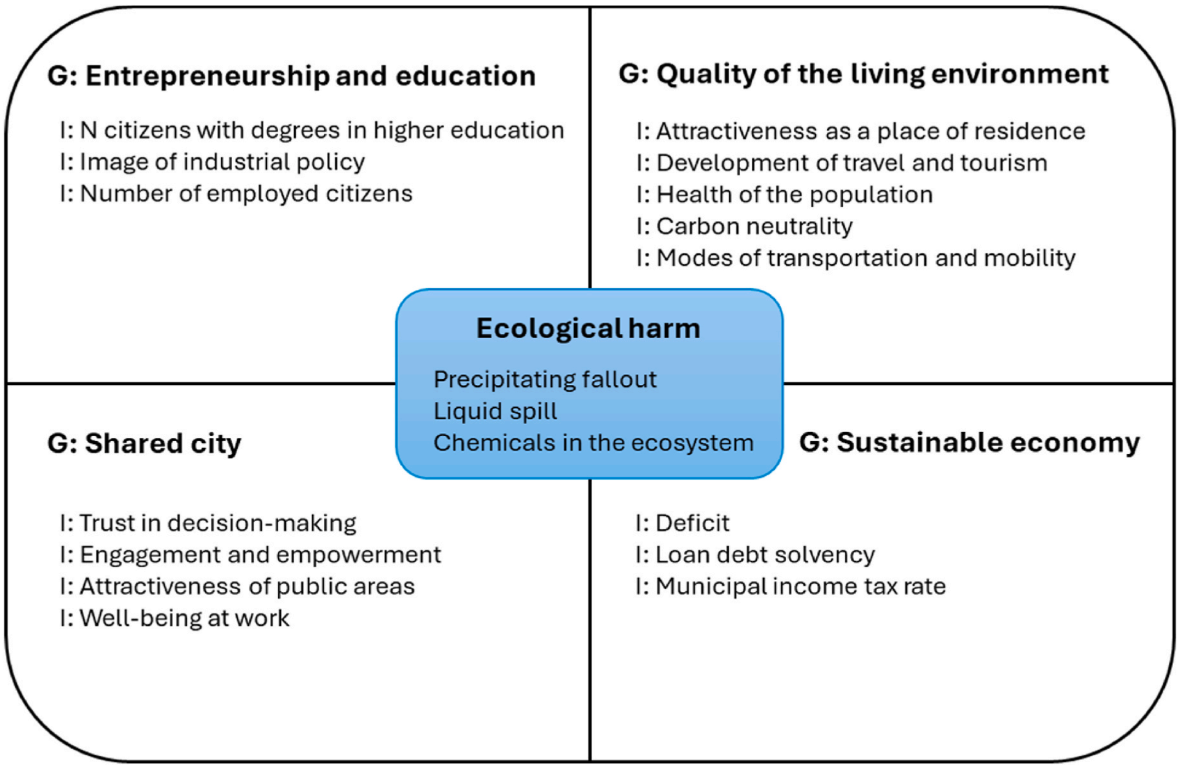

Fig. 2. An example four-field to support the discussion regarding the long-term consequences of the acute ecological harm to the four strategic objectives of the case study city. Each field represents one strategic goal (G) with its indicator metrics (I). The keywords under "Ecological harm" are the factors (parent variables in BN terms) that together make the "Ecological harm" in the researchers' acute model. Corresponding four-fields were applied for the two other harm perspectives.

variables of this case study system could be represented basically with three different node types: (1) the nodes describing change or development of something, (2) the nodes defining whether some event materializes or not, and (3) the nodes representing the level of something. In the resulting BN, Type 1 ("change") nodes can take an integer value between -2 and +2, where the zero state corresponds to the state before the chemical accident, with -1 and +1 indicating moderate and -2 and +2 indicating strong negative or positive changes. Type 2 ("occurrence") nodes, in turn, have the alternative states of true or false. In addition, as the starting point nodes (the output nodes from the researchers' model) we had nodes describing the level of harm. These Type 3 "harm nodes" can take some of the three alternative states representing "low", "moderate" or "high" level of the acute harm (Appendix 1). The final "assessment endpoints", i.e. the strategy objectives and their indicators, were represented as Type 1 nodes, depicting their development in the post-accident future, compared to the pre-accident status.

2.2.3. Third stage: diagnosis of the scenario system (analysing the BN to help the participants identify key leverage points and potential resilience measures)

With the completed numerical model, using the BN software Hugin (Madsen et al., 2005; www.hugin.com), parameter sensitivity analyses were run utilizing the "Sensitivity Set Graph Panel" tool. The analysis tool shows how sensitive a certain state of a selected "hypothesis variable" is to changes of the parameter values (CPT entries) of other variables in the model. By using the city's strategic objective variables (integrating the distributions of their indicator variables) one by one as the hypothesis variables for the analysis, focusing on their worst states (-1 and -2), we identified variables that could potentially play significant role in the failure of strategy implementation. Some of these variables belonged to the acute crisis model, such as weather conditions prevailing during the accident, but the majority were included in the participants' constructed long-term cause-and-effect chains threatening the strategy.

After this, the variables identified as the most relevant were examined more closely, applying both prognostic and diagnostic approaches, to delve deeper into the mechanisms through which these factors could jeopardize the implementation of the strategy. Based on these results, narratives were written by MA and TR, illustrating what the future after the initial situation might look like when each goal reaches its worst possible state.

In the third, final, meeting of the exercise, the participants were presented with the goal-specific worst-case narratives. In addition, they had access to a view of the entire model in Hugin, as well as the probability distributions of variables. Per each strategy goal and the associated narrative, the participants reflected on three questions: Did the narrative seem logical and justified considering their original thinking? What could or should the city, as an organization, do in terms of preparedness? What would hinder or, conversely, facilitate successful preparedness? At the end of each perspective discussion, the researchers also challenged the participants with additional questions based on observations made during the construction of the BN and its analytical use.

The thoughts arising from the discussions were compiled into an action table, detailing actions to reduce the likelihood of a disturbance like the initial one or its consequences and to enhance the city's strategic future resilience. In the end of the last meeting, we had also allocated time for a debrief session to receive immediate feedback from the participants.

In Appendix 1, section 1.3, we provide a closer look at the different roles of the researchers at various stages of the PM process.

2.3. Data and analysis of the participatory modelling exercise

The three meetings with the participants were recorded and transcribed. We listened to the recordings together and shared our immediate observations and thoughts on the exercise and how it aligned with the questions and objectives we had set. Afterwards, TR, NJ and MA conducted a more detailed qualitative content analysis of the transcripts, which AL complemented from the modelling-oriented perspective (section 3.2). The analysis was guided by our two research questions: how the modelling process allowed the participants to identify and elaborate relevant future pathways vis-à-vis the city's strategic goals; and how the exercise approach could be incorporated into strategic risk management more broadly. In addition to the recordings and transcripts, the design and implementation of the exercise produced drawings, BNs, BN analytics, and the final action table. Although presented here on structural (non-content) level only, these served as relevant materials for us as well, when we elaborated the research questions.

The full transcripts and other raw data from the case study cannot be shared to readers as the substance matter of the exercise deals with sensitive issues pertaining to the maintenance of national security and emergency preparedness. Second, our research consortium committed to full confidentiality and anonymity in our interactions with the participating individuals. Third, full confidentiality and anonymity could not be definitively secured even with anonymization because of the risk of revealing individual or organizational identities from the topics discussed and the contexts in which they were discussed. Fourth, the difficulty of keeping confidentiality and anonymity commitments was amplified by the fact that Finland is a small nation state with close individual and organizational networks. However, publishing the full case study data is not necessary considering the purpose and findings of this article, focusing on documentation and demonstration of a PM approach – not a particular model. The presented excerpts from the transcribed discussions during the exercise are original and serve to verify the interpretation of the authors.

3. Results

In this section, we describe our key observations regarding the immediate outcomes of the PM exercise, i.e. the co-identified management actions (3.1), and present the results of the qualitative analysis of the proceedings of the exercise (3.2), showing how the exercise engendered emergent, collective insights and observations of the possibilities and limitations of the city's strategic thinking and its capacities for managing environmental disasters such as the example scenario.

3.1. Resilience-increasing management action categories

The total number of potential resilience-increasing actions identified in the test exercise was 28. Based on the diagnosis of the co-developed BN model, it could be observed that the identified strategic leverage points were not limited only to the long-term consequence pathway model created by the participants themselves. Some variables from the acute accident model – including uncontrollable weather phenomena – also emerged. Accordingly, in addition to anticipatory pre-accident

Table 1
The format of the action table generated, listing and categorizing the different types of potential management measures to increase the strategic long-term resilience of the city against corresponding environmental hazards. The types of measures that were identified by the participants in our test exercise are indicated in blue font, whereas the ones in grey font are possible, too.

Strategic mission (Strategy objective)	Measures and processes for accident prevention	Measures and processes for mitigating consequences	Cross-cutting resilience-increasing measures and processes
Safeguarding the quality of the living environment	*E.g. measures focused on improving the safety of transport routes and enhancing monitoring and control of shipments*	*E.g. measures that improve readiness for controlling or managing environmental damage during and after an accident*	*E.g. Actions to develop operational communications;* *Proactive risk communication and stakeholder engagement actions;* *Training and education to improve preparedness and capabilities*
Safeguarding the conditions for economic activities	*E.g. measures that support the continuation of chemical transportation while reducing the likelihood of accidents.*	*E.g. measures that ensure the continuity of transportation and other business activities during and after an accident*	
Safeguarding the wellbeing of residents	*E.g. measures aiming to reduce the likelihood of chemical transport accidents while ensuring smooth and safe mobility for residents*	*E.g. measures that safeguard the well-being of residents and minimize health risks during and after an accident*	
Mitigating the costs incurred by the city	*E.g. measures aiming to shift the costs of proactive risk management from the city to the regional or national level, or to private companies*	*E.g. Insurances and other measures aiming to transfer part of the response and recovery costs away from the city in the event of an accident*	

actions, both acute-phase and post-incident measures were recognized as significant from a long-term strategic resilience perspective. Some of the measures aimed at preventing similar accidents were quite specifically tied to the accident site, but interestingly, a clear majority of the identified actions served the city's preparedness and resilience against similar types of environmental hazards more broadly, regardless of the exact location, timing, or chemicals involved. The identified measures were compiled into a table, categorized by the strategic city objectives targeted. Some of the proposed actions could have been assigned to more than one of these categories. We used the primary strategic objective perspective (as it emerged in participants' discussions) of each action as the placement criterion. In addition, several cross-cutting actions not tied to any specific objective were identified and assigned to a separate category (see Table 1). Table 1 shows the resulting action categories and descriptions of the types of actions that fit them. The precise actions were agreed to be kept confidential.

3.2. The content analysis: construction of collective insight and reflexivity in the exercise

3.2.1. From distinct causalities to collective system-wide insights

In addition to the list of potential resilience-increasing actions, the causal mapping and the diagnostic use of the finished model yielded new insights for the participants, emerging as the result of the collective work and the systemic view constructed of the scenario and its ramifications. While the construction of the model required participants to focus on the causal connections between two variables at a time and bracket the rest, once the graphic representation of the model structure was finished at the end of the first exercise session, the participants identified the possibility of a system-wide disaster, something that would cut across nearly all the strategy objectives conceived. Namely, toxic gas clouds might precipitate over the city, in the event of an unlucky timing of rainfall. We quote the conversation at length to illustrate the gradual accumulation of the shared perception.

'Participant 1: Is it possible that with that gaseous cloud, if there's a precipitation event in the city center, we might actually have quite a wide area (…)

P2: Are you referring to the exposure area?

P1: Yes.

P2: There would also be a pollution area and thus an exposure area coming from it. Yes. It is possible. True.

P1: Ouch.

(…)

P1: I'm left wondering, if such a gas cloud precipitates in the city center, how on earth do we clean this up? (…) It's on top of buildings, in all maintenance services, all the staff maintaining that environment, the impact here is (…) it endangers the entire maintenance system due to this environmental hazard. How do we operate there?

P3: This also has broader implications for all other residents and people in that area … the impact is much greater then.

P4: So, this could lead to long-term health issues: cancer effects and others. Those are indeed difficult to measure.

P1: Yes. The harmfulness of this precipitation is quite a serious issue. Because we're talking here about the whole culture of our operation and how it can recover from such a significant environmental hazard (…) Now, this started to hit me as I began to think about how much maintenance activity is related to this precipitation (…) and if it occurred over a larger area in the city center, then we would be talking … this issue would multiply.

P4: Playgrounds.

P1: All of them.

P4: Daycares, school yards.

P1: All of them.

P4: It would depend a lot on the toxin, whether it could be insignificant or serious, how much it dilutes and whether it washes away with subsequent rains or adheres to surfaces.

P1: Regarding food chains, we talked about cumulative effects, but do we have any research data on how this behaves in a built urban environment, how stormwater behaves, whether it ends up at the wastewater treatment plant or remains somewhere in urban areas?

P3: I was also thinking about runoff water – that if it rains and streets flood and it all flows somewhere, a lot of it into the sea.

P4: Into the sea, but it means trouble also in the water treatment process, it can have really serious effects there.

P1: This has now become complicated with this precipitation. But maybe this was the goal [of the exercise], to see at what point we throw in the towel.'

The possibility of toxic rainfall was included in the researchers' acute situation model. However, its complex and uncertain consequences became fully apparent only through the participants' collective thinking, its systemic features being captured in the graphical representation of causal connections. This realization extended beyond just the specific scenario – it also raised questions about the city's overall preparedness and its ability to handle such events. While some participants initially speculated that the purpose of the exercise was to expose the limits of their ability to manage the scenario, they did not "*throw in the towel*". Instead, the city acknowledged the unexpected threat and later, in 2021, organized its own preparedness exercise focused on the issue.

The participants unanimously agreed that the exercise successfully transformed what was initially "*a quite daunting*" scenario with many uncertainties into something more manageable. Beyond the new insights gained during the modelling process and the concrete action points identified in discussions (see Table 1), all participants emphasized the value of the collaborative working format. According to their feedback, the exercise "*definitely enriched our thinking of the big picture*", made "*complexity tangible*" and "*clarified it*", and even "*stripped away preconceived notions*". They also noted that the collective input was an added benefit, typically difficult to achieve in their usual work routines.

> 'This method concretized complex discussions where there are many different topics by making them more tangible. It's like finding a common language. Everyone often talks about the same things, but they have slightly different perspectives. This kind of tool definitely helps to consolidate and ensure that everyone is on the same page, using a shared language and a common way of discussing the matter.'

The exercise introduced a new meta-language that connected different sectors and bridged gaps between specialized areas of expertise within the organization. This supports the idea that collectively mapping causal relationships can reduce communication-related uncertainties and help establish shared, context-specific definitions of key terms. Participants also found the experience valuable because it allowed them to see how individual perspectives came together to form a broader, systemic understanding of cause and effect. As one participant noted, "*something that you thought wasn't a big deal turned out to be quite serious*" – an insight that might not have emerged without this collaborative approach.

A participant elaborated:

> 'The exercise had a dialog-promoting effect, it clarified and calmed down the handling of such a complex problem. It allowed us to get things on paper, to make them gradually visible and adaptable and not set in stone. They are right there for us to see, they gain shape, it's clear how many there are of them, and the important ones begin to stand out. (…) Of course, we rarely have the luxury of focusing on building such large frameworks. But at least when the opportunity clearly presents itself, [this way of working] would definitely be worth using.'

While mapping causal relationships helped simplify the complexity of the disaster scenario, the process also highlighted that it is a "*rare luxury*" to "*carefully consider*" how to achieve such reductions. This suggests that the meta-language achieved in the exercise might be absent in real life, which certainly poses a risk for the city. Officials responsible for implementing strategies and managing resources often do not have the time to engage in such structured reflection in their daily work. Yet, these are the kinds of emergent insights that we suggest that the exercise concept can bring forth. Here we find a dual meaning of reflexivity prominently at play: on one hand, it can lead to discoveries about the operational environment and reveal new leverage points for managing disruptive events. On the other hand, reflexivity also exposes its own limitations, continuously redefining the boundaries of what is known and understood.

3.2.2. From the first-order model and thinking to second-order inquiries

A discussion about the city's goal of fostering a "Shared City" clearly illustrates how the modelling exercise influenced strategic thinking. When reflecting on the metrics for this goal – particularly the indicator of "engaging and empowering citizens" – one participant suggested that, rather than using these measures solely to assess the impact of the city's actions, "*it would be good to do all of this in advance, especially considering scenario events*". Another participant echoed this perspective, noting "*a subtle irony in the fact that the Shared City itself is the one improving resilience*". At this point, participants began to critically reassess their own thinking as it was translated into the model, particularly in light of the sensitivity analysis conducted on the strategy objective in question. They recognized that improving the city's resilience is not merely a byproduct of external resilience initiatives but is instead intrinsically linked to achieving their strategic goals.

At this point, the facilitator stepped in to steer the discussion back to the model. She reminded the participants that, during the initial phases of the exercise, when mapping out causal connections, the metric "Engaging and empowering" had no incoming links. In other words, in the participants' original thinking, it was not seen as influencing the progression of the scenario. This triggered further reflection as a participant remarked that

> 'I think this fits into what we are actually doing here. The scenario hasn't materialized yet. (…) So, we are initiating connections to the metric through this work. To me that seems exactly how it should be. (…) It bothers me a bit though that these strategic perspectives are now being examined one-by-one, whereas there should be some overlap to see how they emerge together.'

Here, the participant takes a step back to reflect on how the modelling process connects to the organization's real-world practices. For them, the value of the exercise lies in finding analogues between the model and the actions that should be taken in real life, i.e. "*initiating connections*" that were not originally present in their thought process. This leads the participant to recognize the broader interconnectedness of the city's strategic goals, which the computational model helped to reveal and make available for critical examination. Another participant returns to the topic towards the end of the session when talking about the planning practices of the city:

> 'It's more about how the decisions made in the past about the structures that we have, could adapt, of how we can create resilience by thinking about adaptation. From there, it should be possible to apply that same idea in more detailed planning. I think we still have a lot to learn in that regard. It's perhaps primarily about interacting across planning levels and organizations to ensure that all objectives are fully realized. But if you were to ask me now for an example where this hasn't succeeded specifically in terms of safety, I wouldn't be able to name one straight away.'

The relationship between planning details and the envisioned comprehensive, integrative approach proves to be complex. The participant realizes that existing planning practices may overlook critical cross-cutting elements, highlighting a potential leverage point for the city to address. The inability to recall past examples suggests that this type of second-order reflexivity expressed upon first-order reflection in the exercise may have been previously absent. While the participant's concern about making "*strategic perspectives emerge together*" brings out the possibility of adjusting planning and management practices, identifying clear entry points into the city organization remains difficult. As one participant reflected after completing the causal mapping in the first phase:

> 'We haven't maybe had discussions like these regarding strategy in this context. This definitely highlighted the limits of our own understanding and the knowledge and skill requirements that arise. At the same time, it raises questions about what kinds of team formations and expertise we need to assess the reliability of our built environment for example. This certainly gave us new perspective on such things.'

In other words, the exercise not only revealed previously unrecognized issues for the city but also highlighted the need to reassess and update the organization's own knowledge processes and operational approaches. The recognition of the collectively constructed systemic view as essential also came with an important realization: this perspective gained was shaped by the composition of the group and the expertise represented within it. Ultimately, the key question is *how well the organization can adapt to these second-order observations – to practice reflexivity regarding its own operations*.

3.2.3. How to bring the observations home: rethinking the city's strategy work

In this final subsection, we trace how the participants saw the exercise concept and its insights as fitting into and potentially reshaping the city's strategic outlook. The approach was seen as having three main applications: first, as *a tool for impact assessment* and for identifying causal chains and their significance; second, as *a decision*-support *tool for navigating complex situations*; and third, as a *method for developing shared strategic insights and assessing progress toward objectives*. Across all three applications, participants valued the exercise for its ability to reveal unexpected factors both in terms of external threats and the adequacy of the organization's existing preparedness and response practices.

We want to emphasize an important nuance in what initially appears as enthusiasm for the kind of thinking the exercise concept inspired. Participants identified multiple ways and contexts in which this novel method could be useful, both within their organization and in collaboration with stakeholders. However, there was also a degree of apprehension about how the method might disrupt existing practices. In the following examples, we begin to see the boundaries of the reflexivity gained through the exercise.

When discussing the finished causal map, a participant suggested several potential uses for the procedure so far:

> 'this could perhaps be used as a tool for strategic development, where we want to get different partners to collaborate. And definitely for impact assessment, which is important in urban planning and likely in many other areas as well. By identifying the chains of impact and their significance, a shared strategic vision could emerge about what we are doing and how our actions might positively or negatively affect it.'

Here the participant sees immediate applicability in the exercise concept, suggesting that it should and could be incorporated into the existing management processes of the city. The model's ability to "*identify chains of impact and their significance*" could enhance the organization's capacity to strategically account for the potential systemic and organization-wide consequences of disasters like the one explored in the exercise. However, the reflexivity observed here is limited to assessing whether the organization's strategic vision leads to the right actions – in positive or negative effects – rather than questioning whether the strategy itself is targeting the right issues. In other words, there is a key distinction between surprising threats always hiding in the blind spots of a strategic vision and the unknown as something that can be shaped through continuous adaptation and refinement of that vision.

During the final, diagnostic, phase of the exercise – when digging into the analytical outcomes of the BN – a mix of reflexive hesitation and excitement emerged as participants grappled with insights that were difficult to fully articulate.

> 'Tools of this kind haven't existed. What we are doing here is creating that … there's something here that's deeply troubling to me, something the significance of we might only see in hindsight. This requires us to take proactive measures in just the right way (…) maybe proactive work like this is the only option.'

Another participant expounded in the same vein in an earlier session:

> 'I would have liked to have this same discussion to ground some working groups of ours because this is what we should make happen. It feels frustrating to me, like, "Oh man, I have to start the work all over again now," but on the other hand now that I see the tool, this is how it should be done. This brought up many questions and discontent, but also joy from the fact that something new has been created.'

The discomfort of realizing that existing practices may have lacked a crucial foundation, whose necessity only becomes clear in hindsight, illustrates well the dual nature of reflexivity outlined earlier. It is not only the unforeseen threats that could be better anticipated when explored with a method similar to the exercise, but also the unknowns in the practice of creating strategic perspectives to such foresight. This unease was, however, overridden by the participants' various ideas about how and where the method could be applied:

> 'I was left waiting you to tell us when this will become an application where we can examine connections and the probabilities of impacts. I found that really interesting, and I'm wondering if at some point this can become a somewhat user-friendly system, where users could be presented with questions and [the model/tool] could run probability calculations. In principle, even in an accident scenario, we could think about putting in a scenario and within a few hours we'd get some feedback on what to prepare for.'

> 'Since we have this regional security forum and they conduct an annual regional risk assessment, those scenarios could very well be tested with a model like this to see how it impacts the city's operations. It would be very interesting to see how that would work.'

> 'When the city board wants a business plan for [a construction project] with certain stakeholders, this is comparable to business plan that involves a multi-strategic goal and requires compiling the inputs from different organizations onto a shared platform [such as the BN-model] where you can see how the inputs measure up to various multi-scenario situations.'

Again, the effects to be modelled are seen as one-directional: external factors impacting the organization, rather than the organization needing to model its own capacities, limitations, and ways of understanding itself. A step toward the latter, broader perspective, was taken by a participant who speculated on how the modelling process could be integrated into the city's management systems:

> 'There is a kind of self-assessment emerging from this (…). With this systematic approach, I think it would be interesting to see when we analyze what we have here in our CFI [Critical Factor Index] analysis, in our information systems, in our management systems, and the likes, we could gain insights on where the critical points are, based on our collaboration here.'

The city has established methods of self-monitoring and assessment, and the exercise raised the expectation that certain "critical points" might have become visible only now, through the modelling process. In the excerpt above and across this subsection, participants explored the idea of integrating the exercise concept or the co-constructed model into the city's existing administrative processes. This implies that the insights generated by the model could be transferred to other contexts. In other words, the thought process initiated by the PM exercise could be adaptable and applicable beyond its original setting.

4. Discussion

By utilizing a BN model as an AI tool for knowledge co-production and analysis, we have developed a novel PM approach to build understanding about connections between acute socio-environmental disruption events and their long-term strategic consequences, to increase the resilience of municipalities. Our research questions were: (1) what kind of participatory BN-modelling process could support our goals, and (2) to what extent such an approach could be incorporated into strategic risk management in contexts beyond the exercise. We formulated these questions against the background of the previous POR exercise (Järvensivu et al., 2021), which identified the focal needs of: a) more engaging, i.e. socio-politically more relevant and complex, scenarios; b) less narrow and more realistic strategy (decision) options; c) more comprehensive involvement of the participants to the planning and development of the exercise; and d) possibility to turn back and reconsider the impacts of alternative scenarios. We addressed these needs through 1) co-designing with city representatives a locally relevant starting point scenario and providing them with evidence-based inputs in the form of a BN model; 2) guiding the participants to continue the causal prognosis from acute impacts towards the longer-term strategic goals of the city and create as rich a scenario space as they wanted (in the given time frame); 3) translating the participants' mental model to a computational, but also visual, BN that could provide them with numerous possible situational pictures and allowed them quite literally see where different pathways would have taken them; 4) guiding the participants themselves, through collective deliberation of the systems analytical outputs provided by the resulting BN, identify the realistic resilience-increasing strategy options.

We hypothesized that a BN could be used to represent the participants' collective thinking about the disruption consequence pathways, and that the modelling process might further provide them with a new level of observation that could help identify critical leverage points and interventions to the city's strategic resilience to corresponding incidents. At the end of the exercise, the participants could recognize the relevance and opportunities of both reactive (acute response), and follow-up (post-treatment) interventions, as well as proactive anticipatory management actions, in controlling the probability of corresponding incidents, and their short- and long-term consequences.

The collective construction and use of the BN thus did make visible the connections between operative management of acute accidents, the evolution of long-term causal pathways, and how the consequences finally materialize on the strategy level. The exercise format enabled a dual layer of reflexivity, revealing strategic resilience as a product of this process. In other words, strategic thinking itself became part of the system under examination. This aligns with the core principle of a self-organizing system: the cognitive act of defining the system should itself fall within the system's own domain (Maturana and Varela, 1980; Luhmann, 1989; Hukkinen, 2014). In practice, the exercise allowed the participants to draw the causal connections between the city's strategic goals and all the factors affecting them. The finished computational model further presented this interconnectedness as an adjustable system that represented the participants' own collective thinking. Therefore, we argue that **unknowns** that became gradually and provisionally knowable during the exercise, were **reflexive**. That is, *they remained open for constant reflection, and this very act of reflection was seen to influence the entire system under observation and the inevitable partiality of that observation, blind spots and all.*

In their systematic analysis on BN applications in the context of environmental risk assessment, Kaikkonen et al. (2021) concluded that so far, the actual stakeholder involvement in the BN modelling processes has been fairly limited. The models are mostly focusing on quantifying risks, whereas Kaikkonen et al. (2021) ask whether they could be seen also as platforms for reflection and discussion among teams. We have demonstrated how our experimental approach to PM with BNs can engage participants in a reflexive process. What is more, the approach shows that quantification and qualitative judgment can be inherently intertwined (cf. Callon and Law, 2005). We see a difference here between risk quantification and our approach in the way expert judgment is made to matter in BN modelling. In risk quantification, expert knowledge is elicited to complement and verify existing observational data to enable a computationally robust representation of the system at hand. Our approach, on the other hand, deliberately keeps the boundary between the quantitative and the qualitative open for interpretation and updating. We suggest that this allows also a greater degree of freedom in drawing a variety of scenario paths in the model as the quantified information retains the contingency with which it was conceived.

Regarding the knowability of the system being modelled, the mutual construction of the quantitative and qualitative dimensions enables inquiries into how the composition of the system itself – what is included in it and what left out – becomes something to be known. Thus, through the gradual accumulation of insights throughout the exercise, a sense of **reflexive unknowns** emerged. By the concept, we mean observations that could be articulated only thanks to the emergent systemic mechanisms that would have remained completely unknown without the BN model and the collective work put into constructing it. A system and its connections and interactions with its environment obviously always exhaust our knowledge of them. Reflexivity towards unknowns can, however, extend and refine the scope of our knowing with the proviso that blind spots will always persist. In other words, the analysed PM exercise helped to recognize a discrepancy between newly identified “strategic risks” and the prevailing lack of management strategies to control them. The term “reflexive unknown” thus combines two different meanings of reflexivity: self-reflection that ponders the bounds of one’s own thought, the “how and why” of the way we think and act (cf. Archer, 2010), and reckoning with the risk that comes with settling on a certain kind knowing and thus foreclosing the scope of what can be known (cf. Beck, 1992). In relation to existing typologies of emergent surprises, this represents a new category of its own (cf. Ekengren, 2024). Characteristic of these reflexive unknowns is that they can *only* be reached through the diagnostic loops of second order observation that the BN-modelling process enabled – in a way that integrates quantification and qualitative judgment.

An example of a reflexive unknown is the prospect of a toxic fallout and its short and long-term implications for the built and natural environment of the city (see sections 3.2.1 and 3.2.2). The possibility of a fallout first appeared as a result of uncontrollable weather phenomena (a combination of precipitation and wind in the event of a gaseous expulsion) in the starting point model provided by the researchers. The graphic representation of causal connections discussed by the participants further led to a new collective awareness of previously unthought chains of events, whose impacts extended to three of the city’s four strategic objectives. The sensitivity analyses run on the completed model further clarified the boundaries of the collective awareness regarding the potential disastrous fallout scenario. In other words, the modelling process helped the participants to see both the limitations and possible leverage points within the city’s strategic thinking – as constructed and represented in the model. The exercise enabled moving from first-order realization of the possibility of a fallout to second-order, reflexive considerations about the city’s capacity to strategically anticipate and manage such a scenario. Ultimately, the toxic fallout actualized as the subject matter of a subsequent preparedness exercise in the city in 2021. Thus, the fallout evolved from being a complete blind spot to an initial acknowledgement of its possibility, and finally to a prospective reality with which the city organization had to contend and modify their strategic actions.

The participants’ discussion on model migration and its modular transferability to contexts beyond the exercise brought out the limitations of our approach and the reflexivity that could be achieved with it. Admittedly, although the acute accident model developed by the research team is bound to a certain location and timing, the co-developed part that linked the acute harmful effects to the long-term consequences to the city’s strategy is more general, and modules of it could potentially be applied to other contexts and data sources. What this notion of model migration and modularity (cf. van der Heide, 2020) risks neglecting, however, is the very objective of designing the exercise concept: to include and take seriously the inherent social and qualitative aspects of producing participatory BN models.

The collective, emergent insight that has gone into a modelling process should not be detached and black-boxed away from the effort and the specific composition of expertise that produced the model (e.g. Svetlova and Dirksen, 2014). Even though the algorithm to translate the participants’ descriptions of the type, strength, and uncertainty of causal links into probabilistic form is not bound to a specific context, the system and its dynamics, and especially its blind spots, certainly are. In this regard, particularly the time constraints of the exercise and our objective to make the modelling procedure as accessible as we could for the participants, the potential of the completed model remained partly untapped. The BN would have allowed a variety of further analyses and queries that the participants could have made. For the model to be able to migrate and find application in other contexts so that the participatory, collective work that it took to construct could also travel to other sites, a user interface much more intuitive and accessible for non-expert modellers would be needed.

During the exercise, the participants also expressed a desire to link observational data to the co-produced BN model. This idea is well aligned with the concept of Bayesian knowledge updating (Bayes and Price, 1763). Thus, technically speaking, the conditional probability distributions of the BN, in this exercise representing the participants’ priors: their degrees of belief, could be updated based on other types of data. However, the data should correspond precisely to the meanings of the variables and their interconnections as defined in the model. In practice, the mental and data-based representations of the same system or issue may appear quite different, as the available data inevitably influences the model’s structure (see e.g. Luoma et al., 2021, 2022). Therefore, updating the expert-driven priors with observational data would fundamentally alter the nature of this exercise. The structure defines the framing of the model and thus also the scenario space that can be explored with it, making its definition a critical task.

Eliciting the probability distributions required for the probability tables from experts and stakeholders is an art of a kind (e.g., O’Hagan et al., 2006). When dealing with Bayesian networks, elicitation becomes particularly demanding as the number of distributions contained in the conditional probability table of each variable grows exponentially with the number of parent variables and/or the number of alternative states of those parents (in discrete models). For this reason, models that rely on expert-elicited probabilities are typically kept simple, or alternatively, only part of the variables are elicited, while the remaining ones are based on simulated or observed data (e.g. Lehikoinen et al., 2013, 2015). This limits the structural complexity of the model, meaning it also constrains the scope of the research questions and scenario spaces that can be addressed.

In the presented exercise, one key goal, while constructing the model, was to avoid imposing restrictive constraints on the participants’ thinking. The model structure aimed to faithfully represent the group’s reasoning, with link elicitation based on a broad classification of relative strengths and uncertainty levels. The translation of these thought processes into probability distributions was entrusted to the modeler, who developed a novel algorithm specifically for this purpose (Appendix 1).

This approach is justified because the primary focus of the modelling exercise was not on estimating the actual likelihoods of risk events but rather on understanding the systemic mechanisms that shape the post-accident future and how to influence them most effectively. In this context, resilience emerges from understanding *how* the values (i.e., strategic objectives) of the city are at risk, and how to navigate within the risk landscape – by managing and adapting to it.

Given this focus, model validation hinges on ensuring that the model accurately captures the participants' thought processes. Specifically, when one part of the BN is modified (e.g., by setting a variable to a known state), the system's response should align with the participants' shared mental model (see Pitchforth and Mengersen, 2013) – in other words, reflect their collective reasoning in a coherent manner.

This type of validation is highly relevant as part of the presented PM approach and was conducted at multiple stages of our test exercise:

1. After the first meeting, the initial model structure (a graphical representation synthesizing the three discussions) was shared with the participants, allowing them to review and provide feedback.
2. After the second meeting, a refined version of the model – now incorporating link strengths and uncertainty markers – was distributed for participant comments.
3. During the third meeting, the quantified BN was demonstrated, along with narratives derived from its results. Participants discussed these findings, jointly confirmed their logical coherence, and collaboratively formulated potential resilience-building management actions.
4. Following the exercise, the modelling process and its outcomes were documented and shared as an internal report within the city organization, supplemented with additional analytics generated from the model. Lately, the exercise's impact was confirmed in discussions with city representatives, who reported that it had directly influenced the contents of the city's new safety strategy and preparedness exercises.

This multi-stage validation process ensured that the model remained aligned with participant reasoning and that its insights were both credible and actionable within the city's decision-making framework.

5. Conclusions

Not all uncertainties in strategic environmental management need to remain unknown or only be identifiable afterwards as lessons learned. By eliciting collective expertise, we have shown that participatory modelling with BNs, as applied in the presented exercise, can generate new insights into both the external operating environment and the internal organizational knowledge and practices that shape the management of that environment. The implication of our findings is that the exercise concept helps experts – be they urban planners, environmental specialists, risk managers, or policymakers – achieve new collective insights on their strategic planning practices and risk awareness across sectoral divides. The exercise concept allows managers to try out, assess and monitor how different compositions of social and environmental factors, organizational actions, and goals, might affect performance. Crucially, the method we propose and have tested enables practitioners to critically examine the underlying causal logic on which their collective judgment is based. By making these causal relationships visible, the approach develops systemic understanding, improving the capacity to anticipate and respond to complex socio-environmental risks.

The exercise highlights that urban crisis resilience is linked not only to infrastructures and action plans but also to the underlying values that should sustain even in the face of severe socio-environmental disruptions. Moreover, it underscores that resilience is fundamentally shaped by how actors within the city organization collaborate: how they share knowledge, foster new forms of collective understanding and awareness, and create spaces for joint reflection on potential threats. This process requires assessing not only their operational preparedness but also their underlying assumptions, decision-making processes, and ways of thinking. Through such structured reflection, city actors can uncover weaknesses and strengths – and, most crucially, previously unrecognized blind spots – and define actions that safeguard shared values proactively, during crises, and in their aftermath.

CRediT authorship contribution statement

Annukka Lehikoinen: Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Tapio Reinekoski:** Writing – review & editing, Writing – original draft, Validation, Methodology, Investigation, Data curation, Conceptualization. **Nina Janasik:** Writing – review & editing, Writing – original draft, Validation, Methodology, Investigation, Data curation, Conceptualization. **Marko Ahvenainen:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Janne I. Hukkinen:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This study has received funding from the Strategic Research Council of Finland (grant nos. 312623, 312624, 312625, 336253, 365651), Research Council of Finland (grant no. 338553, 338557), and NextGenerationEU instrument via the Research Council of Finland (grant nos. 353057, 353058).

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jenvman.2025.125373.

Data availability

The data that has been used is confidential.

References

Ahvenainen, M., Janasik, N., Lehikoinen, A., Reinekoski, T., 2021. Tulevaisuusresilienssi ja strateginen ennakointi: kriisinkestävyyden harjoittelua bayeslaisella kausaalimallinnuksella. Futura – journal of the Finnish Society for Futures Studies 4/2021. Tulevaisuustietoisuus ja tulevaisuusresilienssi, ISSN 0785-5494.

Andriatsitohaina, R.N.N., Celio, E., Llopis, J.C., Rabemananjara, Z.H., Ramamonjisoa, B. S., Grêt-Regamey, A., 2020. Participatory Bayesian network modeling to understand driving factors of land-use change decisions: insights from two case studies in northeast Madagascar. J. Land Use Sci. 15 (1). https://doi.org/10.1080/1747423X.2020.1742810.

Archer, M.S., 2010. Introduction. The reflexive re-turn. In: Archer, M.S. (Ed.), Conversations about Reflexivity. Routledge, Taylor & Francis Group, London.

Barton, D.N., Kuikka, S., Varis, O., Uusitalo, L., Henriksen, H.J., Borsuk, M., de la Hera, A., Farmani, R., Johnson, S., Linnell, J.D.C., 2012. Bayesian networks in environmental and resource management. Integrated Environ. Assess. Manag. 8 (3). https://doi.org/10.1002/ieam.1327.

Bayes, T., Price, R., 1763. An essay towards solving a problem in the doctrine of changes. Phil. Trans. 53. https://doi.org/10.1098/rstl.1763.0053.

Beck, U., 1992. Risk Society: towards a New Modernity. Sage, London.

Bromley, J., Jackson, N.A., Clymer, O.J., Giacomelloa, A.M., Jensen, F.V., 2005. The use of Hugin to develop Bayesian networks as an aid to integrated water resource planning. Environ. Model. Software 20. https://doi.org/10.1016/j.envsoft.2003.12.021.

Callon, M., Law, J., 2005. On qualculation, agency, and otherness. Environ. Plann. Soc. Space 23 (5). https://doi.org/10.1068/d343t.

Carmona, G., Varela-Ortega, C., Bromley, J., 2011. The use of participatory object-oriented bayesian networks and agro-economic models for groundwater management in Spain. Water Resour. Manag. 25 (5). https://doi.org/10.1007/s11269-010-9757-y.

Carriger, J.F., Dyson, B.E., Benson, W.H., 2018. Representing causal knowledge in environmental policy interventions: advantages and opportunities for qualitative

influence diagram applications. Integrated Environ. Assess. Manag. 14 (3). https://doi.org/10.1002/ieam.2027.
Carriger, J.F., Parker, R.A., 2021. Conceptual Bayesian networks for contaminated site ecological risk assessment and remediation support. J. Environ. Manag. 278 (2). https://doi.org/10.1016/j.jenvman.2020.111478.
Castelletti, A., Soncini-Sessa, R., 2007. Bayesian Networks and participatory modelling in water resource management. Environ. Model. Software 22. https://doi.org/10.1016/j.envsoft.2006.06.003.
Düspohl, M., Frank, S., Döll, P., 2012. Review of bayesian networks as a participatory modeling approach in support of sustainable environmental management. J. Sustain. Dev. 5 (12). https://doi.org/10.5539/jsd.v5n12p1.
Ekengren, M., 2024. Why are we surprised by extreme weather, pandemics and migration crises when we know they will happen? Exploring the added value of contingency thinking. J. Contingencies Crisis Manag. 32 (1). https://doi-org.libproxy.tuni.fi/10.1111/1468-5973.12515.
Esposito, E., 2024. Can we use the open future? Preparedness and innovation in times of self-generated uncertainty. Eur. J. Soc. Theor 27 (2), 209–224. https://doi.org/10.1177/13684310231224546.
Fazey, I., Schäpke, N., Caniglia, G., Patterson, J., Hultman, J., Wyborn, C., 2018. Ten essentials for action-oriented and second order energy transitions, transformations and climate change research. Energy Res. Social Sci. 40. https://doi.org/10.1016/j.erss.2017.11.026.
Fischer, J., Riechers, M., 2019. A leverage points perspective on sustainability. People and Nature 1. https://doi.org/10.1002/pan3.13.
Helle, I., Ahtiainen, H., Luoma, E., Hänninen, M., Kuikka, S., 2015. A probabilistic approach for a cost-benefit analysis of oil spill management under uncertainty: a Bayesian network model for the Gulf of Finland. J. Environ. Manag. 158. https://doi.org/10.1016/j.jenvman.2015.04.042.
Hichert, T., Biggs, R., de Vos, A., Peterson, G., 2021. Scenario development. In: Biggs, R., de Vos, A., Preiser, R., Clements, H., Maciejewski, K., Schlüter, M. (Eds.), The Routledge Handbook of Research Methods for Social-Ecological Systems, first ed. Routledge. https://doi.org/10.4324/9781003021339. 2021.
Hukkinen, J., 1993. Bayesian analysis of agricultural drainage problems in California's San Joaquin valley. J. Environ. Manag. 37, 183–200. https://doi.org/10.1006/jema.1993.1015.
Hukkinen, J.I., 2014. Model of the social–ecological system depends on model of the mind: contrasting information-processing and embodied views of cognition. Ecol. Econ. 99, 100–109. https://doi.org/10.1016/j.ecolecon.2014.01.017.
Hukkinen, J., Eronen, J.T., Janasik, N., Kuikka, S., Lehikoinen, A., Lund, P.D., Räisänen, H., Virtanen, M.J., 2022. The policy operations room: analyzing path-dependent decision-making in wicked socio-ecological disruptions. Saf. Sci. 146. https://doi.org/10.1016/j.ssci.2021.105567.
Hukkinen, J.I., Vihma, P., Janasik, N., Mäkelä, A., Reinekoski, T., Järvensivu, P., Lähde, V., Parviainen, T., Lehikoinen, A., Salomaa, A., 2024. Dashboard for narrating to policymakers the urgency of preparedness for chronic climate crisis (preprint). https://doi.org/10.2139/ssrn.4792781.
Janasik, N., 2021. Plausibility designs, or a probabilistic and predictive take on scenario effectiveness. Futures: The journal of policy, planning and futures studies 127 (9). https://doi.org/10.1016/j.futures.2020.102670.
Järvensivu, P., Räisänen, H., Hukkinen, J.I., 2021. A simulation exercise for incorporating long-term path dependencies in urgent decision-making. Futures 132. https://doi.org/10.1016/j.futures.2021.102812.
Kaikkonen, L., Parviainen, T., Rahikainen, M., Uusitalo, L., Lehikoinen, A., 2021. Bayesian networks in environmental risk assessment: a review. Integrated Environ. Assess. Manag. 17 (1). https://doi.org/10.1002/ieam.4332.
Korb, K.B., Nicholson, A.E., 2010. Bayesian Artificial Intelligence. CRC Press.
Kuikka, S., Hildén, M., Gislason, H., Hansson, S., Sparholt, H., Varis, O., 1999. Modelling environmentally driven uncertainties in Baltic cod (Gadus morhua) management by Bayesian influence diagrams. Can. J. Fish. Aquat. Sci. 56. https://doi.org/10.1139/f98-206.
LaMere, K., Mäntyniemi, S., Vanhatalo, J., Haapasaari, P., 2020. Making the most of mental models: advancing the methodology for mental model elicitation and documentation with expert stakeholders. Environ. Model. Software 124. https://doi.org/10.1016/j.envsoft.2019.104589.
Lehikoinen, A., Luoma, E., Mäntyniemi, S., Kuikka, S., 2013. Optimizing the recovery efficiency of Finnish oil combating vessels in the gulf of Finland using Bayesian networks. Environ. Sci. Technol. 47. https://doi.org/10.1021/es303634f.
Lehikoinen, A., Helle, I., Klemola, E., Mäntyniemi, S., Kuikka, S., Pitkänen, H., 2014. Evaluating the impact of nutrient abatement measures on the ecological status of coastal waters: a Bayesian network for decision analysis. Int. J. Multicriteria Decis. Mak. (IJMCDM) 4 (2). https://doi.org/10.1504/IJMCDM.2014.060426.
Lehikoinen, A., Hänninen, M., Storgård, J., Luoma, E., Mäntyniemi, S., Kuikka, S., 2015. A Bayesian network for assessing the collision induced risk of an oil accident in the Gulf of Finland. Environ. Sci. Technol. 49 (9). https://doi.org/10.1021/es501777g.
Ludwig, D., Hilborn, R., Walters, C., 1993. Uncertainty, resource exploitation, and conservation: lessons from history. Ecol. Appl. 3 (4). http://www.jstor.org/stable/1942074.
Luhmann, N., 1989. Ecological Communication. Polity Press, Cambridge.
Luoma, E., Nevalainen, L., Altarriba, E., Helle, I., Lehikoinen, A., 2021. Developing a conceptual influence diagram for socio-eco-technical systems analysis of biofouling management in shipping – a Baltic Sea case study. Mar. Pollut. Bull. 170. https://doi.org/10.1016/j.marpolbul.2021.112614.
Luoma, E., Laurila-Pant, M., Altarriba, E., Nevalainen, L., Helle, I., Granhag, L., Lehtiniemi, M., Srėbalienė, G., Olenin, S., Lehikoinen, A., 2022. A multi-criteria decision analysis model for ship biofouling management in the Baltic Sea. Sci. Total Environ. 852. https://doi.org/10.1016/j.scitotenv.2022.158316.
Luoma, E., Parviainen, T., Haapasaari, P., Lehikoinen, A., 2024. Sustainability as a shared objective? Stakeholders' interpretations on the sustainable development of marinas in the Gulf of Finland. Ocean Coast Manag. 254. https://doi.org/10.1016/j.ocecoaman.2024.107197.
Madsen, A.L., Jensen, F., Kjaerulff, U.B., Lang, M., 2005. The Hugin Tool for probabilistic graphical models. Int. J. Artif. Intell. Tool. 14. https://doi.org/10.1142/S0218213005002235.
Mäntyniemi, S., Haapasaari, P., Kuikka, S., Parmanne, R., Lehtiniemi, M., Kaitaranta, J., 2013. Incorporating stakeholders‘ knowledge to stock assessment: central Baltic herring. Can. J. Fish. Aquat. Sci. 70. https://doi.org/10.1139/cjfas-2012-0316.
Maturana, H.R., Varela, F.J., 1980. Autopoiesis and Cognition: the Realization of the Living. D. Reidel, Dordrecht, Boston, London.
Mayfield, H.J., Eberhard, R., Baker, C., Baresi, U., Bode, M., Coggan, A., Dean, A.J., Deane, F., Hamman, E., Jarvis, D., Loechel, B., Taylor, B.M., Stevens, L., Vella, K., Helmstedt, K.J., 2023. Designing an expert-led Bayesian network to understand interactions between policy instruments for adoption of eco-friendly farming practices. Environ. Sci. Pol. 141. https://doi.org/10.1016/j.envsci.2022.12.017.
Moglia, M., Alexander, K.S., Thephavanh, M., Thammavong, P., Sodahak, V., Khounsy, B., Vorlasan, S., Larson, S., Connell, J., Case, P., 2018. A Bayesian network model to explore practice change by smallholder rice farmers in Lao PDR. Agric. Syst. 164. https://doi.org/10.1016/j.agsy.2018.04.004.
Müller, L., Czymai, M., Blättel-Mink, B., Döll, P., 2024. How to assess conditions for the acceptance of climate change adaptation measures by applying implementation probability Bayesian Networks in participatory processes. Environ. Model. Software 181. https://doi.org/10.1016/j.envsoft.2024.106188.
O'Hagan, A., Buck, C.E., Daneshkhah, A., Eiser, J.R., Garthwaite, P.H., Jenkinson, D.J., Oakley, J.E., Rakow, T., 2006. Uncertain Judgements: Eliciting Experts' Probabilities. Wiley.
Parviainen, T., Goerlandt, F., Helle, I., Haapasaari, P., Kuikka, S., 2021. Implementing Bayesian networks for ISO 31000:2018-based maritime oil spill risk management: state-of-art, implementation benefits and challenges, and future research directions. J. Environ. Manag. 278 (1). https://doi.org/10.1016/j.jenvman.2020.111520.
Parviainen, T., Kuikka, S., Haapasaari, P., 2022. Enhancing science-policy interface in marine environmental governance: oil spill response models as boundary objects in the Gulf of Finland, Baltic Sea. Mar. Pol. 135. https://doi.org/10.1016/j.marpol.2021.104863.
Pihlajamäki, M., Helle, I., Haapasaari, P., Sarkki, S., Kuikka, S., Lehikoinen, A., 2020. Catching the future: applying Bayesian belief networks to exploratory scenario storylines to assess the long-term changes in Baltic herring and salmon fisheries. Fish Fish. 21 (4). https://doi.org/10.1111/faf.12461.
Pitchforth, J., Mengersen, K., 2013. A proposed validation framework for expert elicited Bayesian Networks. Expert Syst. Appl. 40 (1). https://doi.org/10.1016/j.eswa.2012.07.026.
Penman, T.D., Cirulis, B., Marcot, B.G., 2020. Bayesian decision network modeling for environmental risk management: a wildfire case study. J. Environ. Manag. 270. https://doi.org/10.1016/j.jenvman.2020.110735.
Peterson, G.D., Cumming, G.S., Carpenter, S.R., 2003. Scenario planning: a tool for conservation in an uncertain world. Conserv. Biol. 17. https://doi.org/10.1046/j.1523-1739.2003.01491.x.
Reinekoski, T., 2025. Governing as if the future obliged the present: a scenario exercise in prefigurative climate policy. Polit. Stud. 0 (0). https://doi.org/10.1177/00323217241311352.
Salomaa, A., Reinekoski, T., Salminen, H., Krivochenitser, K., Hukkinen, J.I., Lehtonen, T.-K., 2025. Disjointed modes of building resilience to socio-environmental crises. Ecol. Soc. 30 (1), 6. https://doi.org/10.5751/ES-15766-300106.
Schindler, D.E., Hilborn, R., 2015. Prediction, precaution, and policy under global change. Science 347. https://doi.org/10.1126/science.1261824.
Singto, C., Fleskens, L., Vos, J., Quinn, C., 2020. Applying Bayesian belief networks (BBNs) with stakeholders to explore and codesign options for water resource interventions. Sustainable Water Resource Management 6 (23). https://doi.org/10.1007/s40899-020-00383-x.
Sperotto, A., Molina, J.-L., Torresan, S., Critto, A., Marcomini, A., 2017. Reviewing Bayesian Networks potentials for climate change impacts assessment and management: a multi-risk perspective. J. Environ. Manag. 202 (1). https://doi.org/10.1016/j.jenvman.2017.07.044.
Star, S.L., Griesemer, J.R., 1989. Institutional ecology, 'Translations' and boundary objects: amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. Soc. Stud. Sci. 19 (3). https://doi.org/10.1177/030631289019003001.
Svetlova, E., Dirksen, V., 2014. Models at work – models in decision making. Sci. Context 27 (4). https://doi.org/10.1017/S0269889714000209.
Uusitalo, L., Kuikka, S., Kauppila, P., Söderkultalahti, P., Bäck, S., 2011. Assessing the roles of environmental factors in coastal fish production in the Northern Baltic sea: a Bayesian network application. Integrated Environ. Assess. Manag. 8 (3). https://doi.org/10.1002/ieam.180.
van der Heide, A., 2020. Model migration and rough edges: British actuaries and the ontologies of modelling. Soc. Stud. Sci. 50 (1). https://doi.org/10.1177/0306312719893465.
Voinov, A., Bousquet, F., 2010. Modelling with stakeholders. Environ. Model. Software 25 (11). https://doi.org/10.1016/j.envsoft.2010.03.007.
Voinov, A., Jenni, K., Gray, S., Kolagani, N., Glynn, P.D., Smajgl, A., 2018. Tools and methods in participatory modeling: selecting the right tool for the job. Environ. Model. Software 109. https://doi.org/10.1016/j.envsoft.2018.08.028.
Zellner, M., Campbell, S.D., 2015. Planning for deep-rooted problems: what can we learn from aligning complex systems and wicked problems? Plann. Theor. Pract. 16 (4). https://doi.org/10.1080/14649357.2015.1084360.