Talking about the likelihood of risks:
An agent-based simulation of discussion processes in risk workshops

A.1 The ODD+D for the ABM model of a risk workshop

Outline		**Guiding questions**	**ODD+D Model description**
I) Overview	I.i Purpose	I.i.a What is the purpose of the study?	The study is a theoretical exploration of the drivers of effectiveness of risk assessments in risk workshops regarding the correctness and required time. Specifically, we model the limits to information transfer, incomplete discussions, group characteristics, and interaction patterns and investigate their effect on risk assessment in risk workshops.
		I.ii.b For whom is the model designed?	The model aims to guide facilitators of risk workshops in understanding the design choices and trade-offs they face. The model also provides a blueprint to use ABM for simulating discussion processes in organizations.
	I.ii Entities, state variables, and scales	I.ii.a What kinds of entities are in the model?	This ABM has 9 agents, all of whom are participants in a risk workshop. In each simulation run, one risk task is generated in the form of a Bayesian network. The participants perform a discussion, during which they exchange information about the risk. During the discussion, one agent (chosen at random) takes a special role as a leader, in addition to the role as a participant. Additionally, during the discussion process, the role of a facilitator is needed. However, as the role of the facilitator only requires information available to all participants or any external observer, the role of the facilitator is not implemented as an agent entity itself but only as a function. Overall, the following entities are present in the model: • Risk workshop participants that discuss a risk in order to enable a correct risk assessment; • A leader, who is one of the participants in the risk workshop but makes the final decision; • A facilitator, who makes decisions regarding the proceeding of the workshop; • A risk, modeled as a Bayesian network that is assessed in the risk workshop.

		I.ii.b By what attributes (i.e. state variables and parameters) are these entities characterized?	• Risk workshop participants (including the leader):
			  ○ Initial information on risk
			  ▪ Participants are provided with some information about the risk during the initialization. The pool of information that is assigned to the participants corresponds to the information nodes in the figure below. If participants are provided information, they learn the true state of the corresponding information node (they also can know of the existence of an information node without knowing the true state, c.f. II.vi.c). To provide some overlap in expertise, 18 out of the 27 pieces of information are provided twice. Thus, 45 pieces of information are provided in total to the participants, and each information is provided to at least one participant. Participants are considered experts regarding the information they are initially provided.
			  ○ Knowledge of the risk structure
			  ▪ Participants have a Bayesian network (Pearl, 2008) as their mental model of the risk. The Bayesian network of the participants is initially limited to structures related to their expertise (c.f. II.vi.c). After participants receive new information, they update their individual Bayesian network either by adding new nodes or by changing the probabilities of the states of nodes.
			  ○ Level of hierarchy
			  ▪ Participants are randomly assigned to one of three levels of hierarchy: low, medium, or high. There are always 3 participants for each level of hierarchy. The leader is randomly chosen from the participants with a high level of hierarchy.
			• Risk:
			  ○ Risk structure
			  ▪ A risk is a Bayesian network with an overall risk assessment node and nodes representing 3 domains, 9 issues (3 per domain) and 27 information (3 per issue). See the figure below for the Directed Acyclic Graph of the Bayesian network used to represent a risk task.
			  ▪ Each node has three states (low, medium, high) with an associated likelihood for each state (see, e.g., Kabir *et al.*, 2015 for a similarly structured risk network).
			  ○ Information about the risk

			▪ Each of the 27 information nodes has a true state (low, medium, or high). The true states of the information nodes are used to calculate the benchmark risk assessment. The benchmark risk assessment is the state of the overall risk assessment node with the highest activation when all information nodes are in the state that is considered "true" for the risk task (each node of the network has three states corresponding to low, medium and high. Each of these states has an 'activation' between 0 and 1, so that the sum of the activations of all three states is 1). See A.3 for an explanation of how the Bayesian network is calibrated and A.4 for examples of how the overall risk assessment is derived from the states of the information nodes. 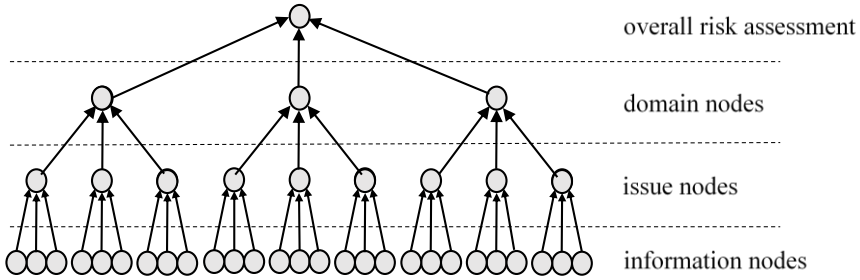

		I.ii.c What are the exogenous factors / drivers of the model?	Several simulation experiments are conducted with the model. Depending on the experiment, the following attributes are systematically varied:
			
			• Decision rules used by the leader to make decisions.
			• Initial distribution of information among the risk workshop participants
			  ○ Either all participants get the same amount of information or information is assigned unequally.
			  Specifically, each information from the pool of information that are initially assigned to the group are distributed one after another. For each information, each participant has a specific probability of receiving this information. If information is distributed equally, each agent's probability of receiving a piece of specific information is 11.1% (i.e., 100%/9). Otherwise, the probabilities are chosen so that the best-informed participant has twice the chance of receiving

			any information as the second-best informed participant (factor 2), and so on. Participants are chosen at random regarding their place in the knowledge distribution (i.e., if they are more or less well informed). • Consideration of hierarchy by the participants: Participants can weigh decisions higher or lower, depending on the position of the sender in the hierarchy relative to their own. • Presence of transactive memory within the group If transactive memory is present, participants have the information if the sender is an expert in the information being sent and can weigh the input higher or lower (depending on their own expertise). • Interaction pattern within the group
		I.ii.d If applicable, how is space included in the model?	Space is not included in the model.
		I.ii.e What are the temporal and spatial resolutions and extents of the model?	One time step is one discussion round, meaning that a participant is chosen to speak, the speaker provides a piece of information, and all other participants process the information provided to them. See I.iii.a for the steps that happen during one discussion round. The discussion is simulated for 140 rounds of discussion. However, most discussions reach a stable state way before this threshold.
	I.iii Process overview and scheduling	I.iii.a What entity does what, and in what order?	The participants repeatedly share information about the risk. The facilitator chooses one participant as a sender. All other participants become receivers. The discussion continues until the leader's decision and termination rule is activated.

			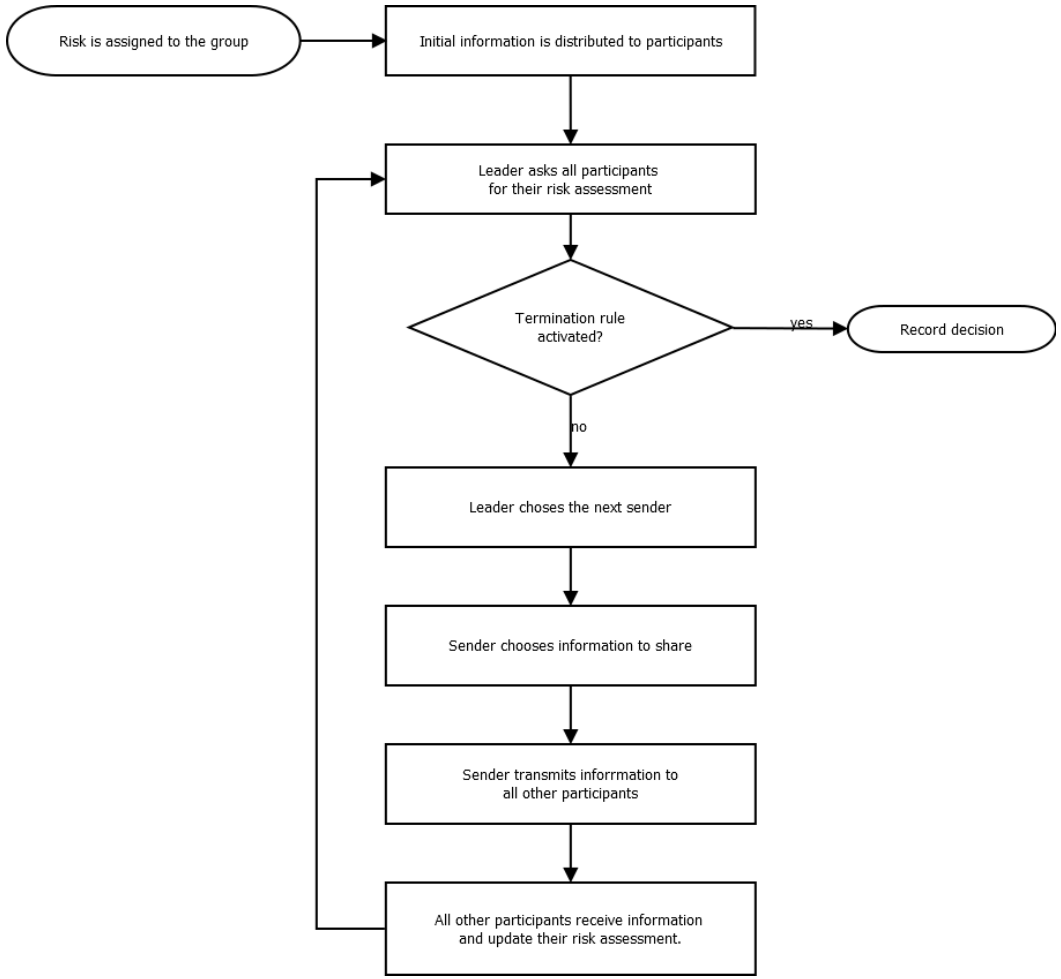
II) Design Concepts	II.i Theoretical and Empirical Background	II.i.a Which general concepts, theories or hypotheses are underlying the model's design at the system level or at the	The model has been developed to investigate the impact of choices made during the facilitation of a risk workshop. While some results are specific to decision-making regarding risk assessment, the general concept can be transferred to other settings where a group shares distributed knowledge in a discussion to make a decision.

		level(s) of the submodel(s) (apart from the decision model)? What is the link to complexity and the purpose of the model?	The experiments concern deviations from an ideal speech situation (Habermas, 1982) that serves as a baseline for the effectiveness of the assessment process. See Section 2.2 of the paper for a discussion of the deviations from this ideal speech situation included in the model. The interaction of the participants is built upon the idea of participants forming a transactive memory system (Wegner, 1987); all information is available to the group, but the group needs to correctly make use of the available information.
		II.i.b On what assumptions is/are the agents' decision model(s) based?	The participants' decisions are based on heuristics that aim to reasonably account for all available information relevant to the respective decision. The participants are either following simple mathematical formulas for their decisions (e.g., when the leader determines if the discussion should continue), or make random choices, where the probability of each option to be chosen is determined by model parameters (e.g., when participants decide what information to share with the others, c.f. III.iv.a).
		II.i.c Why is a/are certain decision model(s) chosen?	As we model human behavior, there are no clear-cut rules available for how our participants will decide to act during the discussion. Thus, when participants need to make decisions, we allow participants to choose randomly between all possible options (e.g., when they decide who should be the next sender). However, the probability of each choice is influenced by reasonable heuristics (e.g., when the group is aware of the hierarchy, the facilitator should be more likely to choose agents with a high position in the hierarchy as the next senders).
		II.i.d If the model / a submodel (e.g., the decision model) is based on empirical data, where does the data come from?	The model is not based on empirical data.
		II.i.e At which level of aggregation were the data available?	Not applicable
	II.ii Individual Decision Making	II.ii.a What are the subjects and objects of decision-making? On which level of aggregation is decision-making modeled? Are multiple	The facilitator makes one decision: • Each round, one participant is chosen by the facilitator to be the next sender. The leader makes two decisions:

		levels of decision making included?	• After each round, the leader decides if the conditions for the decision and termination approach are met, and therefore, a decision can be made; • The leader makes the final decision on how to assess the risk. Participants make two types of decisions: • What information to talk about when they are chosen to be senders for a discussion round; • How to weigh the input they receiver from the sender if they are chosen to be receivers for a discussion round.
		II.ii.b What is the basic rationality behind agents' decision-making in the model? Do agents pursue an explicit objective or have other success criteria?	Participants make their best effort to gain a correct understanding of the risk in order to reach an accurate risk assessment. The objective of the leader is to reach an accurate risk assessment in the shortest possible time. However, the criteria to end the discussion is not chosen by the leader, but is a model parameter.
		II.ii.c How do agents make their decisions?	Participants make their decisions by random choice; however, the probability of each decision might not be equal. For example, participants are more likely to talk about information they are experts on.
		II.ii.d Do the agents adapt their behavior to changing endogenous and exogenous state variables? And if yes, how?	They do not.
		II.ii.e Do social norms or cultural values play a role in the decision-making process?	Social norms influence the next sender's selection (e.g., prioritize participants based on hierarchy) in some simulation experiments.
		II.ii.f Do spatial aspects play a role in the decision process?	No.

		II.ii.g Do temporal aspects play a role in the decision process?	The leader might decide to end the discussion and make a decision if the group has not progressed for some time.
		II.ii.h To which extent and how is uncertainty included in the agents' decision rules?	As there is no clear-cut rule governing the participants' decisions, participants' uncertainty regarding the best choices is accounted for by implementing decisions as (partly) random (c.f. III.iv.a).
	II.iii Learning	II.iii.a Is individual learning included in the decision process? How do individuals change their decision rules over time as consequence of their experience?	The agents understanding of the risk under assessment is modeled as a Bayesian network. Participants update their understanding of the risk, both concerning the risk structure and information about the risk, based on the input they receive during the discussion from other participants. Their understanding of the risk determines their individual risk assessment. They do not learn beyond the discussion of an individual risk, i.e., there is no interaction spanning over several simulated discussions.
		II.iii.b Is collective learning implemented in the model?	By exchanging their individual knowledge, the participants' understandings of the risk move towards a joint understanding.
	II.iv Individual Sensing	II.iv.a What endogenous and exogenous state variables are individuals assumed to sense and consider in their decisions? Is the sensing process erroneous?	None
		II.iv.b What state variables of which other individuals can an individual perceive? Is the sensing process erroneous?	In simulation experiment 3, participants know other participants' relative position in the hierarchy and other participants' expertise regarding topics. This knowledge, if available, is not erroneous. The facilitator and the leader know the overall risk assessment of all participants and can chose the next sender and decide when to end the discussion based on this knowledge (see also the flowchart in I.iii.a).

			Knowledge and information about the risk is explicitly exchanged in a simulated discussion.
		II.iv.c What is the spatial scale of sensing?	Not applicable
		II.iv.d Are the mechanisms by which agents obtain information modeled explicitly, or are individuals simply assumed to know these variables?	Participants are assumed to know the expertise and hierarchy of other participants (if the corresponding experimental condition is present, c.f. simulation experiment 3). Participants learn about other participants' understanding of the risk via the simulated discussion.
		II.iv.e Are costs for cognition and costs for gathering information included in the model?	The cost of cognition and information gathering is accounted for by making a decision as soon as specific decision and termination criteria are met, instead of continuing the discussion potentially infinitely. However, the participants make no conscious decision on whether to invest cognitive resources: whenever they get new input, they update their knowledge about the risk and their risk assessment.
	II.v Individual Prediction	II.v.a Which data uses the agent to predict future conditions?	Participants do not predict future conditions.
		II.v.b What internal models are agents assumed to use to estimate future conditions or consequences of their decisions?	Not applicable
		II.v.c Might agents be erroneous in the prediction process, and how is it implemented?	Not applicable

	II.vi Interaction	II.vi.a Are interactions among agents and entities assumed as direct or indirect?	The participants directly interact with each other by exchanging information about the risk during the discussion.
		II.vi.b On what do the interactions depend?	In each discussion round, one participant is chosen to talk to all other participants (c.f. I.iii). The other participants are only receivers for that discussion round but can become senders themselves in subsequent rounds.
		II.vi.c If the interactions involve communication, how are such communications represented?	Participants communicate by exchanging information about their individual knowledge about the risk. Usually, they will share the activation of each state of the information node they have decided to share. However, initially, participants do not know the complete structure of the risk Bayesian network. They are only aware of information nodes provided to them and all issue and domain nodes that these information nodes belong to. Additionally, for all issue nodes, domain nodes, and the overall risk assessment node that are now present, all missing direct connected nodes are added so that each issue node is connected to three information nodes, each domain node is connected to three issue nodes, and the overall risk assessment node is connected to three domain nodes. For example, if a participant is initially provided with only one piece of information, the Bayesian network of the participant will contain 3 information nodes (for the information provided and the other two information from the same issue), 3 issue nodes (one as the parent of the information node, and the two other issues that from the same domain), all domain nodes and the overall assessment node. The two information nodes that represent information that is not provided will be initialized with the default values (derived from parameter BN_factor_lmr), as will all other nodes with no descendants (in this example, the two additional issue nodes and the two additional domain nodes). Therefore, if not all receivers are aware of the existence of an information node the sender wants to talk about, the sender will use the discussion round to communicate information about the risk structure instead.

			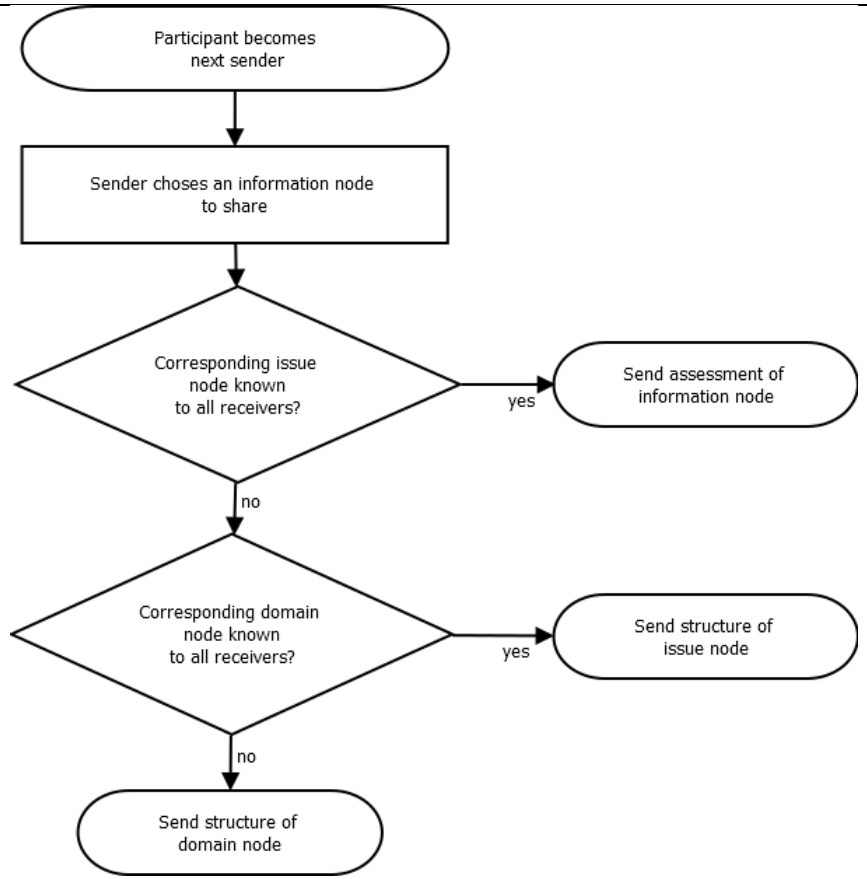
		II.vi.d If a coordination network exists, how does it affect the agent behaviour? Is the structure of the network imposed or emergent?	The setting for the simulation is one discussion by all participants. Therefore, each participant can send information to all other participants if chosen to be the sender.

	II.vii Collectives	II.vii.a Do the individuals form or belong to aggregations that affect, and are affected by, the individuals? Are these aggregations imposed by the modeller or do they emerge during the simulation?	The individual participants form a group. The risk assessments of the individuals can be aggregated to a group opinion, like a consensus, a majority vote, or an average vote (used to determine if the group opinion is moving over time). The leader decides to end the discussion based on such aggregates.
		II.vii.b How are collectives represented?	The collective(that is, the group of all participants) has no agency by itself and is only a conceptual component of the model. Decisions are made by the individual participants.
	II.viii Heterogeneity	II.viii.a Are the agents heterogeneous? If yes, which state variables and/or processes differ between the agents?	The group participants differ in their initial knowledge and expertise. They might differ regarding their position in a hierarchy if the corresponding experimental condition is present. Notably, the participants initially have heterogeneous mental models, as the risk structure of their mental models depends on the information provided to them.
		II.viii.b Are the agents heterogeneous in their decision-making? If yes, which decision models or decision objects differ between the agents?	Agents are not heterogeneous in their decision-making.
	II.ix Stochasticity	II.ix.a What processes (including initialization) are modeled by assuming they are random or partly random?	• For each run, a randomly generated risk is chosen for the group to assess. • The decision by the facilitator of who is the next sender in a discussion round is random. The probability of each participant becoming the next sender can be unequal, depending on experimental settings. • The decision by the sender what to share is random. The probability of each node to be chosen for sharing can be unequal, depending on experimental settings.

	II.x Observation	II.x.a What data are collected from the ABM for testing, understanding, and analyzing it, and how and when are they collected?	At the beginning of each run, the independent variables are recorded. After each discussion round, the activation of all states of the 'overall assessment' node is recorded for each participant. Furthermore, it is checked and recorded if a decision and termination rule (see chapter 3.3.2 of the paper) has been activated during the most recent discussion round and what decision by the leader it would result in.
		II.x.b What key results, outputs or characteristics of the model are emerging from the individuals? (Emergence)	We can observe how limitations to the ideal speech situation impact the effectiveness of risk workshops, measured as the share of high and low risks that get assessed correctly, as well as the time it takes to make a decision.
III) Details	II.i Implementation Details	III.i.a How has the model been implemented?	The simulation is written in Python. The *pgmpy* library is used to create risks by sampling from the Bayesian reference network (c.f. A3 for the Directed Acyclic Graph of the network and its calibration). The Bayesian networks are calculated in R using the *bnlearn* library.
		III.i.b Is the model accessible and if so where?	The model is be published on the OpenABM platform.
	III.ii Initialization	III.ii.a What is the initial state of the model world, i.e., at time t=0 of a simulation run?	<u>The benchmark assessment and the benchmark process</u> The simulation requires a benchmark assessment that the assessments reached in the risk workshop can be compared against. The benchmark assessment is determined by calculating a Bayesian network with full information about the true risk. This is the assessment that a group would reach given an ideal speech situation (the benchmark process). <u>Initial information distribution</u> In the beginning, information on the risk is provided to the participants. After this initial seeding of (true) information, participants only get new input from other participants. All participants are able to assess the overall risk based on the limited knowledge they are provided initially. Each information about the risk is available to at least one participant so that, in an ideal speech situation, a correct risk assessment is achievable. *initial_distr_info* bits of information are distributed among the participants, either equally or unequally (see III.ii.b). <u>Knowledge about each other's hierarchical position</u>

			Each participant is assigned one of three hierarchical statuses (low, medium, high). If the corresponding experimental condition is present, participants are aware of the hierarchical status of the other participants and consider it when including sender input into their mental model. <u>Knowledge about each other's expertise</u> The information provided to the participants initially constitutes their expertise regarding the risk. Depending on the experimental condition, participants might be informed about each other's expertise, i.e., who is an expert on which information.
		III.ii.b Is initialization always the same, or is it allowed to vary among simulations?	<u>Risk task provided to the participants</u> A new risk is generated for each simulation run, and information is assigned to the participants in a randomized process. <u>Initial information distribution</u> Depending on the experimental condition, information is provided equally or unequally to the participants. If the information is provided equally, each participant has the same probability of receiving any bit of information. If the information is provided unequally, the probability (P) of participant *i* (out of n participants) to be provided information *j* is: $$P(i,j,n) = \frac{2^i}{2^n - 1}$$ <u>Knowledge about each other's hierarchical position</u> Depending on the experimental condition, participants are provided with information about each other's position in the hierarchy, allowing the participants to consider it during the inclusion of new information. <u>Knowledge about each other's expertise</u> Depending on the experimental condition, participants are provided with information about each other's expertise, allowing the participants to consider it during the inclusion of new information.

		III.ii.c Are the initial values chosen arbitrarily or based on data?	The initial values are chosen arbitrarily within the constraints set by the model parameter.
	III.iii Input Data	III.iii.a Does the model use input from external sources such as data files or other models to represent processes that change over time?	The model does not use external sources for input.
	III.iv Submodels	III.iv.a What, in detail, are the submodels that represent the processes listed in 'Process overview and scheduling'?	<u>Sender selection</u> The sender is chosen by the facilitator in a random draw that depends on the experimental condition. Only agents who have any expertise in any node are eligible to become senders for technical reasons. See chapter 3.3.4 for the different interaction patterns investigated. Random: In the baseline model, the next participant to speak is chosen at random, with an equal probability for each participant. Priority given to concern: The probability of each participant to be chosen is weighted by the value a participant assigns to the 'high' state of the overall risk assessment node. Priority given to dissent: The probability of each participant to be chosen is weighted by the distance of their individual risk assessment from the average group risk assessment. This distance is calculated as the sum of the absolute differences between the activation of the low-, medium-, and high states of the overall risk assessments by the individual participants and the average group risk assessment. Priority given to hierarchy: Participants are more likely to be the sender if they are assigned a higher hierarchical position. The probability of each participant to be chosen is weighted by *weight_h_low*, *weight_h_medium*, or *weight_h_high*. Priority given to homogeneity: The probability to be the next sender is higher if the participant's risk assessment is close to the average group risk assessment. The distance is calculated in the same way as when priority is given to dissent, however the probability of each participant to be chosen is weighted by the reciprocal of the distance to the group assessment.

			<u>Selection of sender output</u> The senders chose one of the information nodes available to them. The chance of choosing a specific node is weighted by the sender's expertise regarding the information node. The expertise is *weight_e_expert* for nodes initially assigned to the sender and *weight_e_nonexpert* for all other nodes. The sender will provide the detailed state of the information node to all other participants, who are the receivers in the given round. Furthermore, participants prioritize talking about nodes that have not been talked about often before. The previously determined weight for each node ($node_i$) based on the participant's expertise regarding the node is multiplied by a weight reflecting how much it has been talked about by the group already, relative to the node that was talked about most often ($node_{max}$): $weight(node_i) = weightExp(node_i) * 2^{discussionCounter(node_{max}) - discussionCounter(node_i)}$ Here, weightExp($node_i$) is either *weight_e_expert* if the sender is an expert on $node_i$ or *weight_e_nonexpert* otherwise. $discussionCounter(node_i)$ gives the number of times $node_i$ has already been discussed in this simulation run. If the sender wants to transmit an information node that is not part of the network of all receivers, the sender will instead send information on the first parent node that is not available to all receivers (see II.vi.c). <u>Inclusion of sender input</u> All receivers update their believes about the node the sender talks about by assigning the input a weight relative to their own, prior believe. Depending on the experimental condition, the weight assigned to the input reflects the following aspects: • difference in expertise regarding the specific node between sender and receiver; • difference in hierarchy between sender and receiver. For example, the receiver will weigh the input higher if the sender is an expert and the receiver is not, or if the sender has a higher position in the hierarchy compared to the receiver.

			For example, if a non-expert assesses the likelihood of an information node to be high as 0% and received the information of an expert that assesses the likelihood to be 23%, the updated belief of the non-expert will be $$23\% * (\text{weight_e_expert} - \text{weight_e_nonexpert}) - 0\% * (1 - (\text{weight_e_expert} - \text{weight_e_nonexpert}))$$ The model also allows weighing the input based on the trust the receiver has in the sender. For this study, however, the trust is not varied and identical between all participants.
		III.iv.b What are the model parameters, their dimensions and reference values?	See A.2
		III.iv.c How were submodels designed or chosen, and how were they parameterized and then tested?	The submodels where designed to reflect a prototypical implementation of the processes that happen during a risk workshop. The experimental conditions were derived as deviations from an ideal speech situation and implemented as simply as possible. All decisions made by agents are either random or influenced by simple heuristics.

A.2: Parameters of the model

Parameter	Explanation	Default value(s)
nr_agents	Number of participants	9
rounds_max	Maximum number of simulated rounds	140
theshold_change continue	Maximum deviation in the group risk assessment for the discussion to be considered stagnant.	0.01
Risk model		
BN_nr_children	Number of children each parent-node has in the reference risk network	3
BN_factor_lmr	The factor for how much more likely the true state of information is to be low rather than medium or medium rather than high	10
Participants		
weight_h_low	Weight used for weighting the input from senders if the hierarchy is to be considered	0.25
weight_h_medium		0.50
weight_h_high		0.75
weight_e_expert	Weight used for weighting the input from the sender if transactive memory is present	1.0
weight_e_nonexpert		0.1
initial_trust	The model allows for varying the trust participants have in each other (note: for the experiments included in the paper, trust is kept constant at the default value)	1
initial_distr_info	Number of information to be distributed initially	45
Experimental conditions		
information_distribution	Determines whether all agents have the same probability of receiving a piece of information in the initial distribution	equal / unequal
weight_hierarchy	Participants consider the difference in the hierarchy when weighting sender input	yes/no
weight_trans_mem	Participants consider the difference in expertise when weighting sender input	yes/no
mode_limited_trans	Participants consider their prior beliefs when integrating new information	yes/no
decision_rule	Decision-making rule used by the leader	See chapter 3.3.2
Interaction_pattern	Interaction pattern used in the group	See chapter 3.3.4

A.3: Calibration of the Bayesian network

The Bayesian network that is used to model the risk needs to be calibrated to provide a plausible risk assessment from the information that constitutes the risk assessment task.

The overall risk assessment is derived from the state of the overall risk assessment node (see Figure 2 in the paper) based on the highest associated probability. We have consciously chosen a calibration for the Bayesian network that will always result in a “high” or “low” overall risk assessment. In the Bayesian network used in this study, from aggregation level to aggregation level the likelihood of “medium” states decreased strongly making it de facto a binary overall risk. Overall, this makes it possible to clearly identify the assessments made as a result of the risk workshop as correct or incorrect, streamlining the communication and discussion of the simulation results.

When Bayesian networks are used for real-world applications, the calibration of the network can be derived using machine learning algorithms on a real-world dataset. Lacking such real-world data, we make plausible assumptions on how the network should aggregate information provided at the information nodes. Technically, the aggregation nodes (the topic nodes, the domain nodes, and the overall assessment node) can be thought of as a lookup table that provides a probability value for the ‘low’, ‘medium,’ and ‘high’ state based on the probabilities of states that feed into the aggregation node.

Figure A.3.1. Bayesian network

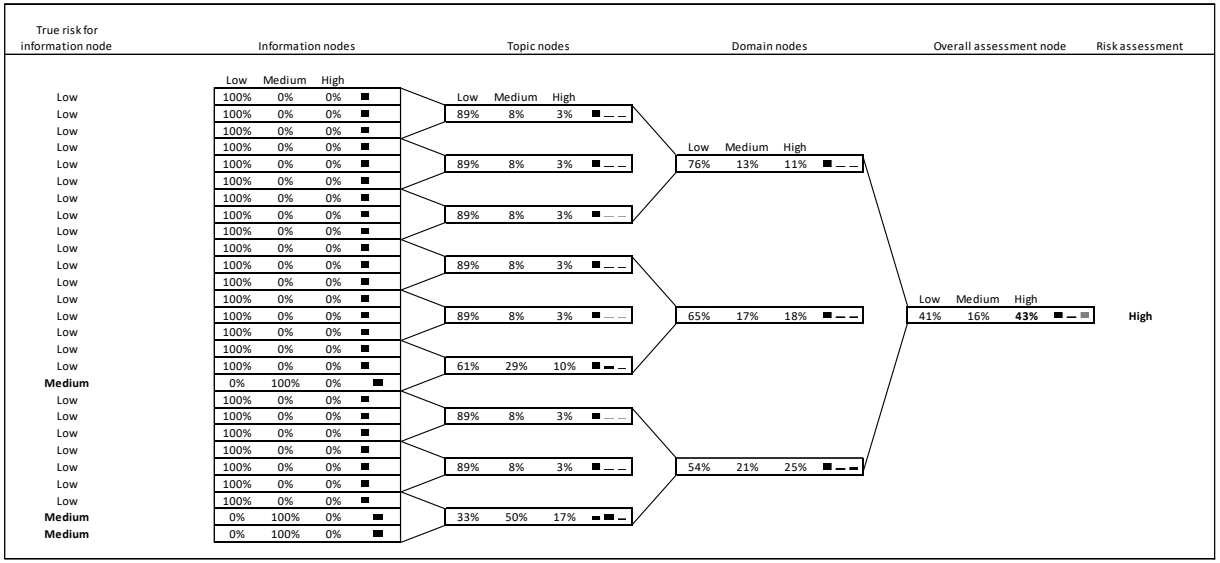

Figure A.3.1 provides an overview of how a benchmark assessment is calculated. It is also an example of how, even with some ‘medium’ information provided, the ‘low’ and ‘high’ states of the overall assessment nodes have higher probabilities than the ‘medium’ state. While for one topic node (the last one

at the bottom of Figure A.3.1), the ‘medium’ state has the highest probability, at each further level, more information is aggregated, and the ‘low’ and/or ‘high’ states become more dominant.

The aggregation of the nine values that feed into each aggregation node (three for each connected node, with one likelihood for the low state, one for the medium, and one for the high) is implemented in a two-step process:

- First, the input is aggregated into a “risk score” in the range 0 to 1. Here is an example from nodes in Figure A.3.1: This example is the case in which all three nodes to be aggregated carry a “low risk” information (e.g., the aggregation of the first three information nodes at the top of Figure A.3.1). A large probability in the ‘low’ states will result in a low risk score.

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} * \begin{bmatrix} 0.01 & 0.09 & 0.90 \end{bmatrix} * \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix} = [0.01]$$

- Second, this aggregate risk score (0.01) is now passed to a lookup table that translates each risk score into probability values for all three states. Figure A.3.2 shows how the aggregate risk score translates into state probabilities. The values were chosen to reach plausible aggregates for some plausible inputs. For example, an input of ‘medium’ on all input nodes will provide the highest value for ‘medium’ at the aggregate node. However, it was important to calibrate the network so that even a single high probability for a ‘high’ state translates into a high probability of the ‘high’ state of the aggregate node, as we assume that most information nodes are always in the ‘low’ state, and the task of the risk workshop is to correctly handle a low number of ‘high risk’ or ‘medium risk’ information.

The configuration of the aggregate node that is derived this way for the topic nodes is equally applied for all aggregate nodes.

Figure A.3.2. Graph on the translation of the aggregate risk score into state probabilities

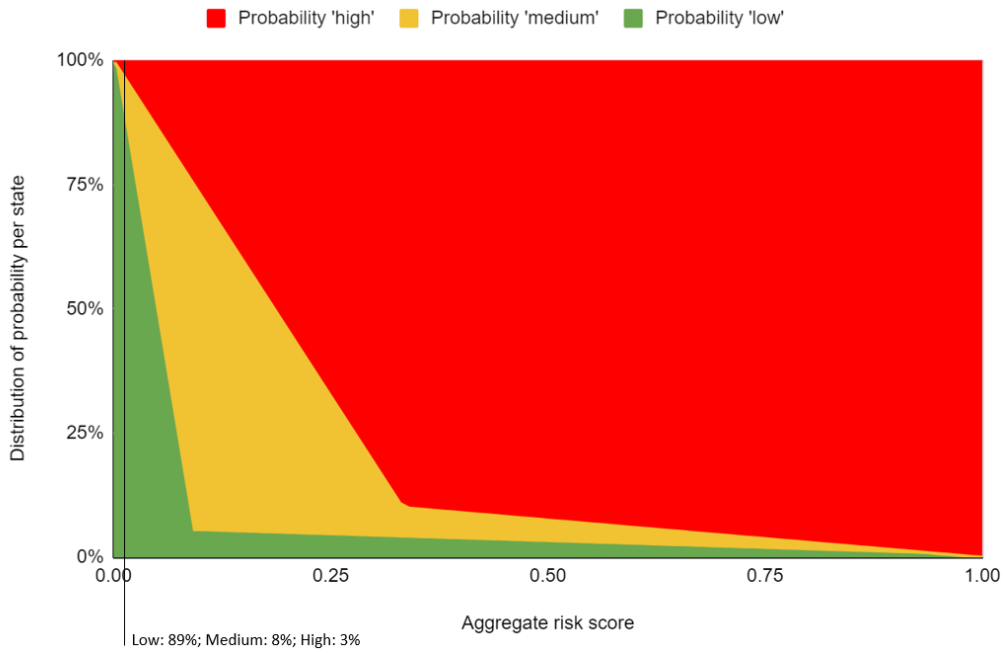

A.4: Information processing in the Bayesian network

The mental model of the participants, implemented as a Bayesian network, allows deriving an overall risk assessment from a set of beliefs about information relevant to the risk. If all information is provided without any noise (as in the end of a discussion with the ideal speech situation), the assessment of a participant is identical to the benchmark assessment. The participant is right for the right reasons (Figure A.4.1).

Figure A.4.2. Correct assessment with all information

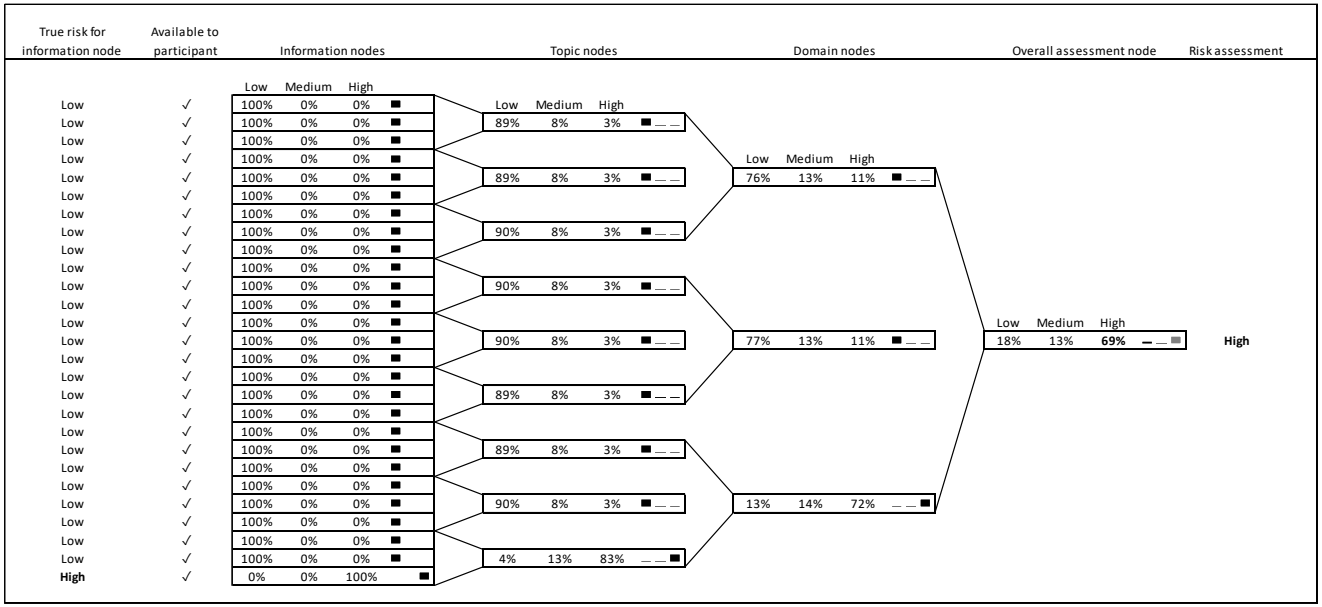

The participant can also reach a correct assessment with much more limited information (Figure A.4.2). However, in this case the assessment is correct only due to uncertainty about information: The ‘high’ information is not available to the participant – the participant is right for the wrong reason.

Figure A.4.2. Correct assessment with few information

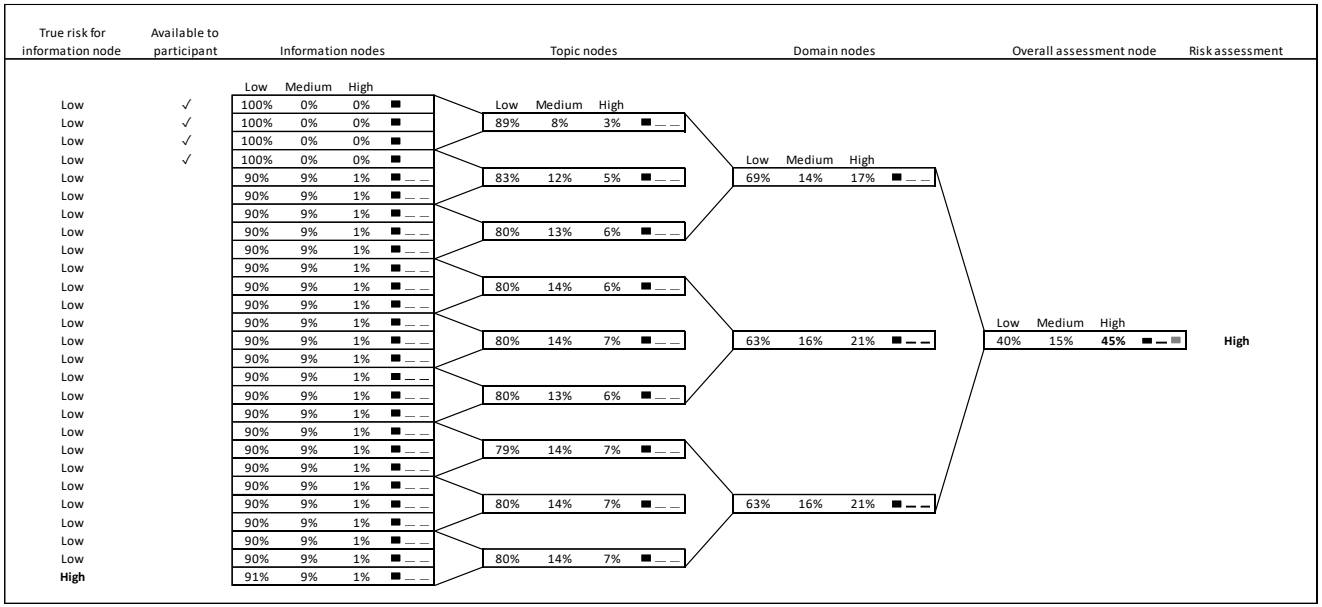

As the participant gains access to more information (again, assuming a unhindered discussion as in the ideal speech situation), the overall assessment of the participant switches to ‘low’ (Figure A.4.3). By learning more (correct) information, the assessment becomes wrong.

Figure A.4.3. Wrong assessment with more information

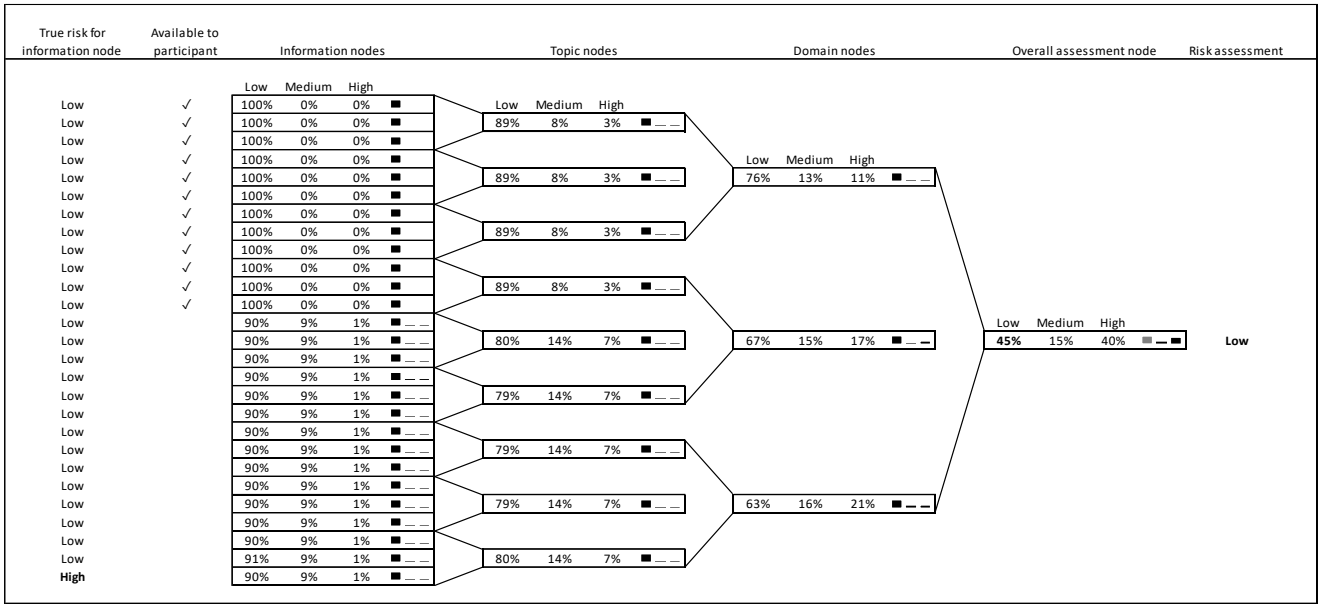

References

Habermas, J. (1982), *Theorie des kommunikativen Handelns*, 2nd ed., Suhrkamp, Frankfurt am Main.

Kabir, G., Tesfamariam, S., Francisque, A. and Sadiq, R. (2015), "Evaluating risk of water mains failure using a Bayesian belief network model", *European Journal of Operational Research*, Vol. 240 No. 1, pp. 220–234.

Pearl, J. (2008), *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Rev. 2. print., 12. [Dr.]., Kaufmann, San Francisco, Calif.

Wegner, D.M. (1987), "Transactive memory: A contemporary analysis of the group mind", in Mullen, B. and Goethals, G.R. (Eds.), *Theories of Group Behavior*, Springer New York, New York, NY, pp. 185–208.