The potential and risks of data-driven risk-based inspection for regulatory oversight

How can a data-driven risk-based inspection approach be designed and implemented at the ANVS to improve efficiency and accuracy under uncertainty?

by

Annabel Verspeek

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Wednesday February 11, 2026 at 15:00.

Student number:	5085241	
Project duration:	September 1, 2025 – February 11, 2026	
Thesis committee:	Prof. dr. M. E. Warnier,	TU Delft, Chair, First supervisor
	Dr. H. G. van der Voort,	TU Delft, Second supervisor
	H. van Haaster,	CGI The Netherlands, External supervisor
	L. K. Janssen-Pinkse	ANVS, External supervisor

GitHub:
https://github.com/annabelverspeek1/abm-data-driven-risk-based-inspection

An electronic version of this thesis is available at
http://repository.tudelft.nl/.

Preface

This thesis was written as part of my Master's programme in Engineering & Policy Analysis at the Technical University of Delft. The research was executed in collaboration with CGI Netherlands and the Autoriteit Nucleaire Veiligheid en Stralingsbescherming (ANVS) and focuses on the design and implementation of a data-driven risk-based approach under uncertainty.

I want to thank my chair and first supervisor, Prof. Dr. M.E. (Martijn) Warnier, for our weekly meetings filled with knowledge and targeted advice. Every time I came with a list of questions and topics to discuss, I left with very valuable new insights and confidence for the next part. These meetings gave me structure throughout the whole research. In addition, I want to thank my second supervisor, Dr. H.G. (Haiko) van der Voort. Although we did not meet weekly, the insights from his different point of view were extremely important. It kept me focused on the target, with strong definitions and a focus on the value of the case system.

To everyone in the CGI iAMLAB, I would like to thank you for your great interest and valuable advice, filled with technical insights and diverse perspectives. In particular, I would like to thank Henk for being my supervisor on behalf of the IAMLAB at CGI. A special thank you to Linda from the ANVS for giving me the opportunity to do this research and for working with her and her team to gain insight into the valuable work you do as inspectors.

Last but not least, I would like to thank my family and friends for their support throughout the entire Master's programme, and especially during the final months of this thesis. Thank you for studying together, reading my drafts, discussing insights, and sharing good coffee breaks with much-needed distractions.

The past few months have brought me academic growth and confidence. It has been a great and challenging experience. I hope you enjoy reading this thesis!

Annabel Verspeek
Delft, January 2026

Executive summary

In high-impact sectors such as nuclear safety and radiation protection, inspections play a crucial role in regulatory oversight. In the Netherlands, the Nuclear Safety and Radiation Protection Authority (ANVS) is responsible for planning and carrying out these inspections. Due to limited inspection capacity and a large number of licence holders, the ANVS applies a risk-based inspection (RBI) approach, in which licence holders with higher expected risks are prioritised. Inspection planning is currently informed by risk profiles, contextual factors, and inspectors' professional judgement. While both experience and quantitative indicators are valuable, current risk profiles are partly based on unstructured information and a limited set of measurable indicators, which may introduce blind spots, especially under changing external conditions.

Data-driven approaches have the potential to make risk-based inspections more efficient and structured, provided that the information used for inspections is accurate. Although risk-based inspection is used across many sectors, the design is highly context-specific. Each regulatory domain faces different risk dynamics and therefore requires its own set of indicators and inspection strategies. In practice, however, many of these important indicators are not available in data, instead, they are hidden in inspectors' experience. Therefore, the goal of this research is to investigate how a data-driven risk-based inspection approach can be designed for the ANVS context, where inspections have limited inspection capacity, and where uncertainty in prioritisation can have high-impact consequences. The focus is on how inspection decisions based on this uncertain data affect inspection outcomes over time, and how data and professional judgement can be combined to improve both efficiency and accuracy.

Therefore, this study aims to answer the following main research question:

How can a data-driven risk-based inspection approach be designed and implemented at the ANVS to improve efficiency and accuracy under uncertainty?

The main findings show that inspectors' tacit knowledge can be externalised into a structured set of risk factors that represent the risk of licence holders. These factors combine measurable characteristics, external developments, and behavioural dynamics, allowing inspection planning to focus on individual licence holders rather than broad groups or entire branches. This focus supports more targeted resource allocation to optimally use the inspection capacity. When these factors are incorporated into a simulation model, licence holders' risk evolves over time in response to inspection history and external conditions. Uncertainty in risk prediction leads to variation in inspection rankings. Under favourable external conditions, this variation has a limited effect on inspection outcomes, while under less stable conditions, it increases the likelihood of missing high-risk inspections. Worst-case analysis shows that, despite this uncertainty, risk levels remain bounded and do not result in extreme outcomes.

Further analysis shows that uncertainty mainly affects inspection accuracy through temporary changes in prioritisation rather than through structural blind spots. Periods of unfavourable external conditions are characterised by external developments that increase risk levels across many licence holders simultaneously, such as an increase in unemployment. In such periods, the number of high-risk cases increases, and prioritisation becomes more sensitive to uncertainty, even though predicted risk scores at the individual level remain relatively distinct. This indicates that differentiation between licence holders is still present, but that higher overall risk levels increase the likelihood of missed high-risk cases.

To better understand how the likelihood of missed high-risk inspections can be reduced under such conditions, the study examined how different policy implementations in inspection prioritisation affect inspection outcomes. The results show that no single change is sufficient to fully reduce missed high-risk inspections. Increasing inspection capacity has a strong short-term effect, but requires additional resources and does not improve efficiency. Improvements in data quality have the most potential. It reduces sensitivity to external changes and strengthens inspection performance over time. Partly random inspections broaden inspection coverage and may indirectly support data collection, although their direct effect remains limited within the current model.

This thesis concludes that a data-driven risk-based inspection approach can support inspection planning at the ANVS when uncertainty in data and judgment is explicitly considered. Data-driven RBI can be introduced within existing inspection practices and used as structured input in annual planning without changing the role of inspectors. The value lies in making risk considerations more explicit and comparable across licence holders, while remaining sensitive to the dynamic and high-stakes regulatory context and implementing policies strategically. While the application of data-driven RBI remains context specific, the method used in this thesis shows how inspection experience and data can be combined and annually updated over time. This allows inspection planning to improve accuracy while including professional judgement in the final decisions of the inspection planning.

Contents

List of Figures

List of Tables

Abbreviations

SQ = Sub-Question

RBI = Risk-Based Inspection

ABM = Agent-Based Model

ANVS = Autoriteit Nucleaire Veiligheid en Stralingsbescherming (Dutch Authority for Nuclear Safety and Radiation Protection)

KPI = Key Performance Indicator

1 Introduction

In high-impact sectors such as nuclear energy and radiation protection, inspection is an important aspect of regulatory oversight. The Nuclear Safety and Radiation Protection Authority (ANVS), in the Netherlands, is responsible for maintaining safety standards in organisations that handle radioactive materials (ANVS, 2022). Currently, ANVS already applies risk-based inspection (RBI) to its annual inspection planning. Rather than distributing inspections equally among all licence holders, RBI prioritises the highest-risk cases. Risk is defined as the likelihood that an incident occurs at a licence holder, combined with the severity of the consequences. However, given the limited inspection capacity and the large number of licence holders, there is increasing recognition that this approach can be strengthened by improving the allocation of inspection resources. This prioritisation supports a more efficient and targeted allocation of inspection resources, while maintaining the focus on licence holders with a high risk.

In the current RBI implementation at the ANVS, inspection planning is informed by the risk profile of a licence holder, as well as other internal and external factors and the inspectors' experience. This is shown in figure 1.1 below. However, the information behind these risk profiles is limited and not fully structured. It is mainly hidden in the experience-based knowledge and a small number of measurable indicators, such as the number of radiation sources and their intensities (ANVS, 2022). Although inspector experience and these measurable factors are valuable indicators, relying just on this and experience can introduce cognitive biases, leading inspectors to follow familiar patterns, overlook critical signals, or be disproportionately influenced by recent events (Coffeng, 2022). At the same time, limiting RBI to a small set of quantitative factors risks ignoring other relevant aspects, such as organisational culture, historical incidents, or recent changes in staffing and operations.

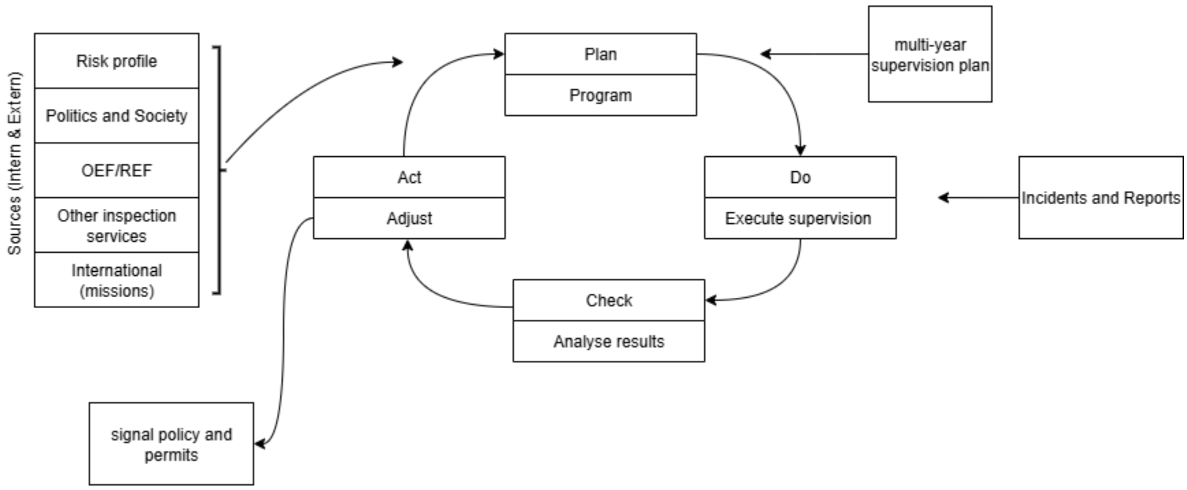

Figure 1.1: Supervision strategy ANVS (adapted from ANVS (2023)).

In the nuclear and radiation domain, where the consequences of failure are highly impactful (World Nuclear Association, 2024), there is a concern that RBI systems, both data-driven and experience-based, might overlook possible risks or strengthen blind spots due to biases. Data-driven risk-based inspection (RBI) relies on quantitative models to prioritise which organisations should be inspected first. However, these models often simplify reality and may overlook relevant factors (Black & Baldwin, 2010). This raises an important question: how accurate are RBI strategies when key input variables are missed, and when parts of inspectors' experience-based knowledge are difficult to formalise in data? How inspectors interpret the inspection context is mainly based on their experience. When these hidden factors are not represented, risk assessment may appear accurate, while missing critical potential hazards.

This is a challenge for ANVS, which aims to improve its RBI approach to increase efficiency. Therefore, this research assesses the risks arising from the use of data-driven RBI in practice. Rather than focusing on increasingly complex algorithms, the research explores a simple prediction approach, combined with simulation, to better understand the effects of data-driven inspection planning over time. By extracting measurable input factors that reflect inspectors' judgment, the research examines how these factors interact with different inspection dynamics and how risk levels and inspection outcomes develop over time. In this way, the study provides insight into how uncertainty in the underlying factors can influence long-term inspection planning.

The results will provide a grounded recommendation for designing an accurate inspection strategy in a high-stakes environment, such as radiation safety. The final focus is on identifying which factors matter most, how sensitive the outcomes are to uncertainty, and what policy implementations limit the risk of missing high-risk licence holders.

The structure of the thesis is as follows. Chapter 2 presents a literature review on risk-based inspection, data-driven methods in regulatory supervision, and the role of inspectors' experiential knowledge. Based on this review, the chapter identifies the main knowledge gaps and the main research question addressed in this research. Chapter 3 describes the research approach and explains how qualitative insights and quantitative analysis are combined to study data-driven RBI under uncertainty. Chapter 4 provides the system description and outlines the organisational context of inspection planning at ANVS, including the role of inspectors and the annual planning cycle.

Chapter 5 addresses the first sub-question (SQ1) and explores which risk factors inspectors consider important and how this knowledge can be structured for further analysis. Chapter 6 focuses on the second sub-question (SQ2) and examines how the identified risk factors can be translated into model inputs. Chapter 7 addresses the third sub-question (SQ3) by analysing how different inspection dynamics and assumptions influence inspection outcomes over time. Chapter 8 answers the fourth sub-question (SQ4) and investigates how uncertainty in the identified factors affects long-term risk assessment and inspection planning. Chapter 9 addresses the fifth sub-question (SQ5) by evaluating different policy options and their impact on inspection accuracy.

Chapter 10 presents the discussion of the results in relation to the literature and the practical implementation of data-driven RBI at the ANVS. Chapter 11 concludes the thesis and provides recommendations for inspection practice.

This thesis is relevant to the Master of Engineering and Policy Analysis (EPA) as it combines technical modelling with identifying policies in a high-risk regulatory environment. It shows how simulation and uncertainty analysis can support inspection planning in a responsible way, while recognising the value of inspector expertise. This closely aligns with the EPA's goal of designing innovative solutions for complex policy challenges.

2

Previous research on (data-driven) Risk-Based Inspections

Accurate inspection is essential for regulatory control in high-risk areas such as radiation protection. The ANVS aims to improve its RBI approach, which currently focuses primarily on inspectors' experience, a small set of quantitative factors, and risk alerts. However, this can lead to cognitive biases (Coffeng, 2022) and the overlooking or ineffective consideration of hard-to-quantify factors, such as company culture and recent operational changes. This chapter reviews previous research on RBI, with a focus on integrating multiple knowledge sources, such as advanced analytics and tacit expertise, into data-driven RBI approaches.

2.1. Literature review method

This literature review explores how risk-based inspection (RBI) can be improved by combining different sources of knowledge within a data-driven inspection approach. The review was executed in several steps. First, literature on tacit knowledge was reviewed to clarify how inspectors' experience and professional judgement influence inspection decision-making.

Second, a structured literature search was done to identify practical research on RBI frameworks and inspection planning. The search focused on data-driven RBI approaches, inspection methodologies and the role of expert judgement, using the following query in Scopus: TITLE-ABS-KEY ((”risk-based inspection” AND (”framework” OR ”methodology” OR ”model” OR ”inspection planning” OR ”regulation” OR ”compliance” OR ”safety”) OR (”risk-based inspection” AND (”data-driven” OR ”machine learning” OR ”artificial intelligence” OR ”predictive modelling” OR ”big data” OR ”temporal data” OR ”time series” OR ”dynamic risk assessment” OR ”real-time data”)) OR (”risk-based inspection” AND (”tacit knowledge” OR ”expert judgement” OR ”inspector experience” OR ”expert knowledge”))). The following papers were found using forward and backwards snowballing searches on the initial results. see figure 2.1 below.

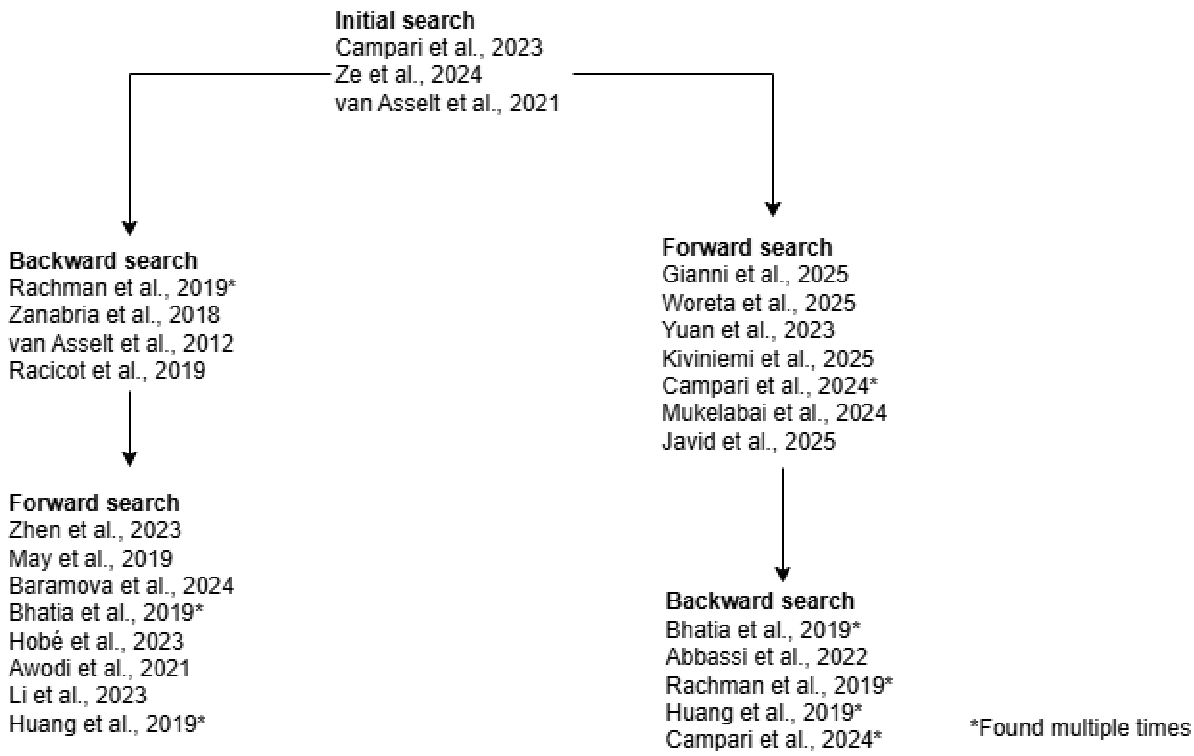

Figure 2.1: Literature method structure

Third, during this review, uncertainty was identified as a key risk in data-driven RBI, due to incomplete data and subjective judgment. As this is a relevant concern for ANVS, additional literature on uncertainty in RBI was included.

Based on this process, the literature was structured around three themes: (1) tacit knowledge in inspections, (2) data-driven risk-based inspection planning, and (3) uncertainty in data-driven RBI. These themes form the basis for identifying knowledge gaps and evaluating how RBI could be improved in the context of radiation safety supervision.

2.2. Tacit knowledge in inspectors' experience

First explained by Polanyi (1966), tacit knowledge refers to the skills, intuitions and insights that people possess but cannot easily communicate. In the context of inspections, this could include the ability to recognise subtle warning signs, capture how an organisation functions, or identify potential hazards before they cause serious damage. Such knowledge is gathered over many years through work and learning from others (Collins, 2010). The problem for inspectors is that tacit knowledge is difficult to document, share or use within information systems. This can result in irregularities in inspectors' decision-making or even the loss of valuable knowledge when experienced inspectors retire (Collins, 2010). Therefore, there is strong potential for ANVS to make use of inspectors' tacit expertise by combining it with more structured methods for potential data-driven inspection prioritisation. However, this is challenging, as these forms of knowledge differ in nature and are not easy to combine.

The spiral of knowledge creation proposed by Nonaka and Takeuchi (1995) provides a valuable framework to address this challenge. Their model describes how tacit knowledge can be externalised into explicit knowledge, combined with other information, and internalised again through the use of the shared information. In the context of inspections, this means that inspectors' experiences and knowledge can be systematically captured, formalised and used as input for inspection planning. Individual experience can be translated into shared organisational knowledge that supports more consistent decision-making if this spiral is applied in an iterative manner.

Currently, ANVS uses RBI to determine which radiation licence holders will be inspected each year. This system relies heavily on the experience of inspectors and a limited number of simple indicators. While expert judgement is useful, if it is the main source, it can also create bias and inconsistency (Coffeng, 2022). For instance, a licence holder with an extensive history of compliance may receive fewer inspections, even when recent staff or organisational changes increase the risk. Focusing on easily quantifiable factors, such as the number of radiation sources, may mean that other important factors, such as safety culture, staffing or previous incidents, are overlooked. Another issue is the quality of the data. Records are sometimes incomplete, outdated, or not standardised. This makes it difficult to plan inspections fairly and transparently.

The same problems affect other sectors. In the field of nuclear energy, for example, Awodi et al. (2021) demonstrate that, in the absence of accident data, regulators rely on expert judgement to identify and prioritise potential risks. While this provides valuable insights, it can also introduce subjectivity and inconsistency. The authors argue that expert knowledge should be applied more systematically. This situation is similar to that faced by the ANVS. In the mining industry, Huang et al. (2019) demonstrate that even when experts use structured methods such as the Analytic Hierarchy Process (AHP), their findings must be verified using data to mitigate bias. These studies demonstrate the importance of tacit knowledge, however, current ANVS RBI methods rely too heavily on it. They are neither standardised enough nor making full use of the available data.

For the ANVS, tacit knowledge is the main value and should form the basis for any further development of RBI. This knowledge is essential for judging risks in this specialised sector, however, it is subject to cognitive bias. More structured methods can help reduce these biases, but the externalisation of tacit knowledge into formal models is complex, requires ongoing iteration, and introduces additional biases through the use of data in these models. When combined with a data-driven approach, these biases have the potential to partly cancel each other out, to reduce inaccuracy.

2.3. Towards data-driven, risk-based inspection planning

Research from the food safety, chemical and energy industries demonstrates how RBI can be improved. A key lesson is that regulators should consider a broader range of indicators and apply them consistently and transparently. van Asselt et al. (2021) point out that there is no single RBI method in food safety and that the prioritisation of inspections largely depends on how consistently risk factors are selected and used across companies. This demonstrates that the value of RBI lies not only in identifying the right factors, but also in applying them in a clear and systematic manner van Asselt et al. (2021). Such an approach reduces bias, makes inspection planning more predictable and increases trust among stakeholders since both regulators and licence holders can understand how and why certain inspection decisions are made.

One of the strongest predictors of future non-compliance is compliance history. Studies in food inspection ((Zanabria et al., 2018); (May & Nikiforova, 2019)) demonstrate that the systematic inclusion of past violations and recalls leads to a more efficient allocation of inspection resources. Kiviniemi et al. (2025) also highlights the importance of treating repeat offenders differently, businesses with repeated violations should immediately be categorised as high risk, resulting in more frequent inspections and stricter enforcement. The size and scope of an organisation also matter. While larger businesses have the resources and expertise to comply with regulations and manage complex and potentially hazardous operations, smaller businesses often lack these capabilities. Hobé et al. (2023) demonstrate how consumption data, pollution levels and production volumes can be combined to identify products causing the greatest risk, enabling inspectors to prioritise their resources effectively. Similarly, van Asselt et al. (2021) highlight that company size and production scale are important factors in food safety inspections as they affect the probability of non-compliance and the potential impact on public health. The same logic applies to the radiation sector, where small clinics may struggle with resources, while large hospitals and industrial users pose a higher inherent risk due to their size and complexity.

In the food sector, Racicot et al. (2019) demonstrate that effective hygiene programmes and well-trained staff reduce the likelihood of non-compliance. Van van Asselt et al. (2021) also highlight that trained personnel and formal management systems are key indicators of compliance. They ensure that risk control is part of daily practice and not only implemented when inspectors visit. The same applies to radiation oversight. The presence of a certified radiation protection officer, written safety procedures and internal audits can all improve compliance. Li et al. (2023) supports this by demonstrating that the presence of certified safety officers is closely associated with improved compliance outcomes. For the ANVS, this suggests that certified radiation protection officers and approved safety systems could be valuable indicators in RBI models.

Although safety culture and organisational behaviour are harder to measure, they are still crucial. Bayramova et al. (2024) have proposed systematic methods for finding important safety indicators that include culture-related factors such as awareness and trust. Huang et al. (2019) developed a simple scoring tool for safety culture, which correlated well with actual compliance outcomes. Woreta et al. (2025) also found that stronger knowledge and attitudes among staff resulted in safer behaviour in water bottling factories. Together, these studies demonstrate that culture, communication and training can be captured in a structured manner and should not be overlooked in RBI.

Further improvements to RBI can be achieved through advanced methods and predictive analytics. Abbassi et al. (2022) provides a detailed overview of quantitative RBI and predictive maintenance techniques. They highlight that proactive, risk-focused inspections can significantly reduce the likelihood of severe accidents. In addition, Bhatia et al. (2019) argue that RBI should not assume that conditions remain constant between inspections, since risks can change quickly. They propose dynamic RBI, whereby inspection schedules adapt in real time based on monitoring data. While such real-time monitoring may not yet be feasible within the ANVS's regulatory context, the insight that inspection planning should incorporate up-to-date information and be responsive to emerging situations is valuable. Yuan et al. (2023) adds to this logic by applying Bayesian methods, a statistical approach that continuously updates risk estimates as new information becomes available, such as sensor readings, test results or incident reports.

Machine learning (ML) is another important tool. Instead of focusing only on predefined risk factors, it can detect patterns within large, complex datasets that are not immediately obvious to inspectors. Campari et al. (2024) demonstrate how ML can improve risk predictions in technical systems, while Mukelabai and Barbour (2024) show how effective it is in situations where traditional models are limited by insufficient data. Similarly, Zhen et al. (2023) demonstrate how ML can identify the most significant early warning signals and areas of weakness, assisting regulators in prioritising their oversight resources.

Rachman and Ratnayake (2019) builds on this by demonstrating how ML can minimise the subjectivity and variability inherent in RBI screening assessments. Traditional inspections rely mostly on qualitative judgement, resulting in inconsistent outcomes. However, by transferring knowledge from past detailed assessments into ML models, Rachman and Ratnayake (2019) shows that screening can become more accurate and less biased. A case study in the oil and gas sector showed that ML models such as gradient boosting and random forests could identify high-risk equipment with over 90% accuracy, outperforming human inspectors and reducing variability between evaluators.

For the ANVS, this means that factors identified in literature on other sectors can be used as a starting point for identifying relevant risk factors. However, the externalisation of tacit knowledge remains important. Combining these initial insights from the literature with the factors inspectors consider important, necessary to make RBI work for the specific context of the ANVS. Afterwards, there is strong potential to use more advanced methods, such as machine learning, to support the prediction or risk scores for licence holders.

2.4. Uncertainty in Data-Driven RBI

One of the main challenges in data-driven RBI is managing uncertainty, particularly when models rely on incomplete data, subjective judgment or prediction of rare events. Several studies show that if uncertainty is not considered, RBI systems risk becoming overly confident, which can lead to biased outcomes or inefficient resource allocation.

Pasman and Rogers (2020) argue that expert judgement is often treated as factual, despite the presence of uncertainty. They recommend the use of formal methods such as evidence theory or fuzzy logic to make this uncertainty visible and prevent over-reliance on individual estimates. In addition, Yazdi et al. (2019) proposes the Z-number to record not only what experts think, but also their level of confidence. This distinction is important when multiple experts disagree or when the available knowledge is limited, which is often the case in regulatory oversight.

From a data perspective, Dyer et al. (2022) shows how missing values and limited available data can be addressed using model and missing data methods. Their work highlights the importance of validating different model types to avoid over-fitting and detect early risk signals. For inspectors, this demonstrates that risk models do not require perfect data to be effective, but they must account for missing data and uncertainty.

Another common challenge is dealing with rare events. In RBI, for example, most licence holders are compliant most of the time, meaning that serious violations are rare. Papadopoulos and Benardos (2021) shows that, without corrections, machine learning can ignore these cases, resulting in inaccurate predictions. They use oversampling to make the model more sensitive to rare but high-risk outcomes. This kind of adjustment is essential when inspections are intended to prevent the worst cases, rather than just follow the average.

While these studies provide robust methods for dealing with uncertainty, they are primarily static. They improve the quality of model inputs or predictions, but do not test what happens after those predictions are made. This is where agent-based modelling (ABM) becomes useful (Macal & North, 2010). ABM can simulate the impact of different types of uncertainty, such as incomplete data, biased factor weights, and shifting risk profiles, over time. It can be used to answer questions such as: What would happen if we prioritised certain factors over others? What if uncertainty causes certain (groups of) licence holders to be overlooked?

In the ANVS context, where inspections must be planned quickly and with limited resources, combining structured uncertainty modelling with ABM offers a practical way to assess the accuracy of the data-driven RBI approach. It also helps to explore the impact of policy choices. In short, ABM does not solve uncertainty, but makes its consequences visible (Bankes, 1993; Edmonds & Meyer, 2017).

2.5. Knowledge gap: Translating inspection knowledge into risk-based inspection

Although RBI frameworks are widely used, with the use of data, there is no consistent approach for translating inspection knowledge into inspection priorities. Within the ANVS, RBI relies heavily on inspectors' tacit knowledge. This knowledge is essential for judging risks in a specialised sector, however, it is difficult to externalise. Following the spiral of knowledge creation (Nonaka & Takeuchi, 1995), translating tacit knowledge into explicit indicators to use quantitative models requires ongoing iteration, which makes it difficult to systematically include it without introducing additional bias, due to the loss of relevant context factors.

At the same time, the literature already provides relevant potential risk factors, such as compliance history, organisational culture, and company size, that have been proven to be relevant in other sectors for RBI. These factors can serve as a useful starting point for the ANVS, however, some risk indicators are context-specific. Without externalising inspectors' knowledge, important ANVS risk indicators may be missed, or existing indicators could be interpreted in a way that does not apply to the ANVS.

This challenge is relevant in a high-impact sector such as radiation oversight, where tolerance for hazards is very low, so inspection accuracy is important, while inspection capacity is limited. Both cognitive bias in expert judgement and bias introduced through data-driven methods affect inspection prioritisation. When combined, these biases may partly overlap and reduce each other. However, it's hard to fully cover all blind spots emerging from biases, therefore uncertainty remains inherent to RBI. It is unclear how this uncertainty influences inspection prioritisation, how it affects inspection accuracy over time, and how RBI should respond to changes in licence holders due to environmental dynamics.

Together, this knowledge gap focuses on how data-driven RBI can be designed so that tacit knowledge and measurable indicators are combined in a way that supports accurate inspection planning under uncertainty. This leads to the following research question:

How can a data-driven risk-based inspection approach be designed and implemented at the ANVS to improve efficiency and accuracy under uncertainty?

3

Research Approach

This chapter provides an overview of the methodological approach used to investigate how the ANVS can incorporate data-driven RBI in their dynamic regulatory environment. The approach builds on the knowledge gap identified in the literature review, which showed that current RBI methods rely mostly on tacit inspector knowledge, use a limited set of quantitative indicators, and do not systematically incorporate past or uncertain data when determining the risk level of a licence holder.

By combining qualitative insights with quantitative modelling, this research aims to identify important risk factors, including tacit insights from inspectors. These factors are then translated into measurable variables, so their influence can be tested in an agent-based model (ABM) and assessed for their potential contribution to efficient and accurate data-driven RBI. The ABM is used to explore whether data-driven RBI can support annual inspection planning for the ANVS and which policy choices help to guide this process, even when predictions hold uncertainty. With this, the ABM provides insight into how inspection priorities shift under different circumstances and levels of uncertainty, thereby answering the main question:

How can a data-driven risk-based inspection approach be designed and implemented at the ANVS to improve efficiency and accuracy under uncertainty?

This research consists of five interconnected components, each linked to one of the sub-questions. First, inspectors' experience and tacit knowledge are explored through interviews and literature research (SQ1). Second, the identified tacit risk factors are translated into measurable quantitative variables to be incorporated into the model (SQ2). Third, these variables are used in the ABM to analyse how risk levels and inspection outcomes evolve when the defined factors interact with the dynamic inspection environment and when different scenarios are applied (SQ3). Finally, the model is extended with uncertainty to assess how the uncertainty of the prediction of variable values affects the inspection strategy and the long-term outcomes (SQ4), and what policies limit the negative outcomes (SQ5).

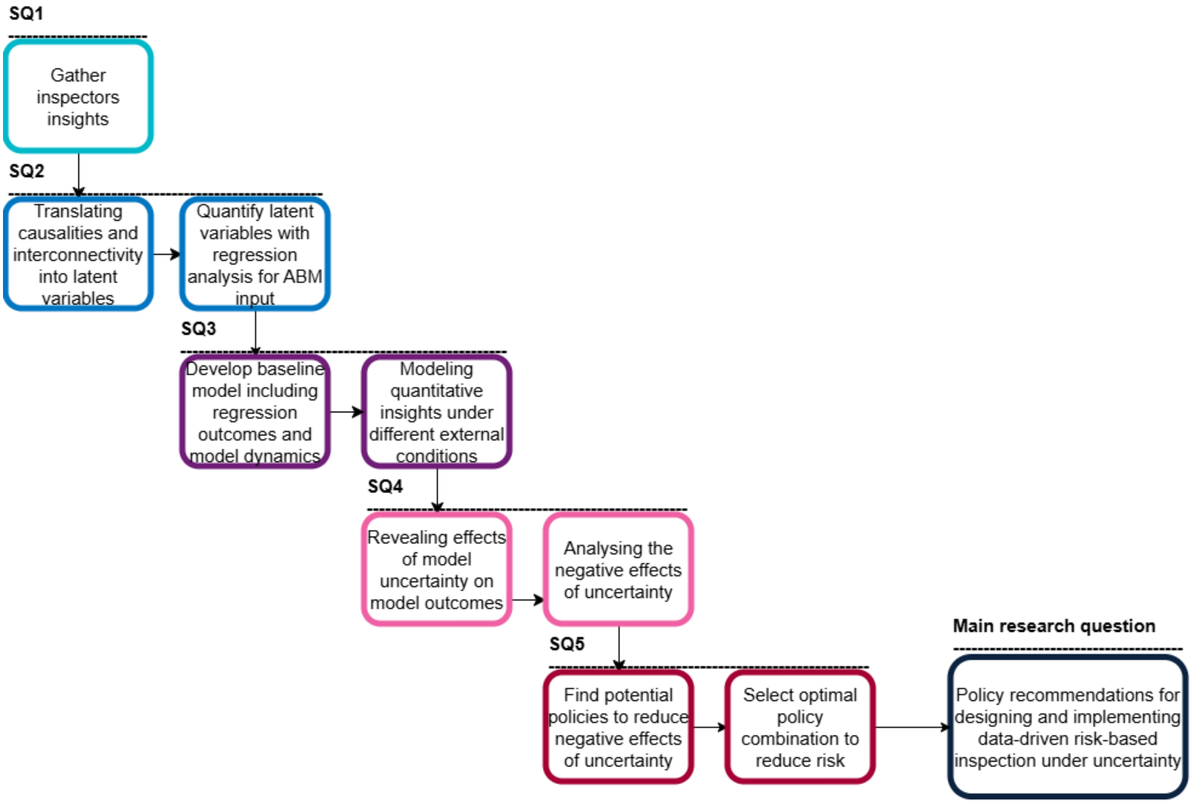

Figure 3.1: Method flow

3.1. Sub-questions

SQ1: Which insights are considered important in determining the risk level of a licence holder?

The first sub-question aims to identify which factors experienced ANVS inspectors consider important when assessing the risk level of a licence holder, to determine if they should be included in the annual inspection planning. Most RBI approaches rely on quantitative information, however, in the ANVS case, inspectors use a wide range of contextual, behavioural and organisational interpretations that are not systematically used because they are hard to measure. SQ1 focuses on uncovering this tacit knowledge and structuring it in a way that can be used in later modelling steps.

This approach, of uncovering inspectors' knowledge, can be explained by using the spiral of knowledge creation model described by Nonaka and Takeuchi (1995). This model shows how experience-based tacit knowledge can be made explicit and shared. In this thesis, inspectors' experiential insights are identified and described to make implicit inspection criteria explicit and suitable for further quantitative analysis.

To identify these insights, semi-structured interviews are conducted with inspectors with expertise in different branches and levels of experience. The semi-structured interview format allows inspectors to describe their reasoning in their own words, while still ensuring that all relevant themes are captured (Johannesson & Perjons, 1996). Inspectors are asked to explain which characteristics licence holders have that they would classify as higher or lower risk, which (external) developments they monitor closely in the radiation field, and which behavioural patterns they identify in practice.

The interview findings are supported by relevant literature on regulatory oversight and tacit knowledge. This helps to ensure that the identified factors not only reflect individual experiences but also consider broader insights from previous research on risk-based inspections. The combined results are used to create a structured overview of risk factors. This overview includes directly measurable characteristics, external developments, and more tacit behavioural or organisational aspects that inspectors consider important but are not directly measurable.

The outcome of SQ1 is a set of risk factors reflecting how inspectors interpret and understand risk in their daily work. These factors are then translated into latent variables and prepared for quantitative and predictive modelling with regression analysis in SQ2.

SQ2: How can important risk factors be translated into an agent-based model (ABM)?

This sub-question explores how the risk factors identified in SQ1 can be translated into quantitative components that can be used in the ABM model. Many of these factors are behavioural or intuitional and are therefore hard to measure directly. To use them in a data-driven RBI approach, they must first be organised and translated into a structured and measurable form.

The process begins by systematically revealing the interconnectivity of the factors derived from the interviews. Each factor is valuable. However, several describe similar underlying ideas or influence one another. To clarify these relationships, a causal relationship diagram will be developed. This diagram shows how individual factors influence each other and how they connect to the overall risk level of a licence holder. It identifies where causal chains begin and how certain factors reinforce each other due to loops. The purpose of this step is to clarify the internal logic of the risk factors, based on how inspectors describe and interpret risk in practice.

From this causal structure, underlying variables become visible, which are captured as latent variables. A latent variable represents a characteristic that cannot be observed directly, but which can be estimated based on multiple observable indicators (Spirtes, 2001). Combining related factors into latent variables makes complex or tacit concepts measurable and suitable for further quantitative analysis.

Once the latent variables are defined, they can be quantified using regression analysis. Regression analysis estimates the relationship between latent variables and the historical risk score of a licence holder, providing numerical coefficients that show the relative influence of each latent variable on predicted risk scores (Wooldridge, 2010). This quantitative step turns the latent variables into usable indicators of the ABM risk model.

The inspection dynamics are simulated as an agent-based model to analyse interactions between licence holders, inspections, and risk development over time. The model is implemented in Python using the Mesa framework (Masad & Kazil, 2015).

In summary, SQ2 translates qualitative inspectors' knowledge into a set of structured, measurable latent variables that can be used to model the development of risk. These latent variables are the basis of the ABM in SQ3, in which their dynamic behaviour and long-term impact on risk scores are explored.

SQ3: How does the model behave when the inspector's insights are included in different scenarios?

The third sub-question investigates how risk scores and inspection outcomes develop when the latent variables defined in SQ2 are implemented. This will be achieved by modelling it in an ABM and exposing it to changing conditions. This step is essential because RBI operates in a dynamic environment. Licence holders' latent variable values change over time due to changes in their environment, and with inspectors influencing their behaviour through their visits. SQ3 therefore examines whether a data-driven RBI approach, including the factors defined by experienced inspectors in SQ1, is a responsible way to determine the annual inspection planning, with the aim of making it more efficient. This will be analysed by how these dynamics prioritise inspections.

In the ABM, each licence holder is represented as an agent with a risk profile that changes over time. The risk score is initially determined by their characteristics included in the latent variables identified in SQ2, this score changes over time due to the changing values of external variables that are also included in the latent variables and several dynamic behavioural patterns. By incorporating these dynamics in the model, the ABM captures both relationships and behavioural trends.

To analyse the behaviour of the model, multiple simulation experiments will be done. Starting with the baseline model, which represents the inspection system under standard conditions. These conditions are the regression outcomes and a calibrated estimation of the pattern strengths. This baseline run is used to understand how risk develops over time in a scenario that is based on the historical values of the external factors, in addition, it provides a reference point for later comparisons. After developing the baseline, the ABM will run under several different scenarios. These scenarios introduce different external changes. By comparing the results across scenarios, the model shows how inspection priorities react to changes in the external environment.

In addition to scenarios, a sensitivity analysis will be conducted. This tests how variations in key assumptions, such as the strength of behavioural effects, affect the model results. Running all combinations of scenarios and sensitivity analysis in combination with the extensive uncertainty analysis would require significant computational power and a large number of results to analyse. Therefore, only the most relevant scenarios are selected for further analysis. This ensures that the results remain manageable while still including the most important sources of variation.

A key concept in SQ3 is inspection variance, which is used as an indicator to assess how the inspection priorities change over time. Inspection variance refers to the difference between risk scores, the number of inspections received, and the years since the last inspection. Analysing these indicators helps to determine whether the model produces stable and realistic inspection patterns.

SQ3 therefore uses the ABM model to analyse how the risk scores of licence holders evolve when inspectors' insights are incorporated and when the inspection environment changes. This provides the basis for SQ4, where uncertainty is added to evaluate whether these indicators show interesting behaviour when the input values of the latent variables are uncertain.

<u>SQ4: What patterns evolve when inspections are missed due to the estimation of risk scores under uncertainty?</u>

The fourth sub-question analyses how uncertainty in the estimation of risk affects the inspection strategy and the long-term development of risk levels in the ABM. The latent variables defined in SQ2 are based on limited information, which means that their estimated values contain uncertainty. As a result, the predicted risk scores used for inspection planning may differ from the actual risk of licence holders. SQ4 investigates how this uncertainty influences inspection outcomes and whether the inspection approach remains responsible when predictions are imperfect.

To analyse this, uncertainty is included in the model by allowing the calculated risk scores of licence holders to vary from the values predicted by the base regression model. The inspection selection, however, continues to rely on the predicted risk scores. This means that the model continues to inspect the licence holders with the highest predicted scores, even if their true risk is different due to uncertainty. This setup makes it possible to explore what happens when high-risk licence holders are not selected due to miscalculated risk scores caused by uncertainty.

The ABM will run multiple times with different uncertainty variations. In these runs, the key performance indicators are calculated to assess the consequences of uncertainty. First, the model measures the number of missed inspections per year. Second, it identifies which of these missed licence holders are at high risk. Third, it analyses the worst-case risk score patterns of individual licence holders to find the long-term patterns and whether they show extreme patterns occur due to missed inspections under uncertainty.

<u>SQ5: What policies can reduce the negative effects of data-driven RBI under uncertainty?</u>

In addition to the uncertainty analysis, several policy implementations are tested to assess whether the negative effects of uncertainty in data-driven RBI can be reduced. These policy implementations include small changes to the inspection selection strategy. Testing these alternatives within the ABM makes it possible to identify which policy implementations help to improve the inspection outcomes when making predictions under uncertainty.

Each policy adjustment is tested under different uncertainty ranges, in different environmental conditions. This makes it possible to assess how policies affect the negative consequences of uncertainty in data-driven RBI. The policies are evaluated using the same key performance indicators, such as the number of missed high-risk inspections and the development of worst-case risk paths of individual licence holders. This makes the effect of the different policies comparable.

By analysing the key performance indicators, SQ5 shows how sensitive the inspection strategy is to uncertainty in the risk estimates. The focus is on relative differences between policies, rather than on identifying a single optimal solution. The results show which policy implementations lead to more stable inspection outcomes and which help reduce missed high-risk inspections under uncertainty. For the ANVS, it helps identify which policy implementations keep the RBI approach efficient while still supporting an accurate annual planning process when information is incomplete, and what this means for inspection planning at the ANVS.

4

System Description

This chapter describes the organisational and regulatory context in which the ANVS operates, with a specific focus on how inspection planning decisions are made in practice. Understanding this context is essential for interpreting how data-driven RBI can be implemented and what this implies for organisational and inspector adaptation.

The centre of inspection planning at the ANVS is the annual planning sessions. These are meetings in which inspectors together review inspection experiences, share observations from previous years, and discuss risks and concerns across licence holders. During these sessions, inspectors critically reflect on past inspections and combine their professional judgement to determine inspection priorities for the coming year. This collective discussion plays a central role in translating individual inspection experiences into shared team priorities.

These annual planning sessions can be seen as an annual process of knowledge creation within the ANVS (Nonaka & Takeuchi, 1995). Tacit knowledge gained during inspections is shared and discussed among inspections throughout the year. This knowledge is made explicit through the planning process of selection and prioritisation of inspections. The outcome of this meeting is the annual inspection plan, which is then applied in practice, where new inspection experiences are generated and fed back into the next planning cycle. In this way, inspection planning functions as a continuous learning spiral rather than a single decision moment.

By describing the legal framework, organisational structure, inspection strategy, and the current role of RBI, this chapter shows how this process of knowledge creation can be supported by data-driven RBI. In this thesis, data-driven RBI is positioned as a tool that supports and structures the annual planning process by improving how available information is combined, while remaining part of the existing spiral of knowledge creation. In this way, data-driven prioritisation and the knowledge gained through inspections strengthen both inspection planning and the learning that takes place throughout the year.

4.1. Organisational structure

The ANVS is the national authority responsible for nuclear safety and radiation protection in the Netherlands. It operates as an independent body under the political responsibility of the state secretary for infrastructure and water management (IenW) (ANVS, 2022). This structure ensures independence, which is essential, as the ANVS must be able to make regulatory and supervisory decisions based on safety considerations, scientific knowledge, and legal requirements.

The ANVS is responsible for the execution of the following tasks:

1. <u>Regulation and licensing</u>: The ANVS develops and maintains rules for nuclear safety and radiation protection in the Netherlands. It assesses license applications, evaluates proposed changes to installations or activities, and sets conditions to maintain compliance with national and international standards. (ANVS, 2022)
2. <u>Supervision and enforcement</u>: The ANVS inspects many different organisations, including nu-

clear installations, research centres, hospitals, industrial facilities, laboratories, and smaller users of radiation sources. Inspectors check compliance with licence requirements, internal safety systems, review documentation, and intervene when necessary. Possible by resulting in interventions, such as written warnings, improvement measures, sanctions or escalation in cases of significant risk. (ANVS, 2022)

3. Advice, communication and crisis preparedness: The ANVS advises the state secretary on policy, participates in national and international expert groups and openly communicates about risks and incidents. The ANVS plays a crucial role in nuclear and radiation emergency responses by collaborating with safety regions, ministries, and international partners to ensure a structured response. (ANVS, 2022)

As this research focuses on the inspection activities of the ANVS, the relevant regulations and organisational processes related to supervision are discussed to understand how data-driven RBI can be implemented in practice. Within the ANVS, there is a clear difference between the nuclear and radiation supervision domains (ANVS, 2025b). Nuclear supervision involves only a small number of licence holders, who are highly detailed and continuous. Given the potentially serious consequences of a nuclear incident, inspections in this domain are mainly based on technical aspects of the installation.

Radiation supervision, on the other hand, covers a much larger and more diverse group of licence holders. Although the consequences of the radiation-related incident are smaller compared to nuclear installations, they remain significant. The large number of radiation licence holders makes it much more challenging to identify high-risk licence holders, and the limited available capacity creates a higher need for data-driven RBI.

4.2. ANVS' inspection strategy without data-driven RBI

The organisation of inspections and annual planning follows a structured process. Each year, inspectors participate in an annual planning session, in which inspectors can share their observations, discuss concerns, and suggest areas that require additional attention. This session is the main moment at which inspection experiences are collectively reviewed and translated into inspection priorities for the next year. This way shared expertise is created in line with the spiral of knowledge creation (Nonaka & Takeuchi, 1995). This corresponds to the socialisation phase shown in figure 4.1.

Based on the outcomes of these discussions, the annual inspection planning is developed. In the current situation, when a specific sector or branch is selected for inspection, a project plan is prepared. This plan outlines the main inspection objectives and specifies the topics or themes to be reviewed. Since it is often not feasible to inspect all possible aspects, project leaders determine the key subject or focus areas for that year's inspection. In figure 4.1, this step corresponds to the externalisation phase, where shared experiences are translated into explicit inspection objectives.

Based on this project plan, the project-specific inspection plans are created, describing what will be inspected. Inspectors conduct the site visits and document their findings in an inspection report. Each inspection question or criterion is evaluated and assigned a result, ranging from violation to good. Together, the development of inspection plans and the execution of inspections align with the combination phase in Figure 4.1, where explicit information is structured and applied within the inspection process.

The overall inspection outcome of a licence holder is determined using the inspection matrix, which combines two dimensions: (1) the degree of cooperation of the organisation; (2) the seriousness of the identified non-compliance. This outcome provides structured feedback on inspection results and supports follow-up actions if necessary. These outcomes create regulatory pressure that stimulates or forces licence holders to take actions to improve their risk management.

Inspection outcomes are not only used for enforcement or follow-up actions, but also for learning within the organisation. Inspectors discuss recurring findings, changes in compliance, and levels of cooperation, and use these insights in the next planning cycle. This way, inspection outcomes help to update inspection priorities and focus areas over time. This learning process represents the internalisation phase shown in Figure 4.1, where inspection results are gathered into inspectors' tacit knowledge.

This inspection strategy ensures that inspection planning and inspection outcomes are closely connected. Information and experience gained during inspections are reflected during the annual planning sessions and used to define priorities. As shown in the knowledge creation spiral in Figure 4.1, inspection planning functions as an iterative process in which experience from practice is translated into shared priorities and applied for defining the annual planning.

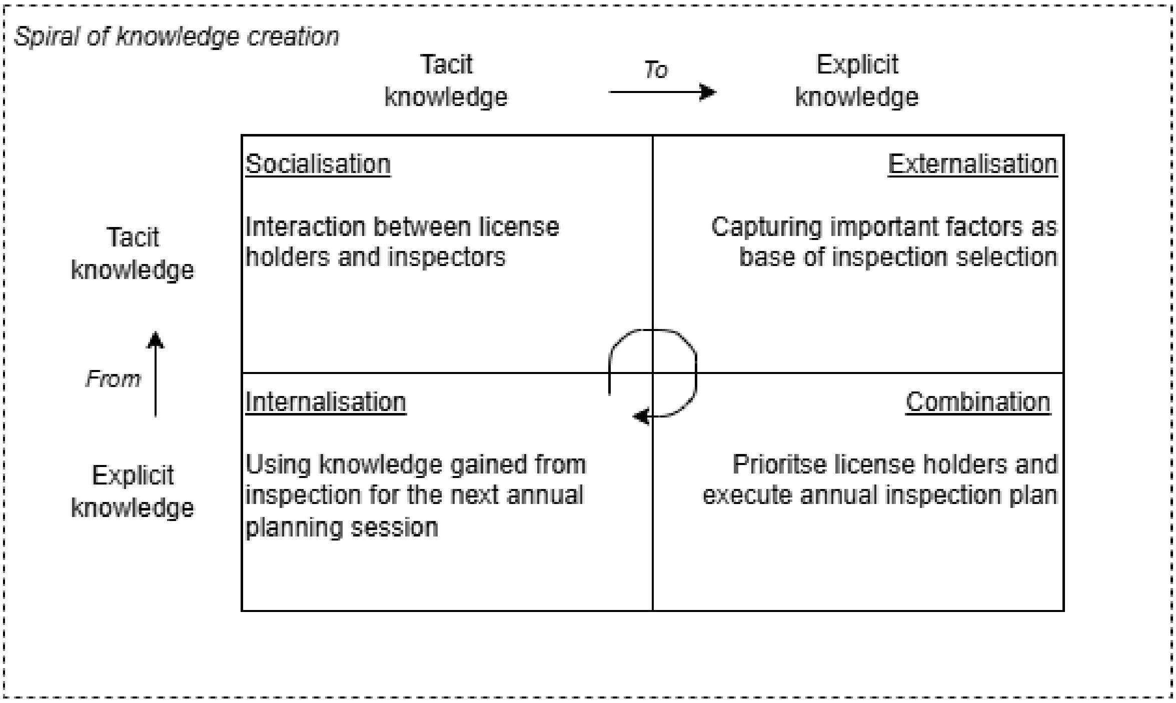

Figure 4.1: Spiral of knowledge creation applied to ANVS case, adapted from Nonaka and Takeuchi (1995)

As discussed in the literature review in chapter 2, knowledge creation based on human expertise and experience is subject to cognitive bias. This creates potential for data-driven approaches. While both expert judgement and data-driven methods come with their own forms of bias, they also have the potential to partly compensate for each other's blind spots. Therefore, the following section integrates both into the current system to study the potential of this combination.

4.3. ANVS' inspection strategy with data-driven RBI

When data-driven RBI is implemented, the process will mainly change in the selection and prioritisation phase of inspection planning. Rather than relying only on professional judgement and discussions at the annual meeting, inspectors will be supported by data-driven risk assessments that identify licence holders, branches or themes with high or increasing risks. As illustrated in figure 4.2, this is an extension for the externalisation step by translating inspection knowledge into formalised factors for data-driven RBI.

The blue arrows indicate two parallel processes. First, inspectors' externalised knowledge is translated into a data-driven RBI tool that supports inspection planning, after which the outcomes are discussed in the annual planning session and can be internalised again through inspection outcomes. Second, externalisation also continues through expert judgement and discussion, with both outcomes brought together in the annual planning session.

These analyses will provide quantitative input for planning and advise inspectors on which inspections can be prioritised in a more consistent way. The data-driven RBI tool operates between the externalisation and combination phases shown in Figure 4.2, by quantifying formalised factors into risk predictions that are used to prioritise licence holders in the inspection plan. The RBI outputs are intended to support inspectors by making risk considerations more explicit, while leaving room for critical reflection and discussion during the planning session. Final decisions on inspection priorities do remain with inspectors and project leaders, who assess the data-driven RBI outcomes in combination with their professional experience and contextual knowledge.

In this context, data-driven RBI functions as a decision-support tool. For RBI to be used in practice, the underlying assumptions and risk factors must be understandable to inspectors, and the results must be interpretable within the context of their inspection experience. This allows RBI outcomes to be internalised, as shown in Figure 4.2, and used as input for future planning meetings. This way, data-driven RBI supports and strengthens the existing inspection approach by making the inspector's work more efficient while also improving accuracy.

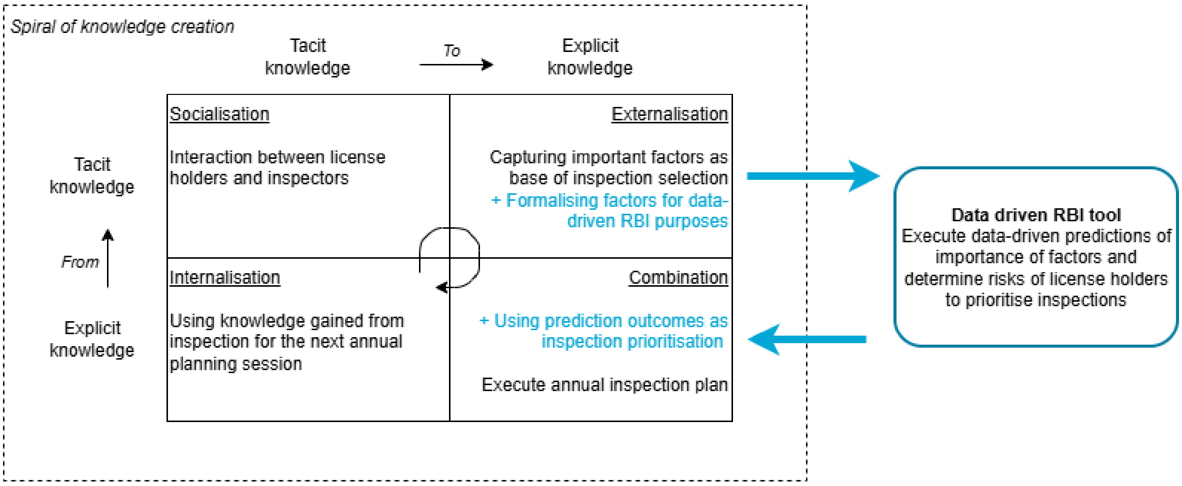

Figure 4.2: Spiral of knowledge creation including data-driven RBI applied to ANVS case, partly adapted from Nonaka and Takeuchi (1995)

4.4. Stakeholders

When developing a data-driven RBI framework, it is crucial to take into account the views of key stakeholders. Each group has a different level of influence over inspection planning and a different degree of interest in the outcomes. Identifying these stakeholders helps to understand how data-driven RBI can be used in practice.

The extension of data in RBI can be supportive. However, it must be accepted by inspectors, as they eventually have to use it. Inspectors play the central role in inspection planning decisions and are the main users of RBI. Data-driven RBI aims to support inspectors by making risk considerations more explicit and comparable across licence holders. To be effective, the output of the data-driven RBI advice must be understandable and easy to use. It should not require increased effort to integrate into the annual planning process. Inspectors should be able to keep their focus on the inspections themselves, including the internalisation and socialisation of knowledge updates. This ongoing knowledge exchange is important for maintaining the accuracy of the annual planning process.

Understanding the interest and behaviour of licence holders is also important. Inspection strategies are influenced by how licence holders respond to regulation. Gathering these insights means that policies can be designed to support a more robust data-driven RBI implementation in the long term. Besides this, licence holders may require a clear reference for why they are being inspected. The inspection strategy must therefore be explainable. Data-driven RBI also affect licence holders as they are directly affected by the inspection planning decisions and generally expect inspections to be fair. A more structured RBI approach can provide clearer insight into inspection priorities. However, differences in inspection frequency may also be perceived as unfair by some licence holders.

In addition, the proposed approach needs to fit within the regulatory context at both the national and European level ("Besluit basisveiligheidsnormen stralingsbescherming", 2025; European Commission, 2025). At both levels, supervisory authorities are expected to apply risk-based approaches to supervision, where the intensity of supervision increases with the level of risk. A consistent RBI approach can support these expectations by showing that inspection capacity is allocated based on risk considerations. These expectations follow from requirements from the European Union and from international guidance such as the International Atomic Energy Agency Fundamental Safety Principles (International Atomic Energy Agency, 2006). Together, these frameworks provide guidance on proportional regulation, continuous improvement, and transparency in supervisory practices.

Finally, for the ANVS as an organisation, this study provides insight into how data-driven RBI can support inspection planning in an accurate way and guides them towards a more data-driven inspection strategy. It highlights both the opportunities, such as improved consistency and efficiency, and the risks, such as blind spots and biases.

4.5. Conclusion - system description

This chapter described the regulatory and organisational context in which inspection planning takes place. It showed how RBI is already tightly incorporated into the organisation of inspection selection. Inspectors create valuable knowledge through their experience and interactions, which is captured through the spiral of knowledge creation and should not be overlooked. At the same time, the current way of defining the annual planning leaves room to further support inspection planning with data-driven methods, if relevant expertise can be carefully captured into a model.

This system description is therefore essential to understand how information should be externalised and how a data-driven approach could eventually be implemented in practice. Key guidelines that need to be considered include risk-based and accurate supervision, transparency in inspection selection, and the continuous role of professional judgement, as required by current national and EU frameworks. In addition, this system description provides guidelines on how structural changes coming with a data-driven RBI approach can be aligned with existing processes. They should support inspectors rather than replace them, and allow for learning, feedback, and adaptation over time, as also highlighted in the literature (Bhatia et al., 2019). This is important for moving towards accurate data-driven annual planning outcomes that remain consistent with regulatory expectations.

5

Risk indicators insights

This chapter addresses sub-question 1 (SQ1): *Which insights are considered important in determining the risk level of a licence holder?*

Determining the risk level of the licence holder is an important step in defining the annual inspection planning at the ANVS. This approach relies on quantitative information, such as historical inspection data and other licence holders' characteristics. However, in practice, ANVS inspectors also use contextual, behavioural, and organisational insights when determining risk levels. These insights are not easily translated into model input variables.

These experience-based insights can be described as tacit knowledge (Polanyi, 1966). To structure and formalise this, this thesis follows the spiral of knowledge creation described by Nonaka and Takeuchi (1995), which explains that tacit knowledge can be externalised into explicit and shareable information. In this study, inspectors' insights are identified and organised for future analysis with interviews.

SQ1, therefore, focuses on identifying which insights inspectors consider important when determining the risk level of a licence holder. The outcome is a structured set of risk indicators that reflects inspection processes and provides input factors for the quantitative modelling steps in SQ2.

5.1. Insights from literature

Before analysing which factors ANVS inspectors consider most important for determining the risk level of a licence holder, a first step is made by summarising which categories of risk indicators are already recognised in the RBI literature. Research across food safety, chemical and energy regulation consistently shows that an effective RBI approach requires multiple indicators to be considered in a systemic and transparent way. As described in the literature review (2), studies highlight four groups of factors:

- Compliance history: past violations are one of the strongest predictors of future non-compliance (May & Nikiforova, 2019; Zanabria et al., 2018). Repeated violations are often a signal for higher risk (Kiviniemi et al., 2025)
- Organisational size and operational scope: Larger and more complex organisations show higher inherent risk, while smaller organisations often struggle with fewer resources (Hobé et al., 2023; van Asselt et al., 2021)
- Expertise and internal safety systems: Trained staff, documented procedures, internal audits and certified safety officers strongly support compliance (Li et al., 2023; Rachman & Ratnayake, 2019)
- Safety culture behaviour: Although less quantifiable, culture, communication and risk awareness are strong predictors of compliance behaviour (Bayramova et al., 2024; Huang et al., 2019; Woreta et al., 2025)

These themes function as the guiding principles in the interviews with the ANVS inspectors to provide context for how their views fit into the general RBI practice.

5.2. Insights from ANVS inspectors

Six semi-structured interviews were conducted with experienced ANVS inspectors to answer SQ1, see table 5.1. A flexible interview guide was used, following the approach of Johannesson and Perjons (1996), which enabled key topics from the literature to be addressed while also allowing the inspectors to explain their reasoning in practice, provide examples, and introduce additional factors.

After six interviews, saturation was reached. The inspectors repeated the same points, and no new categories were found, indicating that the current insights collected were sufficient to answer SQ1.

Table 5.1: Overview of interviewees, appendix references, and roles

Interviewee	Appendix reference	Role
Interviewee 1	Appendix A.1	Inspector – ANVS
Interviewee 2	Appendix A.2	Inspector – ANVS
Interviewee 3	Appendix A.3	Inspector – ANVS
Interviewee 4	Appendix A.4	Inspector – ANVS
Interviewee 5	Appendix A.5	Inspector – ANVS
Interviewee 6	Appendix A.6	Inspector – ANVS

The results are presented in three parts: (1) Internal factors; (2) External factors; (3) Behavioural and communication patterns.

5.2.1. Internal factors

Internal factors describe the characteristics of the organisation and its internal practices. These factors are specific to the licence holders and are mostly stable over time, although their influence may vary depending on organisational circumstances.

Personnel and Continuity

Personnel stability was described as an important factor of organisational risk. Frequent staff turnover, unfilled vacancies, or changes in key positions may indicate limited continuity in safety management (see Appendices A.1, A.2, A.4, and A.5). Organisation size combined with radiation expertise could be an important predictor. Larger organisations often possess sufficient internal expertise to maintain stability, while smaller licence holders tend to rely on external radiation experts, which can result in weaker communication and slower response to safety issues (see Appendices A.2, A.3, and A.4). Although changes in personnel can also display patterns over time, the underlying stability and availability of qualified staff are primarily connected to internal properties of the organisation.

This theme describes internal organisational properties related to the stability of personnel and safety responsibilities over time. Organisation size and radiation expertise influence these dynamics but should be included in combination with other factors in future analysis.

Behaviour, Documentation, and Communication Culture

Behavioural aspects during inspections reflect the overall safety culture but are not easily quantifiable (see Appendix A.1). However, they can be represented by certain factors. The quality and timeliness of documentation are measurable indicators of an organisation's willingness to comply and safety attitude. Incomplete, delayed, or disorganised documentation often correlates to lower commitment to compliance, while structured and detailed documentation suggests more control and awareness (Appendices A.3, A.4, and A.5). Internal communication is also a critical part of safety culture. In settings where radiation experts are contracted from outside the organisation or where language differences within the organisation make communication more difficult (Appendix A.1), ineffective communication can lead to oversight or delays in corrective actions (Appendices A.2, A.3). Transparent and continuous communication supports robust safety management (Appendices A.3, A.4, and A.6). In addition,

Appendix A.6 highlighted that organisational openness during inspections, where inspectors can easily speak with multiple relevant employees, is seen as an indicator of a transparent and well-organised safety culture (Appendix A.6).

The theme represents observable indicators of safety culture and compliance behaviour. Broader behavioural aspects that cannot be directly observed are reflected indirectly through these factors.

Organisational Characteristics

Company profiles, such as financial stability, commercial pressure, and the type of safety service used, strongly influence the level of compliance (Appendices A.3, A.5). Organisations under financial pressure or focused mainly on productivity may give less attention to safety measures (Appendices A.1, A.3, and A.5). The size of an organisation was interpreted differently. Larger companies may have more complex systems and higher potential risks (Appendix A.1), while smaller companies often have less internal expertise and weaker communication (Appendix A.2). Company size, therefore, does not directly determine risk by itself but has underlying, internal aspects such as licence scope and available expertise (Appendices A.4, A.5), however, it has a significant impact in combination with other factors (Appendix A.6). It was therefore suggested to focus more on the size of the licence rather than the company size itself, as the risk is not directly related to company size but to the scope of the licence, including the number and strength of radiation sources and the types of activities carried out (Appendices A.4, A.5). The core business of an organisation also matters. When radiation is part of the main activities, safety is usually better integrated into daily processes than in organisations where radiation plays only a supporting role (Appendices A.3, A.5). In addition, it was also noted that recent contact with the ANVS is a helpful indicator, where organisations with regular communication towards the ANVS show better awareness of regulatory regulations, while long periods without contact create uncertainty about compliance (Appendix A.6)

Radiation Sources

The type and strength of the radiation sources used by a licence holder are important internal factors. Stronger or more complex sources are linked to stricter regulations and a higher chance of serious violations (Appendix A.4). In addition, the age of the licence and therefore also the sources potentially influence the risk score (Appendix A.6). The properties of these sources, therefore, influence both the inspection focus and the overall risk level.

These factors reflect intrinsic properties of the licenced activities that influence regulatory complexity, inspection focus, and potential risk severity.

5.2.2. External factors

External factors are influences that come from outside the organisation. These are often beyond the control of the licence holder but can still have a strong impact on compliance behaviour.

Economic and Technological Environment

Economic conditions can influence how much attention organisations give to safety. Financial instability or strong commercial pressure can lead to fewer resources for safety measures and training (Appendices A.1, A.3, and A.5). Technological developments also play a role. The fast introduction of new radiation devices, especially in hospitals, can create uncertainty about how to handle them safely (Appendices A.1, A.2, A.3, and A.5). In new or fast-growing branches, safety structures often take time to develop, which can temporarily increase the level of risk (Appendices A.3, A.5). When new technologies are introduced, new licences must be requested. This information can therefore be taken into account in risk assessments: if a new licence involves a recently introduced technology, it may indicate a higher probability of rule violations and should be given a higher risk score or inspection priority (Appendices A.4, A.5). In addition, national unemployment levels were mentioned as influencing organisational continuity and staff availability (Appendices A.2, A.4, and A.5).

This theme describes external conditions that influence organisational behaviour but are outside the direct control of the licence holder.

Incident-Driven Compliance

National or international radiation incidents can lead to a temporary increase in awareness. After such events, similar organisations often tighten their procedures and pay more attention to compliance (Appendix A.4). This effect is usually short-term but shows how external developments can influence behaviour.

This theme reflects short-term behavioural changes following external incidents.

5.2.3. Patterns

Patterns describe developments and changes that occur over time. In this research, these patterns are not used as data-based input for the ABM but are important to incorporate into the model to represent these dynamics. They can indicate where data should be corrected or where the output should be interpreted differently to make the results more reliable. Considering these patterns also helps to put unexpected outcomes into context and to discuss them in a structured way.

Compliance and Inspection History

Compliance history was mentioned as one of the strongest predictors of future violations. Repeated violations, often of the same kind, show that problems are not fully solved and may be structural (Appendices A.1, A.2, and A.6). Some improvement usually follows after an inspection, but this effect often fades over time and differs between branches (Appendix A.4). Inspection history shows a similar pattern. When inspections take place irregularly or not for a long time, licence holders tend to become less alert and more relaxed about compliance (Appendices A.1, A.4). Keeping inspection intervals regular helps to avoid this effect. In addition, all interviewees recognised a clear pattern in which the first inspection within a project generally shows lower compliance than later inspections. This is mainly caused by inter-branch communication, as information about the inspection focus is often shared between licence holders. However, some branches have stronger communication, due to the number of licence holders within a branch and the activity of the umbrella organisation. This pattern should therefore be taken into account when determining the risk score.

Based on this extracted information. This theme describes dynamic patterns over time and is therefore not treated as static organisational characteristics.

Temporal and Seasonal Trends

Several interviewees mentioned that certain periods of the year are more sensitive to risk. During holiday periods or times with high staff turnover, oversight tends to decrease, and the chance of mistakes increases (Appendices A.1, A.2, and A.4). These are seasonal changes and can be used to plan inspections more effectively.

Based on this extracted information, this theme captures recurring temporal variations that affect compliance and inspection effectiveness.

5.3. Conclusion sub-question 1

The interviews and literature review result in a structured overview of factors that ANVS inspectors consider important when determining the risk level of a licence holder. These factors include directly measurable characteristics, external developments, and qualitative and behavioural organisational factors, such as safety culture and inspection behaviour. Although some of these factors cannot be directly quantified, they are included to ensure that experience-based insights are systematically considered in later modelling steps.

The process follows the spiral of knowledge creation described by Nonaka and Takeuchi (1995). Inspectors' tacit knowledge, developed through inspection experience (socialisation), is captured and structured through interviews (externalisation). The resulting factors are then organised and combined with insights from literature (combination), which results in a structured representation of inspection knowledge. This can then be applied and tested in quantitative modelling (internalisation), where the outcomes of the model are used at the annual inspection meeting. There, the outcomes will be reflected upon and discussed by the inspectors (socialisation). This process of externalisation can then be repeated to improve the data-driven determination of risk.

Table 5.2 presents an overview of all factors identified in the interviews and literature. This overview forms the basis for SQ2, where a subset of these factors is selected, grouped, and transformed into latent variables for quantitative modelling. The selection and operationalisation of these latent variables is explained in Chapter 6. This way, answering SQ1 with this table represents the externalisation of tacit inspection knowledge into explicit factors that provide the basis for SQ2.

Table 5.2: Overview of all factors mentioned in the interviews

Category	Factor	Information source	Description	Source
Internal factors	Personnel stability	Interviews	Staff turnover, vacancies, and continuity in key positions influence the reliability of safety management.	Appendices A.1–A.5
	Documentation quality	Literature & Interviews	Completeness, timeliness, and organisation of documentation reflect safety attitude and preparedness.	Appendices A.3–A.5
	Communication quality	Literature & Interviews	Internal communication effectiveness; language barriers; communication with ANVS.	Appendices A.1–A.4, A.6
	Organisational openness	Interviews	Willingness of staff to talk openly during inspection; access to relevant employees.	Appendix A.6
	Company size	Literature & Interviews	Larger companies are more complex but have more expertise; smaller companies have fewer resources and may rely on external experts.	Appendices A.1–A.4, A.6
	Licence size/scope	Interviews	Number and strength of sources; number of applications; regulatory complexity of the licence.	Appendices A.4, A.5
	Financial stability	Literature & Interviews	Economic pressure, declining turnover, or bankruptcy risk affect attention to safety.	Appendices A.1, A.3, and A.5
	Core business	Interviews	Whether radiation is a main activity or a supporting function influences integration of safety.	Appendices A.3, A.5
	Radiation source properties	Interviews	Type, strength, mobility, and age of sources influence regulatory difficulty and potential risk.	Appendices A.4, A.6
	Use of external experts	Interviews	Outsourced radiation expertise may weaken internal communication or continuity.	Appendices A.2, A.3
External factors	Economic environment	Literature & Interviews	Sector-wide financial developments influence safety investment and behaviour.	Appendices A.1, A.3, and A.5
	Technological developments	Interviews	Introduction of new radiation devices; transition from sealed sources to X-ray or ultrasound.	Appendices A.1–A.5, A.6
	National unemployment	Interviews	National unemployment rates influence the continuity of organisations	Appendices A.1–A.5
Patterns	Compliance history	Literature & Interviews	Repeated violations signal structural issues and higher long-term risk.	Appendices A.1, A.2, and A.6
	Inspection history	Literature & Interviews	Long intervals between inspections correlate with weaker compliance; improved behaviour after inspection fades over time.	Appendices A.1, A.4, and A.6
	Inter-organisational communication	Literature & Interviews	Companies share inspection experiences; umbrella organisations influence preparedness.	Appendices A.4, A.6
	Seasonal patterns	Interviews	Holiday periods and times of high turnover increase the likelihood of mistakes.	Appendices A.1, A.2, and A.4
	Post-inspection effects	Interviews	Temporary improvement after inspection, followed by a decline as attention decreases.	Appendix A.4

6

Translating risk indicators

This chapter addresses Sub-question 2 (SQ2): *How can important risk factors be translated into an agent-based model (ABM)?*

The factors identified in SQ1 will be translated into a model structure that can be used for quantitative analysis. While in SQ1 the focus was on externalising inspectors' tacit knowledge into explicit factors, SQ2 focuses on selecting, grouping, and operationalising these factors using the data currently available within the ANVS. This chapter steps outside the traditional knowledge creation spiral, as visualised in chapter 4, figure 4.2. The aim is not to represent all detailed qualitative insights, but to provide an abstract representation that captures the main interactions influencing a licence holder's risk score.

To analyse how the identified factors interact and influence risk, a regression analysis is used. This represents the relative contribution of each selected factor to be quantified and used as a basis for modelling how risk scores evolve over time. As several factors overlap or are causally related, a causal relationship diagram (CRD) is developed to structure these relations. Based on this diagram, related factors are grouped into latent variables. These latent variables are not directly observable, however, they represent underlying interactions captured through combinations of measurable indicators (Spirtes, 2001). Together, they form the basis of the ABM. The ABM represents individual licence holders as agents working in the regulatory environment of the ANVS. Agents vary in terms of their organisational characteristics. Risk evolves over time as a result of both model dynamics and external influences.

The chapter starts with the design of the agent-based model. This is followed by the regression analysis, where latent variables based on the interconnected factors defined in chapter 5 are included in a regression model to predict the initial risk scores of the model. Finally, the translation from regression outcomes to yearly updates of model parameters is explained.

6.1. Structuring risk indicators for ABM modelling

6.1.1. Objective of the ABM model

The objective of the model is to analyse how a risk-based inspection approach develops over time when inspection selections are based on data-driven predicted risk scores that are uncertain due to simplified representations of licence holders' characteristics and behaviour. Every year, the model is updated with new information and learns from past experiences. This means that the predicted risk for each licence holder is recalculated every year, and the highest-risk licence holders are selected for inspection. This allows the model to assess how these selections influence the overall goal of the ANVS under uncertainty: reducing risk across the whole regulated environment by maintaining low individual risk levels for all licence holders. This is essential as even a single high-risk licence holder may result in an incident, and such incidents are unacceptable in this high-impact regulatory domain.

The general goal is to understand how this uncertainty affects the distribution of inspection selections. As the predicted scores are based on simplified relationships, individual licence holders may differ from these patterns, which can lead to differences between who is predicted to be high-risk and who actually is. The model, therefore, examines what patterns occur when high-risk licence holders are missed.

The model also tests how policy measures influence these outcomes. By tracking yearly statistics, such as changes in individual risk levels and inspection distribution. This makes it possible to identify which policies make the inspection strategy more accurate under uncertain conditions.

In short, the model evaluates how well data-driven inspection selection performs when predicted risk scores and behaviour developments are uncertain, and which policy choices support effective and low-risk supervision. The interactions between licence holders (agents) and the inspection environment are visualised in figure 6.1.

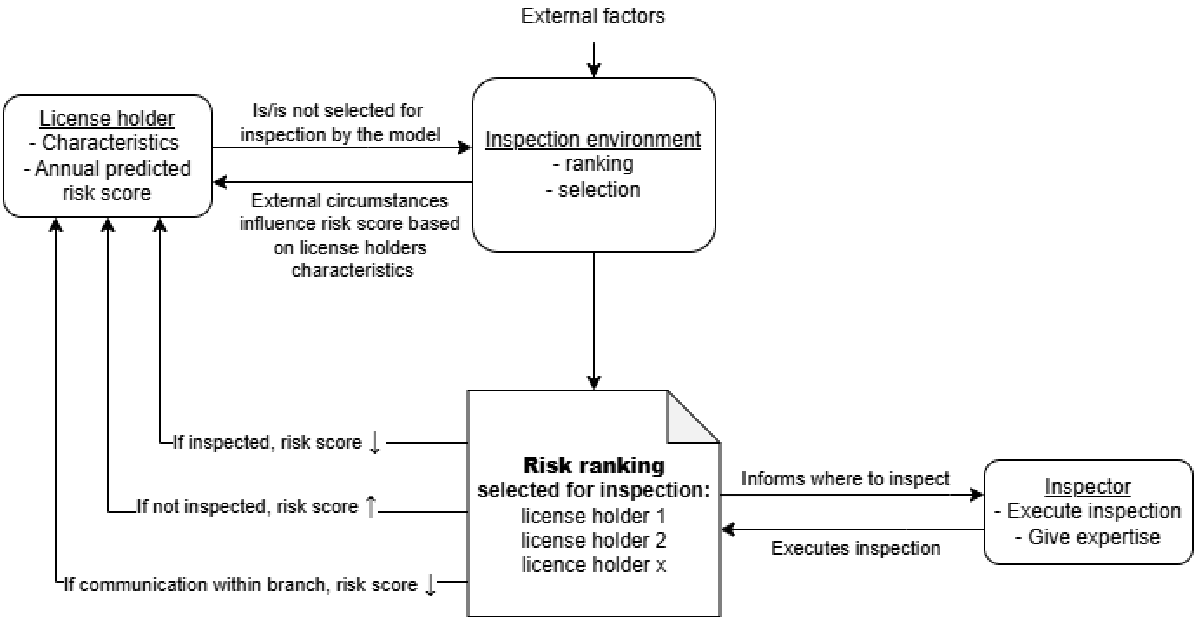

Figure 6.1: Agent-Based Model dynamics visualisation

6.1.2. Agents

The model will represent individual licence holders as agents. Each agent represents an organisation that holds a radiation license and can therefore be inspected by the ANVS. The agents differ in several characteristics that together determine their risk level and how that risk develops over time.

Each agent has a set of attributes that reflect both internal and external factors, based on insights gathered from the inspector interviews. These attributes are grouped into five underlying components that together determine the total risk score of a licence holder. The grouping is constrained by the information that is currently available within the ANVS databases. As a result, the latent variables represent a simplified version of the interview themes that can be operationalised using existing data, rather than a complete representation of all qualitative factors identified in SQ1. The definition of these latent variables is based on analytical modelling choices.

The causal relations, defined in the causal relationship diagram in appendix B.1, show how organisational, behavioural and contextual factors interact to influence risk. licence holders vary in characteristics such as company size, profit orientation, internal/external supervisor, radiation background of employees, and the number and type of sources. These influence five parts of how a licence holder behaves or operates: Financial stability (FS), Compliance behaviour (CB), Continuity (CT), Expertise (EX), and License Scope (LS).

- `Financial Stability (FS)`: influenced by whether the organisation is profit-driven, its size, and the general economic context (e.g. GDP growth).
- `Compliance Behaviour (CB)`: shaped by past notices and the time since the last inspection.

In addition, several behavioural patterns (feedback loops) apply: compliance tends to decrease when a licence holder has not been inspected for a long time. Compliance is higher when inspections occur later in a sequence within the same branch.
- `Continuity (CT)`: reflects organisational stability over time and depends on the type of supervision and trends in the national labour market.
- `Expertise (EX)`: relates to the technical knowledge of the organisation and is influenced by whether radiation work is part of the core activity, the background of employees, and whether a recent inspection has taken place. In addition, expertise improves after an inspection due to increased information from inspectors.
- `License Scope (LS)`: Is determined by the number and type of radiation sources of a licence holder. This is influenced by national innovation levels, which increase both the quantity and diversity of sources over time.

Based on the average of these five components, each agent receives a risk score between 0 and 1. This score is updated each simulation year based on changes in external conditions and behaviour patterns. How this will be updated in every step of one year will be further explained in the conceptualisation of the latent variables.

6.1.3. Environment

The environment in which these agents operate is an abstract, non-spatial decision environment that reflects the regulatory context. It does not simulate physical locations or spatial distance. Agents operate mostly independently, although some peer communication occurs through umbrella organisations concerning the content of the inspections, as defined during the interviews. This communication is modelled through the model environment, as a behavioural effect of the agents.

External developments, such as economic growth, unemployment, and innovation, affect all agents, but their impact depends on each agent's characteristics. The influence of these external variables on the five risk components is described in more detail in the following section.

The behavioural patterns listed below are included in the model based on insights from inspector interviews. They represent key feedback loops that drive agent behaviour over time:

- `Years since last inspection`: Compliance tends to decrease the longer a licence holder is not inspected.
- `Communication bonus`: Within a branch, licence holders that are inspected later in the year often show higher compliance. This is because inspection-related information is shared between licence holders within the same branch, as those inspected later can learn from earlier inspections of similar organisations.
- `Tightened behaviour after inspection`: An inspection results in improved expertise, due to feedback and increased awareness.

These dynamics are essential for going beyond the static risk-based model. They represent behavioural effects that are not visible in the initial risk indicators but play an important role in long-term compliance.

6.1.4. Time

The model runs on an annual time step, reflecting the real-world cycle of radiation inspection planning. Each simulation year represents one calendar year in which licence holders are selected for inspection, inspected and evaluated, and eventually receive an updated risk score. This yearly structure aligns with how ANVS prioritises inspections in practice in the annual inspection planning meeting and allows the model to capture long-term system and behaviour effects.

Agents update their internal state once a year. This includes recalculating the risk score based on their organisational characteristics, external developments, and inspection selections. Behavioural changes, such as increased expertise following an inspection, or declining compliance after years without inspection, are also updated at the end of each year.

Simulations run for a period of 10 years. This time span is long enough to observe how certain behavioural patterns (such as feedback loops or changing risk levels) evolve over time. The runs are used to explore potential long-term effects of changes in the risk landscape or inspection policy. As the model is relatively simple and does not include high-frequency decisions, the yearly time step is efficient while still suitable for the purpose of policy testing.

6.1.5. Technical model design overview

The inspection model is implemented as an agent-based model in which each agent represents a licence holder. Licence holders have a set of fixed characteristics, such as organisational and license characteristics. They also have dynamic variables that change over time, including a yearly risk score and the time since the last inspection.

The model runs in yearly time steps. At the start of each year, external factors are gathered, such as economic conditions and innovation rates. Risk scores are then calculated for all licence holders based on these external factors and their individual characteristics. The risk score is calculated from the predictor variables of the regression model, which represent different latent risk components. The regression model and its predictors are discussed in the following section.

The inspection environment ranks all licence holders based on their risk scores. A limited number is selected for inspection, this is determined by the inspection capacity. After inspections are carried out, inspection feedback is applied. This feedback updates the agent state variables and influences future risk scores.

For the full overview of the algorithms and variables used in the agent-based model, including the yearly update and inspection selection process, see appendices B.2.1 and B.2.2.

6.2. Quantifying inspectors' insights for ABM input

To initiate the regression analysis and provide input parameters for the ABM, the regression is performed using a predefined risk score as the dependent variable. This section first describes how licence holders are selected and how their risk scores are determined for the regression analysis. Second, it presents the regression formulas and resulting parameters that are used as direct inputs for the ABM.

6.2.1. Selection of licence holders as agents

As described in the system overview in chapter 4, annual inspection planning meetings have guided the selection of licence holders to be inspected in recent years. Since inspections are time-consuming, not all branches have been covered in previous years. For the same reason, even when a branch is included in the planning, not all licence holders within that branch are inspected. Under the current system, visit plans are prepared in advance. These are checklists of requirements for licence holders, and a different selection of checks is made for each project. Each project consists of a selection of licence holders from the same branch to be inspected within a given time period.

Over the past five years, the ANVS has implemented a structured system for inspections. Each year, an annual inspection plan is developed, as described in the system overview, along with a specific visit plan for each selected branch. These plans differ per branch. For each branch, a specific set of questions is created based on its characteristics, previous inspection experiences, and past violations. On average, each visit plan contains between 20 and 35 questions. Each question is assigned to a category representing the severity of the violation. The questions can be answered with "yes" or "no." In some cases, a "yes" response may include recommendations or warnings. "No" answers can also be marked as partly or fully corrected if improvements have been made. This results in a four-level scale:

Everything is in order (Yes): score = 0

Yes, with a recommendation or warning: score = 1

Partially violated or violation resolved: score = 2

Completely or almost completely violated: score = 3

To ensure comparability, data from the past five years has been used, including only licence holders for which an inspection report has been completed in line with the standard visit plan for that project. This approach ensures that all included licence holders are treated equally when translating their outcomes into risk scores. Ultimately, 128 licences were selected.

6.2.2. Initialising risk score for regression prediction

The risk score is calculated by reviewing the responses to each question on the inspection checklist. Each question receives a score from 0 to 3, which is multiplied by the severity of the violation, to calculate a weighted risk score. The overall risk score is determined by dividing the total points obtained by the maximum number of points that would be assigned if all rules were violated.

The formula for calculating each risk score is as follows:

$$R = \frac{\sum_{i=1}^{n}(s_i \cdot c_i)}{\sum_{i=1}^{n}(s_{\max} \cdot c_i)} \tag{6.1}$$

Where:

- R = total risk score,
- n = number of questions in the checklist,
- s_i = score for question i (ranging from 0 to 3),
- c_i = category importance or severity for question i (ranging from 1 to 4),
- $s_{\max} = 3$, the maximum category per question.

Since each branch has its own set of specific questions, note that the same type of question may impact the final risk score differently depending on the visit plan. For example, if one branch has 20 questions and another has 35, a single violation in the first branch is divided by 20, while in the second, it is divided by 35.

This difference should be kept in mind, but the weighting by category importance balances out some parts of such effects. More critical questions belong to higher-weighted categories and therefore influence the final outcome more. Furthermore, the difference between 20 and 35 questions is not expected to cause significant bias, however, it should still be considered per branch when interpreting the results.

6.2.3. Operationalising agents' latent variables in the ABM using regression

All five latent variables together determine the predicted risk score, which represents the overall risk score of each licence holder. This score eventually determines whether a licence holder will be inspected in a given year. External factors such as unemployment, economic growth, and innovation rate influence some of these relations indirectly, allowing the model to reflect changes in the environment. A full overview of all causal relations of the latent variables can be found in appendix B.1

To translate these latent components into usable model inputs, each of the five factors is operationalised using the results of the statistical analysis. For four components (`FS`, `CB`, `CT`, and `LS`), this is done using linear regression models that provide an intercept (α) and one or more coefficients (β). These coefficients represent the direction and strength of the relationship between each external variable and the corresponding latent variable, for every combination of categorical characteristics of the licence holder. These are predicted based on the The values of the intercepts and coefficients can be found in appendix B.3.

`Expertise (EX)` is the only component not based on regression. Instead, it is based on average differences: the difference between the overall mean risk score and the mean risk score of a subgroup defined by a licence holder's radiation expertise characteristics.

The following section presents the formulas used to compute each latent variable from the regression parameter of average differences, and describes how these values are applied in the model.

`Financial Stability (FS)`

`Financial Stability` represents the general economic purpose of an organisation. It reflects whether a licence holder has enough motivation, resources, and capacity to maintain safe radiation practices. FS is influenced by basic characteristics such as profit orientation and organisation size, in combination with economic growth.

$$FS_{p,s,t} = \alpha_{p,s} + \beta_{p,s} \cdot G_t, \tag{6.2}$$

- $FS_{p,s,t}$ is the financial stability score for profit orientation p and size category s in year t;
- $p \in \{0, 1\}$ indicates whether an organisation is non-profit (0) or profit-driven (1);
- $s \in \{\text{S}, \text{M}, \text{L}\}$ indicates whether the organisation is small (S), medium (M), or large (L);
- $\alpha_{p,s}$ and $\beta_{p,s}$ are the group-specific intercept and slope parameters;
- G_t is the GDP growth rate in year t.

`Compliance Behaviour (CB)`

`Compliance Behaviour` represents how well a licence holder has complied with rules in the past and how this develops over time. It depends on the number of past notices, for every notice in the past, the risk increases. In addition, this risk score reacts to the model dynamics, for every year of no inspection, a licence holder shows less compliance and therefore the total risk score increases.

$$CB_{n,t} = \alpha_{\text{CB}} + \beta_n \cdot N \tag{6.3}$$

- $CB_{n,t}$ is the compliance-behaviour risk contribution in year t;
- N is the number of past notices (integer);
- α_{CB} and β_n are the regression parameters.

`Continuity (CT)`

`Continuity` reflects how stable an organisation is over time, for example, in staffing, supervision procedures, and overall structure. Instability may increase the likelihood of mistakes or incomplete processes. `CT` depends on both supervision type and changes in labour market conditions, where changes in the labour market have a different influence on different sizes of the licence holder.

$$CT_{sv,s,t} = \alpha_{sv,s} + \beta_{sv,s} \cdot U_t, \tag{6.4}$$

- $CT_{sv,s,t}$ is the continuity-related risk contribution for supervision type sv and size category s in year t;
- $sv \in \{0, 1\}$ indicates whether supervision is external (0) or internal (1);
- $s \in \{\text{XS}, \text{S}, \text{M}, \text{L}\}$ represents the size category;
- U_t is the unemployment rate in year t;
- $\alpha_{sv,s}$ and $\beta_{sv,s}$ are the group-specific intercept and slope parameters.

`Expertise (EX)`

`Expertise` represents the technical knowledge and experience available within an organisation. Higher expertise reduces the chance of mistakes and unsafe situations. `EX` is based on average differences in risk between groups rather than regression. In addition, this risk score reacts to the model dynamics, if a licence holder is inspected recently, they gain expertise, so their total risk score decreases.

$$EX_{s,e} = \overline{R} + \left(\overline{R}_{s,e} - \overline{R}\right), \tag{6.5}$$

- $EX_{s,e}$ is the expertise score for supervision level s and radiation background e;
- $s \in \{0, 1\}$ indicates whether supervision is internal (1) or external (0);
- $e \in \{L, M, H\}$ indicates whether the radiation background is low (L), medium (M), or high (H);
- $\overline{R}$ is the overall mean risk score;
- $\overline{R}_{s,e}$ is the mean risk score of subgroup (s, e).

`License Scope (LS)`

`License Scope` represents the amount and diversity of radiation sources that an organisation works with. A larger or more complex scope increases the likelihood that something goes wrong if controls are not in place. In addition, if the innovation rate is high, there is a higher chance of new sources. As mentioned in the interviews, if there are new sources, there is less experience on how to comply, so this increases the risk score.

$$LS_{\text{src,tp},t} = \alpha_{\text{LS}} + \beta_{\text{src}} \cdot N_{\text{src}} + \beta_{\text{tp}} \cdot N_{\text{tp}} + \beta_i \cdot I_t, \tag{6.6}$$

- $LS_{\text{src,tp},t}$ is the license-scope risk contribution in year t;
- N_{src} is the integer number of radiation sources;
- N_{tp} is the integer number of source types;
- I_t is the innovation rate in year t;
- α_{LS}, β_{src}, β_{tp}, and β_i are the regression parameters.

`Total Risk Score (TR)`

The total risk score combines all latent-variable components into a single yearly value. For a licence holder with profit group p, size category s, supervision level s, radiation background e, notice history n, and licence scope (src, tp), The total risk in year t is given by

$$TR_t = \frac{1}{5}\left(FS_{p,s,t} + CB_{n,t} + CT_{s,t} + EX_{s,e,t} + LS_{\text{src,tp},t}\right). \tag{6.7}$$

6.2.4. From regression to yearly update of agent parameters

The values of the latent variables used in the ABM are based on separate regression analyses for four components (`FS, CB, CT and LS`). These regressions show how the normalised risk scores relate to organisational characteristics and external conditions. The results of the coefficients are used in the formulas in section 6.2.3. `Expertise (EX)` is based on group averages instead of regression.

Some external factors change over time. Therefore, the latent variables are updated at the start of every simulation year using the new values of unemployment, economic growth or innovation. In addition, some changes occur over time due to patterns occurring in the system. In the following section, the regression results will be discussed.

`Financial stability`

For `Financial Stability (FS)`, each influence of every combination of profit orientation and organisation size was predicted. The results show significant variation across groups. Profit-driven licence holders generally show stronger responses to changes in GDP growth. The largest positive effect is observed for large profit-driven organisations ($\beta = 0.030$, $\alpha = 0.307$), meaning that the risk score increases with economic growth. In contrast, small profit-driven organisations show a negative relationship ($\beta = -0.0106$, $\alpha = 0.290$), meaning that the risk decreases with economic growth. For non-profit organisations, only the large-size category shows a small negative coefficient ($\beta = -0.007$, $\alpha = 0.303$), meaning that changes in the economic growth have a less significant effect on non-profit licence holders. The results show how risk scores of licence holders holding different characteristics react to economic growth (appendix B.2, B.3.1.1).

`Compliance behaviour`

The regression for `Compliance behaviour (CB)` reflects the relationship between the number of past notices and the risk score. The estimated coefficient ($\beta = 0.008$, $\alpha = 0.298$) shows a small increase in risk scores in case of previous non-compliance, measured by the number of notices. In the model, this means that when a licence holder has had a notice in the current year or in earlier years, their compliance behaviour is predicted to be worse, resulting in a higher risk score. Since the regression is based only on notice history, the regression captures only the direct relationship between notices and risk score, and therefore a single regression coefficient is sufficient for the model (appendix B.3, B.3.2.2).

`Continuity`

For `continuity`, regressions are estimated for each combination of supervision type and organisation size. The coefficients show that the effect of unemployment varies across different combinations of size and supervision type. Among internally supervised organisations, the strongest effect is found in medium-sized organisations ($\beta = 0.005$, $\alpha = -1.961$). While larger (L) organisations show barely any relation strength ($\beta = 0.001$, $\alpha = -0.085$). For externally supervised organisations, large organisations show the highest coefficient ($\beta = 0.009$, $\alpha = -3.076$). This indicates that groups which are both externally supervised and larger in size are more sensitive to labour market conditions.

`Expertise`

The `expertise (EX)` variable is based on mean group differences instead of regression. The overall mean risk ($\alpha = 0.330$) serves as the baseline, and the group-specific correction (β) indicates how each combination varies from this average. The highest correction is found for internally supervised organisations with a low radiation background ($\beta = 0.090$), while the strongest negative correction is found in externally supervised organisations with a medium radiation background ($\beta = -0.087$) (appendix B.5, B.3.4.2).

`License Scope`

Finally, the regression for the `License Scope (LS)` shows that innovation has the strongest effect on the risk score, with a relatively large coefficient ($\beta = 0.063$, $\alpha = 0.350$). The number of sources shows only a very small negative effect ($\beta = -0.001$), and the number of source types also shows little influence ($\beta = -0.002$). Unexpectedly, these variables are not strong predictors of the risk score. Although their effects are small, the innovation rate in particular provides a basis for representing small differences in groups and for capturing how the introduction of new radiation sources, linked to higher innovation rates, as mentioned in the interviews, affects risk levels (Appendix B.6, B.3.5.2).

Overall, the regression outcomes provide a set of quantitative relationships that operationalise the latent variables for the ABM. The resulting risk score reflects both static organisational characteristics and dynamic external changes, together functioning as a solid basis for modelling how risk evolves over time. However, the coefficients carry uncertainty, as they generalise patterns from the available data and, in some cases, are derived from regression on small group sizes. These limitations will be assessed further through the uncertainty analysis in Chapter 8.

6.3. Conclusion - sub-question 2

SQ2 demonstrates how the identified risk factors in SQ1 can be translated into an agent-based model by translating them into quantitative model components with defined roles and yearly update rules. Cicence holders are represented as agents with fixed organisational characteristics, dynamic state variables, and are exposed to the dynamic environment. Together, these predict how risk develops over time within the model.

The translation from qualitative risk factors to the ABM followed three steps. First, externalised factors were grouped into five latent variables that represent the risk score: `Financial Stability (FS)`, `Compliance Behaviour (CB)`, `Continuity (CT)`, `Expertise (EX)`, and `Licence Scope (LS)`. Second, the regression analysis was used to quantify how these indicators contribute to the risk score and how they respond to yearly changes in external conditions and risk scores of other licence holders.

The regression analysis functions as the input dynamics strengths of this process. For `financial stability (FS)`, economic growth, combined with profit orientation and organisation, shows that risk responds differently to different organisation types. `Compliance behaviour` is mainly driven by inspection history, with past notices leading to higher risks. `Continuity` is most sensitive to labour market conditions, especially for larger and externally supervised organisations. `Expertise` is captured through group-level differences, which show variation in risk associated with radiation knowledge and the type of supervision. `Licence Scope` is mainly influenced by innovation rates, it indicates that the introduction of new technologies plays a larger role in risk development than the number or type of sources alone.

Together, these relations show how static organisational characteristics, external influences, and model dynamics form risk scores over time. By including these insights into the ABM, their effects can be made visual over time, with their interaction and responses to the environment. This translation makes inspector knowledge quantifiable and provides a basis for analysing inspection planning and policy effects in the following chapters.

7

Agent-Based Modelling output results

This chapter addresses Sub-question 3 (SQ3): *How does the model behave when inspectors' insights are included in different scenarios?*

Having defined the latent risk variables and behavioural patterns, this chapter implements them in the ABM to test what patterns and biases develop under different environmental conditions. The analysis starts with the baseline model, which functions as an initial check on the representativeness of the model and serves as a benchmark for interpreting all other experiments. It shows how risk levels, inspection selection, and behavioural pattern responses develop when external conditions remain stable.

Next, the model is tested under different external conditions that vary in economic growth, unemployment levels, and innovation rates. Comparing these scenarios shows which factors strongly influence the model patterns and which influence is more limited. The factors with a strong influence are taken into the uncertainty analysis, as they are most relevant for understanding how uncertainty affects the long-term evolution of the model in extreme circumstances.

Finally, this chapter analyses how the behavioural patterns identified in the interviews respond to small changes in licence holders' behaviour through a sensitivity analysis. To highlight where the model is stable and where small changes can lead to different outcomes.

7.1. Baseline model

The baseline model represents the inspection environment of the last 10 years (2015 to 2025). The purpose is to simulate how a data-driven RBI approach would have evolved if the estimated regression relationships had guided inspection selections on the annual plan in this period. The external variables follow real historical values, which makes the baseline a realistic reference for later scenario, sensitivity, and uncertainty experiments.

The model contains a fixed population of licence holders from different branches. Each agent starts with initial scores for `Financial Stability`, `Compliance Behaviour`, `Continuity`, `Expertise`, and `Licence Scope`, based on their characteristics.

The regression coefficients estimated in chapter 6 are applied directly in the baseline model. This means that each latent component evolves based on the initial state of the average historical relationships.

The simulation runs in yearly steps, each year:

- External variables are updated using historical data;
- Agents update their component scores based on these external conditions, their past behaviour, and previous inspection selections;
- Total risk scores are recalculated for each agent, based on these components;
- The licence holders with the top 20 highest risk scores are selected for inspections.
- Licence holders receive higher expertise scores following an inspection. Licence holders from

the same branch also gain expertise through communication. Meanwhile, compliance scores decrease when a licence holder goes an additional year without an inspection.

7.1.1. Results baseline model

The baseline results show a consistent pattern in all random behaviour inputs (see figure 7.1). This indicates that the base model is driven by the model's underlying relationships rather than by random variations. The average total risk remains stable between 2015 and 2019, followed by a clear increase in 2020. This is a result of changes in the external factors, mostly the peak in unemployment (see Appendix C.1). After 2021, the average risk decreases and stabilises at a lower average risk score, as the external conditions improve.

The standard deviation of inspection shows the same structure, with a peak of uneven inspection count distribution around 2020 and a decrease afterwards. The same pattern can be observed for the distribution of the years since a licence holder had their last inspection. However, this increase has a later reaction. This means that there is a strong relationship between the increase in the risk score and the distribution of inspections.

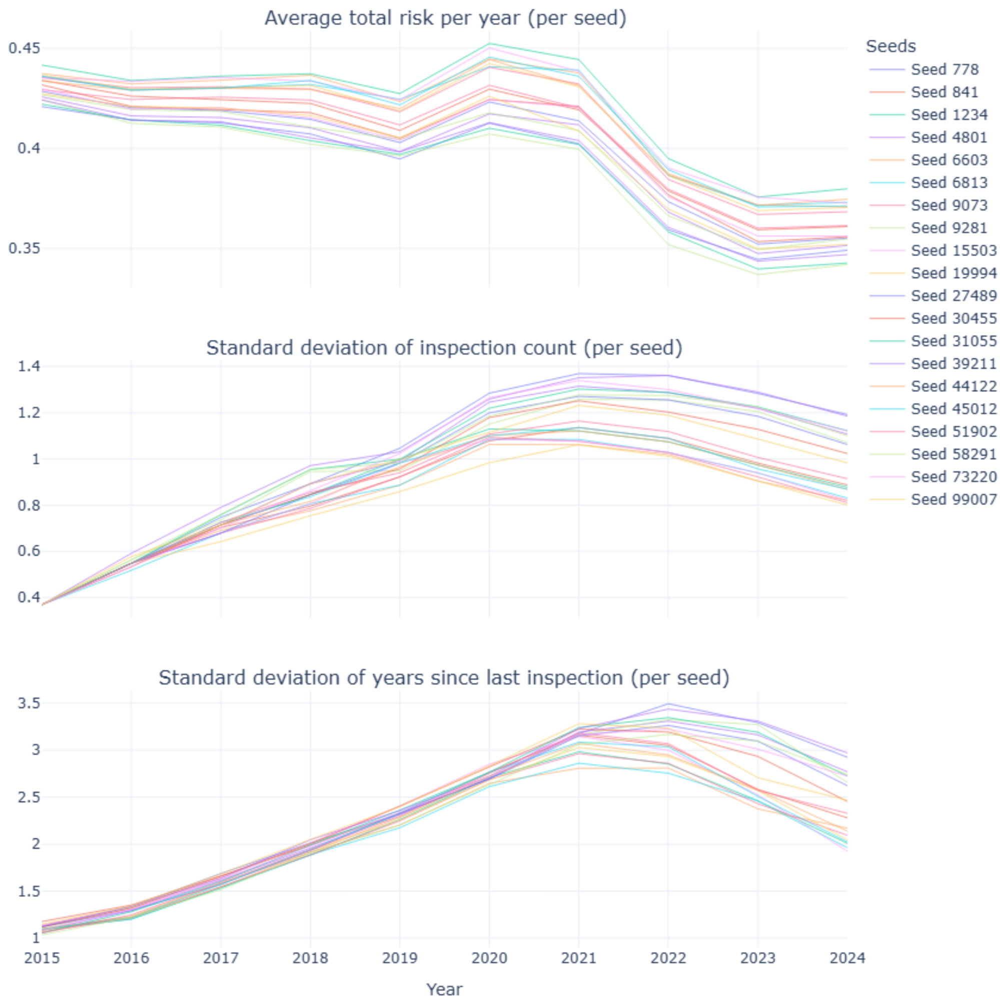

Figure 7.1: Baseline run

7.2. Scenario Analysis

The scenario analysis explores how changes in external conditions influence the development of risk within the model. Based on the interviews and the regression results, three external factors were identified as important drivers: economic growth, unemployment, and innovation rate. Because these factors can vary over time and are beyond the control of the inspection choices, it is essential to understand how different future developments may affect the overall risk levels of licence holders. In this section, several combinations of these external variables are tested to determine which scenarios lead to notable differences in model outcomes and which external conditions are most influential for the changes in risk over time. The scenario analysis is carried out through the following scheme.

Scenario	Economic Growth	Unemployment Rate	Innovation Rate
Scenario 1	Low	Low	Low
Scenario 2	Low	Low	High
Scenario 3	Low	High	Low
Scenario 4	Low	High	High
Scenario 5	High	Low	Low
Scenario 6	High	Low	High
Scenario 7	High	High	Low
Scenario 8	High	High	High

Table 7.1: Scenario overview: all combinations of economic growth, unemployment rate, and innovation rate

From the total set of scenarios, scenarios 1 and 3 are selected for further analysis. Scenarios 1 & 2, 3 & 4, 5 & 6, and 7 & 8 only differ in their innovation rates, which have no significant effect on the total risk score and therefore do not produce distinct patterns in the results (see appendix C.2). Although economic growth influences the final average risk score, the development over time does not significantly change for both high and low GDP growth.

The most relevant differences were found in scenarios 1 and 3, which vary in unemployment levels. This results in a clear difference in both the final risk level and the yearly risk development. Scenario 1 shows a steady decline (see figure 7.2), while scenario 3 decreases more slowly and then drops more sharply in 2022 due to the strong reduction in unemployment (see figure 7.3). As unemployment is an important predictor of continuity, the changes directly affect the risk scores. These two scenarios, therefore, provide the most informative contrast.

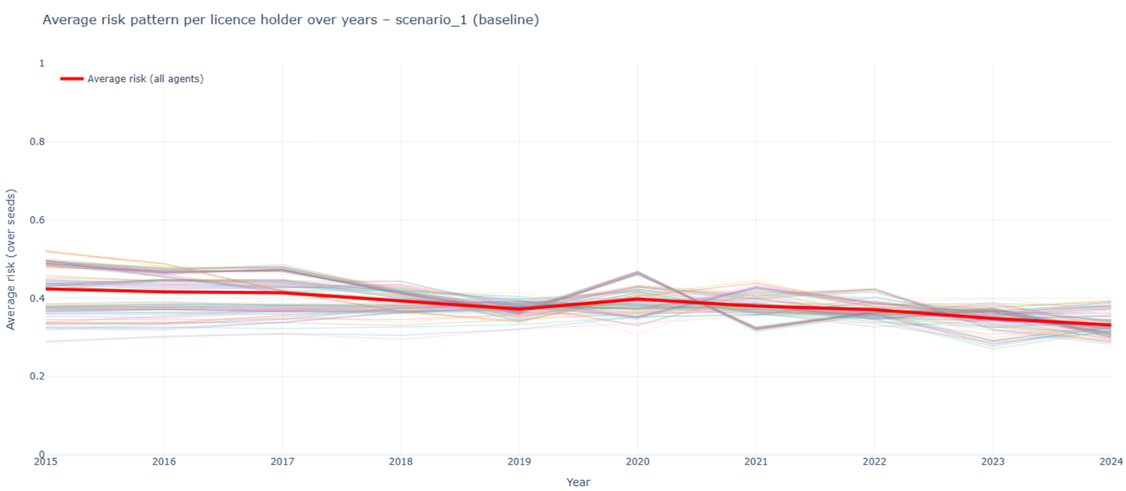

Figure 7.2: Risk patterns per licence holder - scenario 1

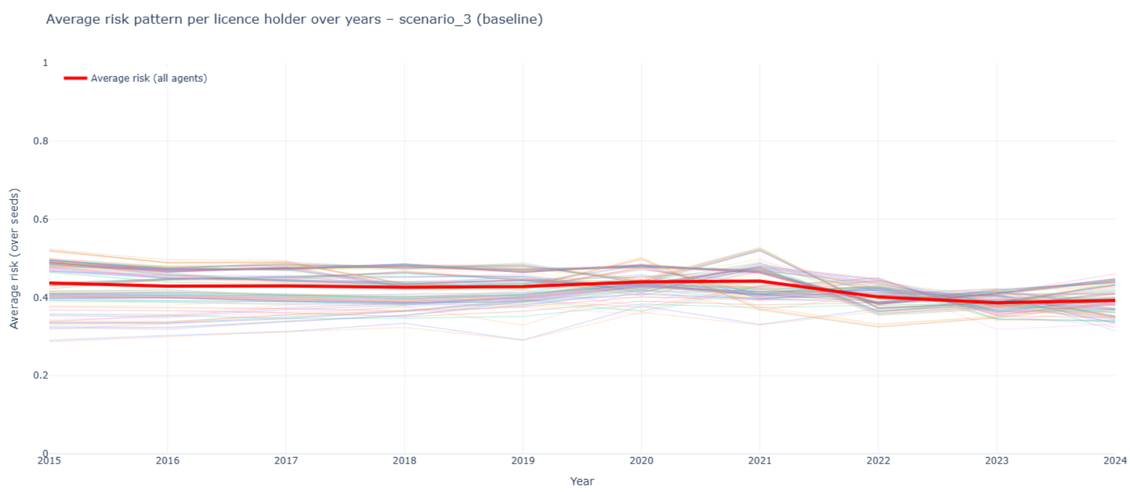

Figure 7.3: Risk patterns per licence holder - scenario 3

To understand how these conditions influence safety outcomes, the final-year risk per licence holder is analysed for both scenarios. In scenario 1 (see figure 7.4), where unemployment is low, the distribution of final risks across licence holders is relatively small. Most licence holders maintain similar levels of risk. This indicates a stable situation in which differences between licence holders stay and no specific group of licence holders stands out as the highest risk.

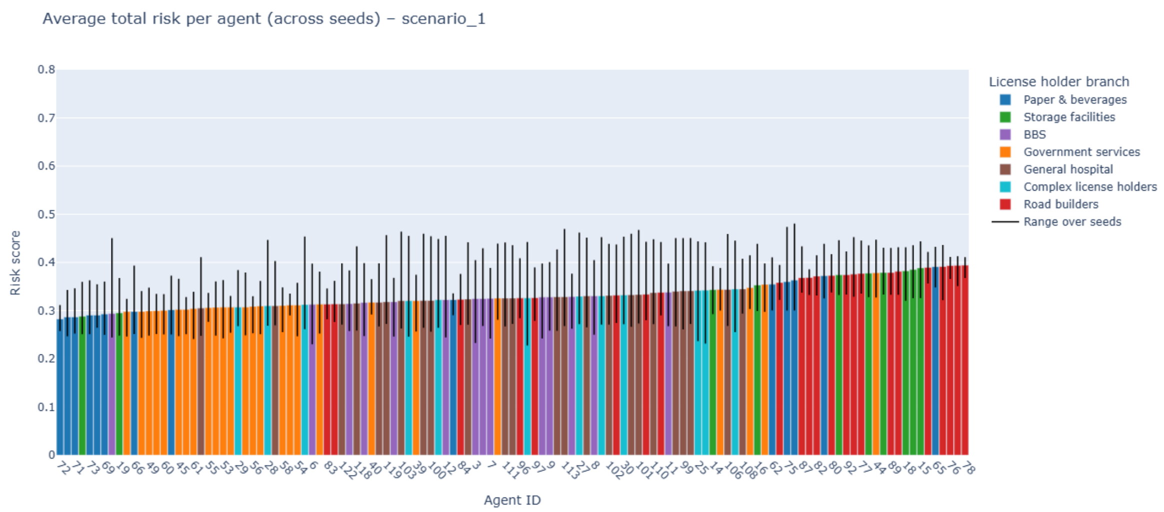

Figure 7.4: Average risk per agent final run - scenario 1

In scenario 3 (see figure 7.5), significant changes can be observed. Higher national unemployment rates lead to a clear increase in the overall risk level, but more importantly, they increase the spread between individual licence holders and groups of licence holders. The final-year risk shows a wider distribution. This demonstrates that unemployment rates not only increase average final risks but also increase differences between organisations, such as governmental services. For RBI, this insight is valuable, as it makes it easier to distinguish which licence holders require inspections. At the same time, it highlights the importance of using characteristics such as unemployment-related continuity, which determine risk scores (such as supervision type and company size), as external factors have a different effect on different groups of licence holders.

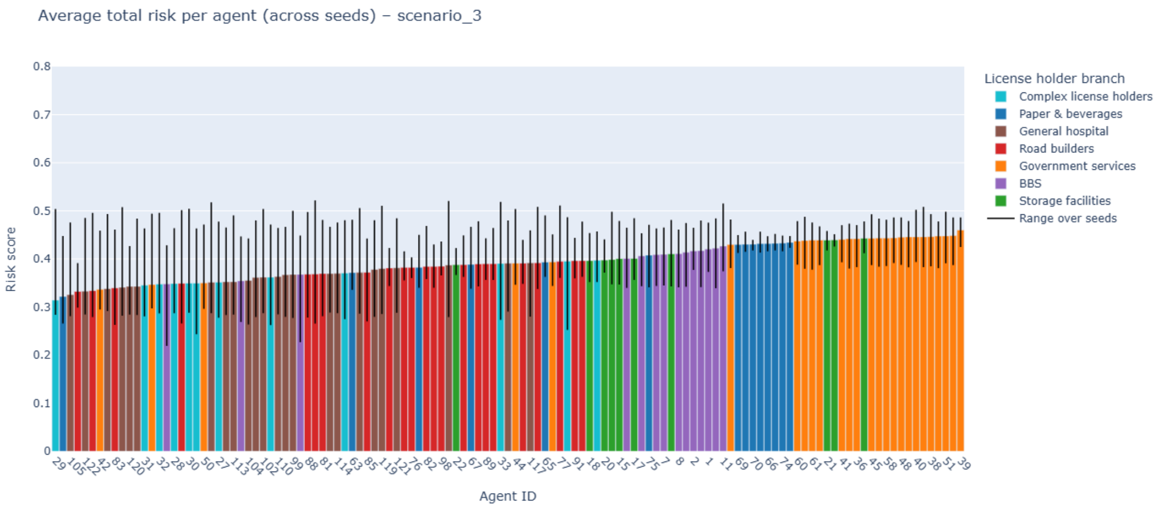

Figure 7.5: Average risk per agent final run - scenario 3

After analysing the final-year risks, the inspection count distributions provide insight into how the inspection selection algorithm responds to these different circumstances. Since the inspection-count plots are sorted by final-year risk, they clearly show how inspection activity clusters around the highest-risk licence holders in the high-unemployment scenario, while remaining more evenly spread in the low-unemployment scenario.

In addition, in scenario 1, the inspection counts remain relatively balanced. Licence holders rotate through the inspection schedule depending on whether they were recently inspected and on how their risk temporarily decreases afterwards. This creates a stable pattern in which the inspection allocation is spread across the licence holders without strong clustering. In figure 7.6, this fluctuation is observed between the highest and lowest risk, licence holders with either low or high risk receive relatively more inspections. This can be explained by the stability of this scenario, in which licence holders hold relatively similar risk scores. Once a licence holder is classified in the top 20 high-risk, it receives an inspection. Due to the fluctuations in external conditions and risk scores across years, this process results in a spread in inspection counts in the final year.

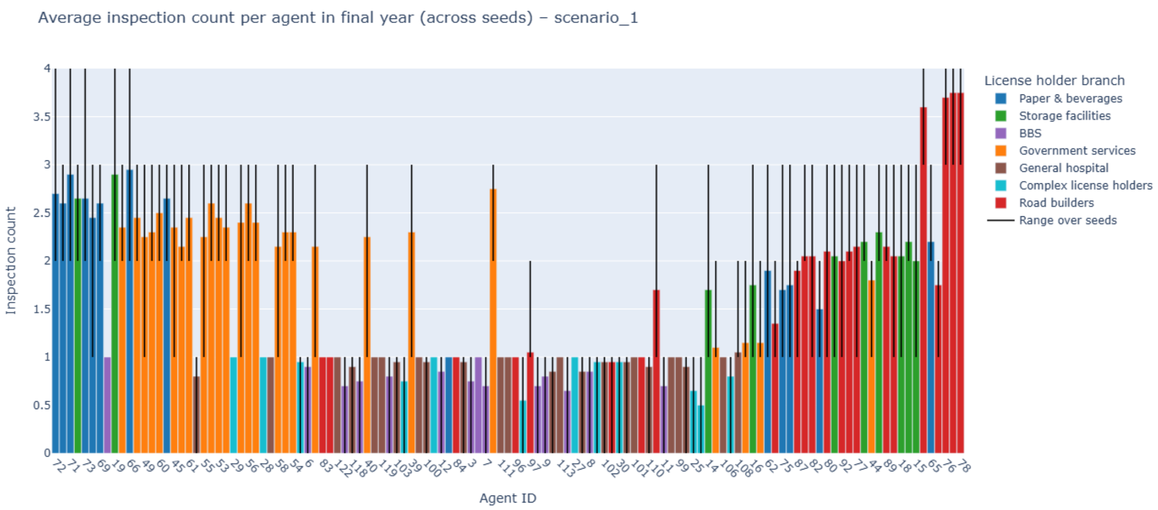

Figure 7.6: Inspection count per agent final runs - scenario 1

In scenario 3 (see figure 7.7), the inspection pattern becomes more clustered. The licence holders with the highest final risk receive significantly more inspections, reflecting the correct prioritisation behaviour of the data-driven RBI system. Simultaneously, some lower-risk licence holders show a rise in risk when they are not inspected for a longer period. The model's feedback mechanism ensures that these licence holders eventually rise in priority and re-enter the inspection selection, to prevent structural no-inspections. This effect is especially visible for specific branches, such as governmental services, which respond clearly to the high-unemployment scenario. This indicates that the model captures branch-specific dynamics rather than assuming uniform behaviour across all licence holders. However, because there are many high-risk organisations that need inspection due to high unemployment, the overall inspection pressure increases, and the distribution becomes more clustered towards the final high-risk licence holders.

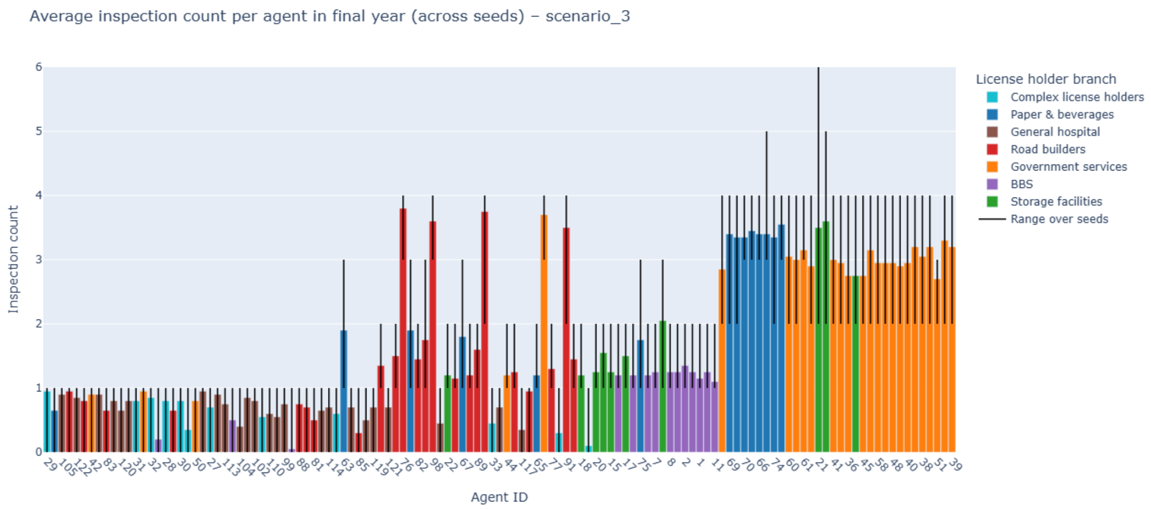

Figure 7.7: Inspection count per agent final runs - scenario 3

7.2.1. Sensitivity Analysis

From the interviews presented in chapter 5, several behavioural patterns were identified that influence how licence holders respond to different inspection dynamics. These include changes in compliance behaviour, improvements in expertise after an inspection, and communication effects within branches, for example, through umbrella organisations. In the base model, the strength of these patterns is calibrated to represent realistic behaviour, however, such behaviour may change over time. As they cannot be quantified directly from available data, their influence in the model involves uncertainty. It is therefore important to find out how sensitive the ABM is to changes in these parameters.

Three behavioural patterns were included in the model based on the interviews:

- The correction applied when a licence holder is not inspected for a long period,
- The increase in expertise following a recent inspection, and
- The influence of communication within a branch.

Each of these patterns is represented by a parameter in the model. In the sensitivity analysis, these parameters are adjusted one at a time by increasing and decreasing their values by 10%, resulting in six sensitivity runs. Building on the scenario analysis, each sensitivity run is applied to the scenarios identified as most relevant (scenarios 1 and 3), resulting in twelve runs. This structure provides a systematic basis for evaluating the robustness of the model outcomes under alternative behavioural assumptions.

Table 7.2: Overview of sensitivity settings

Scenario	No inspection correction	Recent inspection bonus	In-branch communication bonus
Scenario 1	-/+10%	-	-
	-	-/+10%	-
	-	-	-/+10%
Scenario 3	-/+10%	-	-
	-	-/+10%	-
	-	-	-/+10%

Sensitivity Analysis results

The sensitivity analysis shows that the model is most responsive to the behaviour of increased risk with every additional year without an inspection. This parameter directly affects the risk score of every licence holder each simulation year, and because it accumulates over time, even a small inspection correction has a significant impact on the model. As a result, the overall risk patterns show that with an increased negative attitude towards no inspection, the risk scores rise and the variance of risk scores increases. This is an unwanted effect. (See appendix C.3)

The influence of the parameter is also visible in the standard deviation of inspection counts and the standard deviation of years since the last inspection. When the correction is higher, inspection gaps widen more quickly, resulting in increased variability between licence holders. The inspection dynamics react to this by inspecting the highest risks, resulting in decreased variance. A stronger correction makes the inspection selection more fluctuating. Licence holders with many non-inspection years rise more quickly to the top 20, while others consistently remain below them. This results in similar patterns for different sensitivity seeds but with different frequencies and attitudes.

The distribution plots per licence holder in the final simulation year (2024) support the argument for the stability in standard deviation. A higher correction raises the lower risk scores as they do not receive inspections, while the highest remain relatively stable, resulting in a balanced spread. In contrast, the changes in the expertise-increase correction and the communication correction lead to only small differences in the results. This is consistent with the interviews, where the time since last inspection was repeatedly identified as a strong behavioural factor. This is reflected in the model choices: the correction for years without inspection increases linearly each year, the expertise effect decreases exponentially over time, and the communication effect is only applied when an inspection occurs and therefore influences only part of the licence holders. Because this modelling structure follows directly from the interview insights, the strong sensitivity for the long-time-no-inspection pattern is logical and was expected.

Overall, the results show that the years-since-last-inspection pattern is the dominant behavioural driver in the ABM. This pattern is modelled directly according to the insights from the interviews. It increases automatically over time, influences the yearly risk update, and strongly shapes the inspection ranking. As a result, the model is very sensitive to this parameter, which follows logically from the modelling choices based on the interview findings (Chapter 5).

7.3. Conclusion - sub-question 3

This chapter examined how the model behaves when quantified inspectors' insights are included through different scenarios. The scenarios represent variation in external conditions of economic growth, unemployment rates, and innovation rates.

The results show that the unemployment rate has the strongest influence on model behaviour. An increase in unemployment leads to higher overall risk scores and to larger differences between licence holders. These differences arise from variation in organisational characteristics and supervision types. This aligns with the interview results, where inspectors indicated that different licence holders respond differently to external circumstances. As a result, the risk score of licence holders shows more variability, and there is also greater variation in which licence holders are selected for inspection. This supports inspection prioritisation, however, at the same time, the overall increase in risk levels may complicate inspection planning. This effect will be analysed further in the uncertainty analysis.

In these scenarios, risk scores fluctuate over time, however, they do not show extreme or unstable patterns for individual licence holders. This is caused by the defined inspection feedback loop, which limits uncontrolled risk development. This behaviour reflects the importance of including this inspection feedback as described by inspectors in the model dynamics.

Overall, the results show that the inspection strategy remains relatively stable across scenarios. However, under more challenging conditions, such as scenario 3, uncertainty increases the likelihood of high-risk licence holders being missed. This highlights the importance of including uncertainty in a data-driven risk-based inspection planning model. Chapter 8 therefore analyses uncertainty in more detail, followed by an evaluation of potential policy alternatives in Chapter 9.

8

Uncertainty in data-driven RBI

This chapter addresses Sub-question 4 (SQ4): *What patterns evolve when inspections are missed due to the estimation of risk scores under uncertainty?*

In this chapter, the risks of using data-driven RBI under uncertainty are analysed. When annual inspections are planned based on the calculated input variables of the base model, the predicted risk ranking may not fully match the actual risk of the licence holder. As a result, the ranking differs, and some high-risk licence holders may be missed, while others with lower risk may be selected. SQ4 analyses how these patterns develop over time when uncertainty is added to the estimated risk scores.

To answer this sub-question, key performance indicators (KPIs) are defined that focus on the effects of uncertainty in the estimated risk scores. These KPIs measure how often inspections are assigned to licence holders that would not be selected if the true risk ranking were known, and how often inspections of higher-risk licence holders are missed. In addition, the KPIs are used to analyse how the model responds to these missed inspections over time, by identifying patterns such as uncontrolled increases in the risk scores of licence holders. This way, the KPIs make it possible to evaluate whether the efficiency benefits of data-driven RBI outweigh the impact of uncertainty on inspection accuracy.

8.1. Risk of inspection inaccuracy

The goal of data-driven RBI is to increase inspection efficiency while remaining accurate. For the ANVS, this means reducing the average risk and avoiding situations in which high-risk licence holders are missed, since each one could potentially lead to a hazard. Analysing only the average risk scores per year does not capture this variation. The mean may appear acceptable, while individual licence holders still hold high risks. Therefore, both the average and the spread of risk scores must be considered.

To test this, the number of missed inspections is analysed by comparing which licence holders are inspected in the base model and which are inspected in the uncertainty runs. The base model represents the inspection selection based on the true risk scores, while the uncertainty runs include variation in the estimated risk scores. This comparison shows how often high-risk licence holders are overlooked due to uncertainty in the risk estimation. The variability across different random seeds is also included. Each seed represents a different realisation of uncertainty in the risk score estimation, which can lead to different rankings of licence holders and, as a result, different inspection selections.

The result of these missed inspections could be that licence holders show unexpected patterns in their risk scores. Therefore, the worst-risk paths are visualised for each licence holder by taking, for every year, the uncertainty run in which that licence holder had the highest risk. Figure 8.5 and 8.6 show these worst-case paths for scenario 1 and scenario 3, together with the average risk score over all licence holders, to allow comparison.

8.1.1. Found risks

To understand how uncertainty affects inspection decisions, the first step is to analyse the number of missed inspections across all uncertainty runs. Figures 8.1 and 8.2 show for multiple random values of uncertainty (seeds) how many licence holders are selected in the base model, while not selected in the corresponding uncertainty run. Each line represents a different combination of uncertainty in the latent variables, which estimate the risk scores and the resulting inspection selection.

These patterns should be evaluated in the context of different external values for unemployment in scenarios 1 and 3. The change of unemployment over time is shown in figure C.3. Between 2015 and 2018, unemployment decreased steadily, which resulted in lower risk levels in the model. However, despite these relatively stable external conditions, scenario 1 still shows strong fluctuations in the number of missed inspections during this period. This can be explained by the fact that many licence holders have very similar risk scores in this scenario, as shown in figure 7.4 in chapter 7. When risk scores are close together, small differences caused by uncertainty can already lead to changes in the ranking and inspection selection.

Around 2020, unemployment increases temporarily, which leads to a general increase in risk scores, which pushes more licence holders towards the inspection threshold. As a result, the uncertainty in the model has a strong effect on the number of missed inspections of high-risk licence holders. This results in the peaks around 2020. Especially scenario 3 has a very strong immediate reaction. Scenario 1 has a more moderate reaction to this change.

After 2020, unemployment decreases again, and overall risk levels are more stable. In scenario 1, the number of missed inspections decreases, however, the variance between seeds remains visible. In scenario 3, the number of missed inspections decreases very fast and stabilises around two or three missed inspections per year. Together, this indicates that under more negative external circumstances (scenario 3), the model is eventually able to correctly capture the risk licence holders, as shown by the fast decrease after the peak. However, the model does not immediately identify these licence holders when a change in external conditions occurs. The peak, therefore, represents a significant inaccuracy which should be avoided. In more stable and favourable circumstances (scenario 1), the model mainly fluctuates due to a small change in the ranking caused by uncertainty.

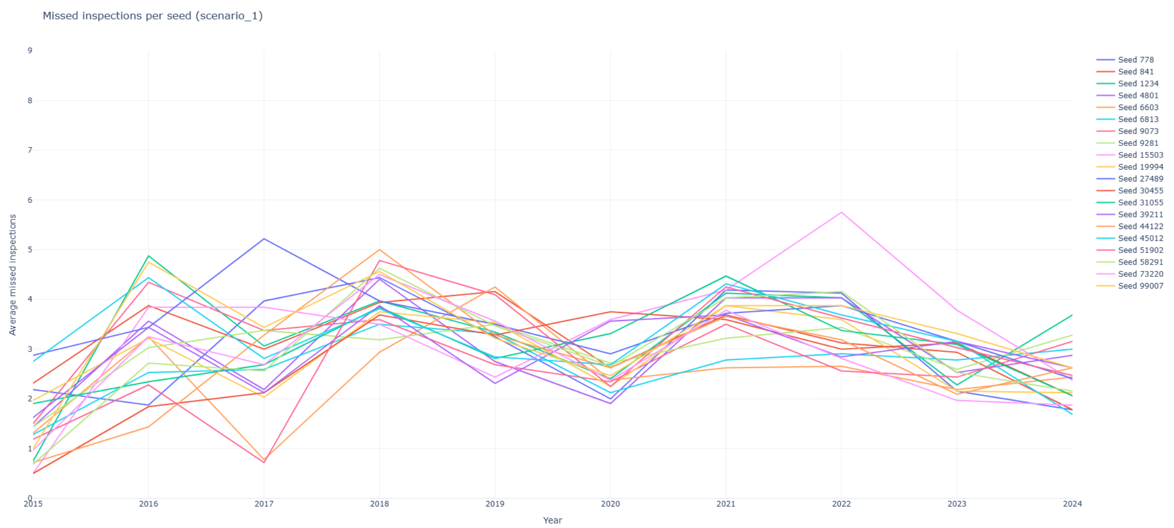

Figure 8.1: Missed inspections - scenario 1

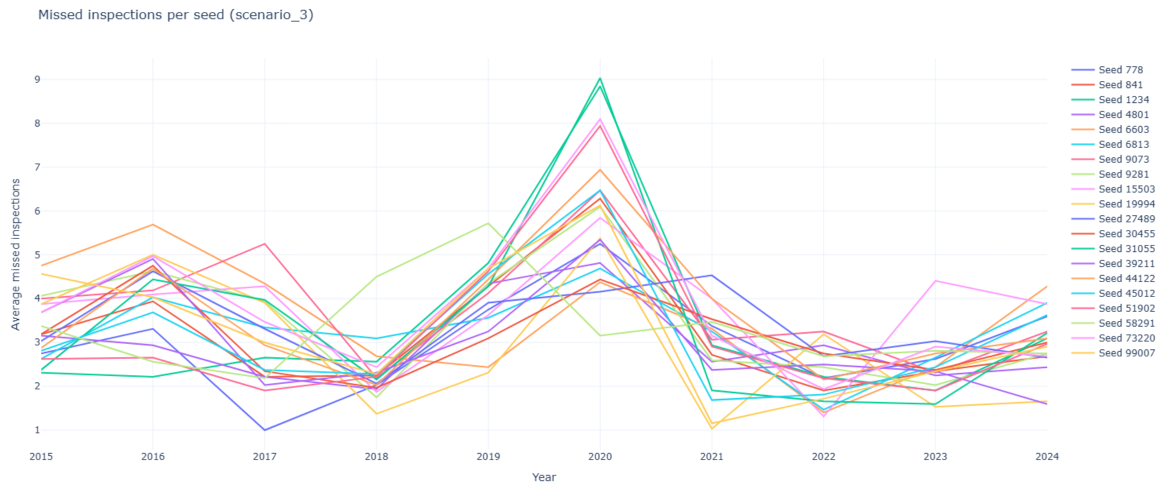

Figure 8.2: Missed inspections - scenario 3

However, many of the licence holder included in the set of missed inspections have low or medium risk scores, meaning that missing their inspection in a specific year is not immediately problematic. To separate those from genuine risks, a threshold is set to define what a high-risk licence holder is. The plots with a threshold of 0.3 include many missed inspections of licence holders with low or medium risk scores (see appendix D.1 and D.2). The 0.5 threshold, in contrast, captures only extreme risks and results in very few missed inspections (see appendix D.3 and D.4). A threshold of 0.4 is therefore used, as it corresponds to the upper one-third of average risk scores and focuses the analysis on relevant missed high-risk inspections. Figures 8.3 and 8.4 therefore show only the missed inspections for licence holders whose risk exceeds this threshold.

8.2. High-risk cases

When focusing only on high-risk cases, the risks become much clearer. In both scenarios, the number of missed high-risk inspections is significantly lower than the total number of missed inspections (Figures 8.3 and 8.4). This shows that licence holders with the highest risk scores are, in most cases, still selected for inspection. As a result, most of the variability in missed inspections is found for licence holders with lower risk scores, where changes in the ranking have less impact on high-risk inspection accuracy.

In scenario 1 (see figure 8.3), missed high-risk inspections are rare, as in many years, no high-risk licence holders are missed at all. This indicates that the high number of missed inspections shown in figure 8.1 is mainly caused by fluctuation in the lower part of the ranking. This supports the suspicion when analysing these results in combination with figure 7.4, where the average risk level remains relatively stable. This is also reflected in the relatively low and stable number of missed high-risk inspections over time.

Scenario 3 (see figure 8.4) shows a different pattern. Due to the higher overall risk levels, more licence holders are pushed towards the high-risk threshold. This leads to multiple years in which one or two high-risk licence holders are not inspected. Around 2020, a clearer peak in missed high-risk inspections can be observed across the seeds. During this period, more licence holders receive similarly higher estimated risk scores, which makes the inspection selection more sensitive to shifts in the ranking, resulting in fast shifting risk rankings under uncertainty.

Even under these more challenging external circumstances, the number of missed high-risk inspections remains stable. Overall, these results show that data-driven RBI is relatively robust in prioritising high-risk licence holders under uncertainty. However, when multiple licence holders receive high-risk score estimations at the same time, uncertainty can still lead to more missed high-risk inspections.

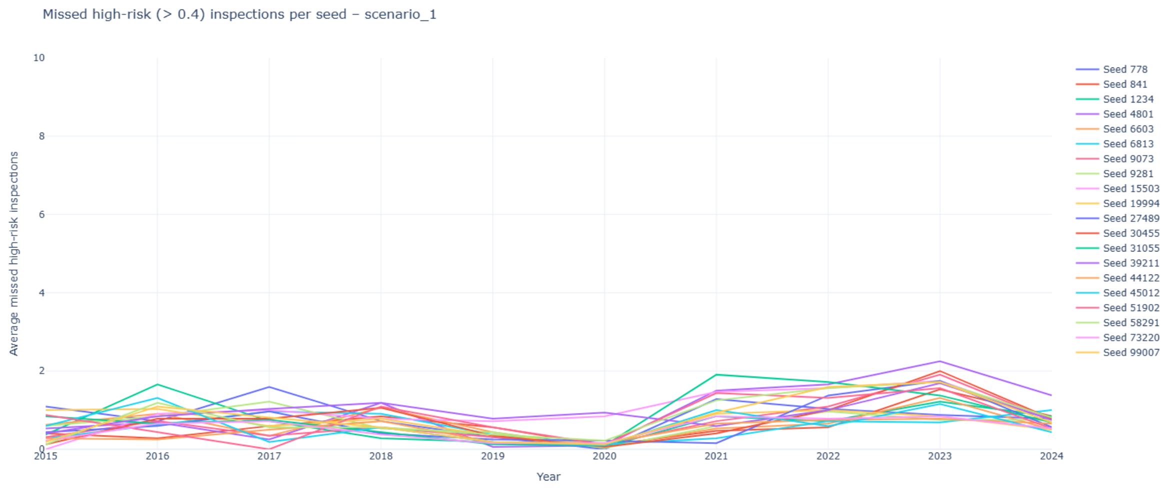

Figure 8.3: Missed high risk (>0.4) inspections - scenario 1

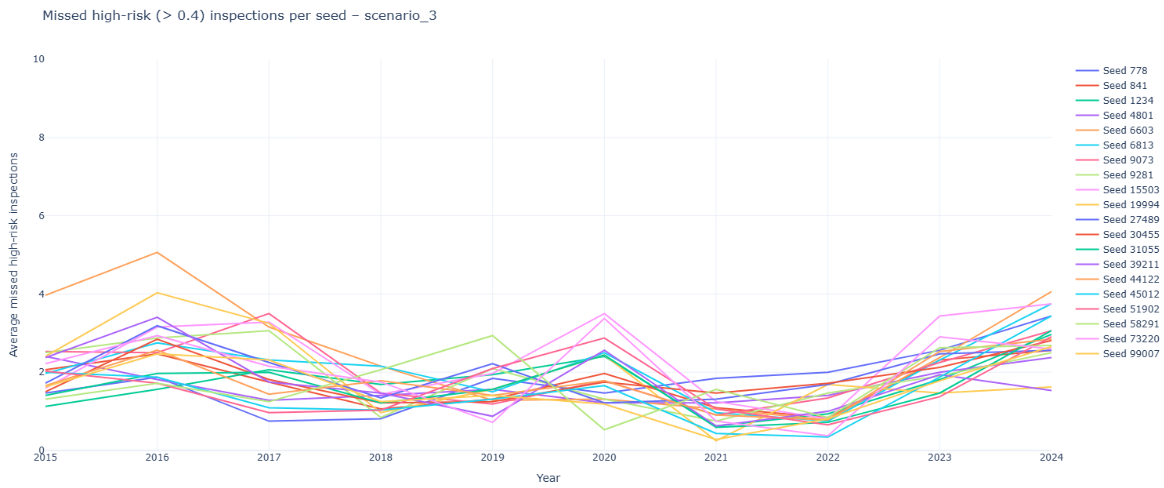

Figure 8.4: Missed high risk (>0.4) inspections - scenario 3

As defined during the interviews and implemented in the model, not inspecting a licence holder's results in an increase in risk score, which reflects reduced compliance over time. Missing a high-risk licence holder is therefore problematic for two reasons. First, it already implies a high immediate risk for that year. Second, when a high-risk licence holder is not inspected for multiple consecutive years, the risk score can increase further, potentially leading to extreme risk levels under uncertainty.

To assess whether such worst-case developments occur, Figures 8.5 and 8.6 focus on the most extreme risk patterns of individual licence holders. These figures show whether, under the current uncertainty variations, there are cases in which licence holders show uncontrolled risk behaviour that could significantly harm radiation safety. This additional analysis considers whether data-driven RBI under uncertainty can occasionally result in missed inspections, as well as whether this can lead to extremely high risk levels in the long term.

The results show that no licence holders get into an uncontrolled risk pattern over time. Although scenario 3 shows a wider spread and slightly higher worst-case risk levels compared to scenario 1, there are paths that remain bounded and do not develop into extremely rapidly increasing risk levels. Although the model does react to the missed high-risk inspections for multiple licence holders (see Figures 8.3 and 8.4), especially around 2020, where in figures 8.5 and 8.6 there is a slight increase in risk scores for multiple licence holders at the same time. These increases do not get out of control and are eventually captured again by the inspection model, which prevents risk scores from increasing to extreme levels. However, this does show that missed inspections result in more variation in risk patterns and high risk scores among multiple licence holders around the same period.

This stability in the worst-case risk patterns helps explain that relatively few high-risk licence holders are missed in the uncertainty runs, as seen in figures 8.3 and 8.4. Even under worst-case uncertainty, most licence holders do not reach extreme or fast-increasing risk levels that would exceed the capacity of the inspection resources (see figures 8.5 and 8.6). Instead, uncertainty mostly causes licence holders with medium risk levels to shift positions within the ranking (see Figures 8.1 and 8.2), resulting in the fluctuations observed in the missed inspection plots. Only in scenario 3, where overall risk levels are higher and more licence holders are close to the high-risk threshold, does uncertainty more frequently lead to high-risk licence holders not being selected for inspection.

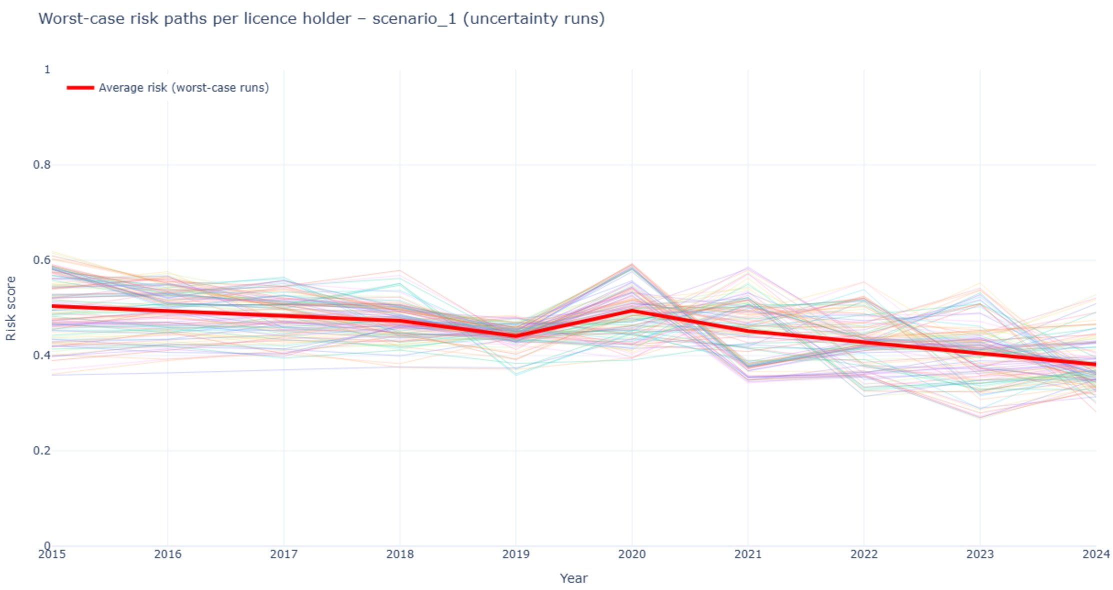

Figure 8.5: Worst risk patterns per licence holder - scenario 1

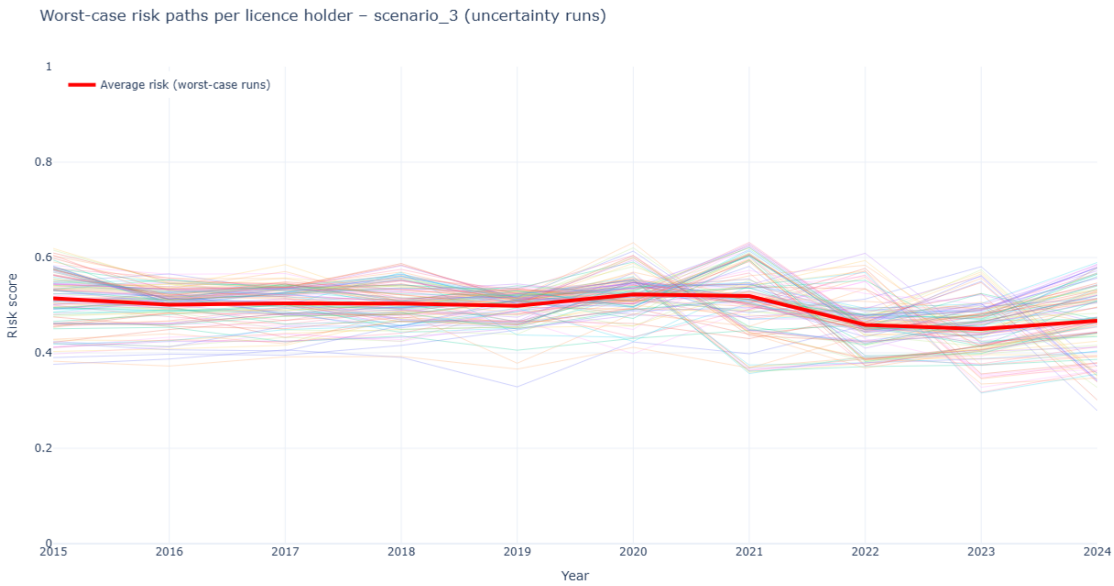

Figure 8.6: Worst risk patterns per licence holder - scenario 3

8.3. Conclusion - sub-question 4

The results show that uncertainty in the predicted risk scores can cause several licence holders to be missed during inspection planning, including a small number of high-risk licence holders. Although the number of missed high-risk inspections is generally low, these misses occur more often during periods with unfavourable and highly influential conditions, such as higher unemployment. In these periods, uncertainty affects the ranking of licence holders, which reduces inspection accuracy temporarily. However, the worst-case analysis shows that these increases remain bounded and do not develop into extreme risk levels. This shows that increases in missed inspections are mainly temporary and result from changes and uncertainty in prioritisation, rather than from structural blind spots in the model that cause high-risk licence holders to be missed.

These findings show that data-driven RBI is relatively robust, but also sensitive to uncertainty, especially when external conditions change. For this reason, chapter 9 evaluates several policy options to reduce the negative effects of uncertainty on inspection selection. These include partly random inspections, improved data quality, and an increase in inspection capacity. The goal of these policies is to reduce missed high-risk inspections while maintaining the inspection efficiency gained with data-driven RBI.

9

Reducing risks of data- driven RBI under uncertainty

This chapter addresses Sub-question 5 (SQ5): *What policies can reduce the negative effects of data-driven RBI under uncertainty?*

Several policies are tested to assess which ones have the highest potential to reduce the negative effects, found in chapter 8, of implementing data-driven RBI under uncertainty. The focus is on policies that limit the number of missed inspections, especially for licence holders with high-risk scores.

Although the results from the previous chapter do not show strongly uncontrolled patterns, the aim is to implement data-driven RBI in this high-stakes environment. Every missed high-risk inspection increases the probability of an accident, and even a small number of accidents is unwanted. For this reason, reducing missed high-risk inspections is essential to implementing accurate data-driven RBI in the future.

Previous studies show opportunities to benefit from the efficiency of RBI while maintaining inspection accuracy. Several of these policy approaches are therefore discussed and implemented in the ABM to evaluate their improvements.

9.1. Potential policies

To reduce the negative effects of data-driven RBI under uncertainty, three policy options are considered in this chapter. The effectiveness of each policy is evaluated based on the ability to reduce the number of missed high-risk inspections, with the best-performing policy being selected as the one that achieves the most reduction.

The first policy focuses on the (1) partly random inspection selection strategy. Within data-driven RBI, inspections are prioritised based on estimated risk scores. Due to uncertainty in these estimates, high-risk licence holders may not always be selected for inspection. During the interviews, it was found that there is communication within branches. This means that when, in a given scenario, inspections are mainly concentrated within one specific branch, other licence holders in that branch indirectly benefit from this through increased communication. As shown in Figure B.1, there is an average feedback loop between the probability of being inspected and the level of expertise. This effect is caused by inspections, but communication within branches also has a positive influence on expertise.

As a result, when inspections focus strongly on only a few branches, mostly in extreme circumstances, licence holders in other branches profit less from this increase in in-branch expertise. One way to reduce this effect is to introduce a share of random inspections in combination with risk-based selection. With this approach, all licence holders have a probability of being inspected and receive the communication advantage within their branch. Previous research indicates that combining risk-based and random inspections can help limit blind spots and bias (Ayaydın, 2023; Heerkens, 2023).

The second policy focuses on (2) improve data quality. Data-driven RBI depends on input data to estimate risk scores. Uncertainty or incompleteness in this data directly affects inspection prioritisation. Improving data quality can therefore reduce uncertainty in risk estimation and reduce the number of missed high-risk inspections (Ayaydın, 2023)

In this thesis, the model is based on a limited amount of data currently available at the ANVS, mainly due to a recent reorganisation. As a result, this policy focuses on improving the structural collection of data over time. During the interviews, several factors were identified as relevant for risk prediction and were included in the regression analysis. However, the regression was set up in a simple way. Latent variables were defined based on interview results and available data, but no statistically significant differences between these variables were tested. In addition, the regression analysis is based on a limited dataset. Although the direction of the regression coefficients was checked, there remains a high level of uncertainty. As a result, the input of the model contains uncertainty. Collecting more data over time and structurally registering licence holder characteristics would reduce this uncertainty. In combination with higher-quality datasets, more advanced prediction methods could be explored to further improve risk estimation.

The third policy focuses on (3) additional inspection capacity. With a fixed number of inspections per year, uncertainty in risk estimates has a larger influence on which licence holders are selected. Increasing inspection capacity allows more licence holders to be inspected each year and reduces the probability that high-risk licence holders are missed (Heerkens, 2023). This policy is straightforward, but difficult to implement. Inspectors require extensive training, and hiring additional inspectors is expensive. For this reason, it is tested whether the effect of increasing inspection capacity is large enough compared to the other policies.

Overall, these policy options are based on previous literature on inspection strategies, which highlights both the benefits of prioritising high-risk licence holders and the limitations of purely risk-based selection under uncertainty (Ayaydın, 2023; Heerkens, 2023).

9.2. Effect of policy implementations

Policy 0 - Base line uncertainty

If there are no policies implemented as seen in chapter 8, a strong change is visible around 2020. In both scenarios, this year corresponds to a sudden change in the number of missed-high-risk inspections (see figures 8.3 and 8.4) and in the worst-case risk paths (see figures 8.5 and 8.6). This follows a relative change in unemployment. Because unemployment is changing in these scenarios, and the model reacts relatively strongly to changes in this, this leads to an immediate response in the inspection selection.

In 2020, the number of missed high-risk inspections decreased strongly. In the years after this event, when unemployment decreases again, the number of missed high-risk inspections increases. This effect is stronger in scenario 3. Variations between seeds increase in the years after 2020. These worst-case paths show a similar response, with some licence holders that remain at higher risk levels for a longer period after 2020, meaning that the model does not immediately have them back into scope and keep being missed for inspections. Overall, the baseline results show limited stabilisation after a disruption in external conditions. This is an important notice to select the best policy.

Policy 1 - partly random inspection selection strategy

Inspections have indirect effects through communication within branches. When inspections spread across more branches, expertise increases more broadly, and average risk scores decrease. Because this communication effect decreases exponentially, when more licence holders from the same branch are inspected, a more diverse inspection distribution can reduce the overall risk more effectively than clustering inspections to a few branches.

With the random inspection policy, the model shows a clear response around 2020 in both scenarios (see figures 9.1 and 9.2). In the years after 2020, missed high-risk inspections increased again. Starting in 2022, this increase becomes strong in Scenario 1 and even stronger in Scenario 3. This indicates that the policy does not lead to a structural reduction in missed high-risk inspections over time.

The worst-case risk paths show a similar pattern (see figures E.1 and 8.6). After a temporary decrease around 2020, risk levels increased again for a subset of licence holders. In Scenario 3, these increases continue for a longer period, indicating the development of more extreme risk patterns in later years.

This behaviour can be explained by the interaction between random selection and changing external conditions. After the disruption around 2020, risk levels began to increase again for a part of the population. Because a fixed share of inspections is allocated randomly, fewer inspections remain available for licence holders with the highest estimated risk. As a result, high-risk licence holders are inspected less frequently in the years after 2020, leading to increasing risk levels and a growing number of missed high-risk inspections. This effect is strong under current uncertainty, as risk estimates are less stable and the impact of random selection becomes larger. As a result, changes in the lower part of the ranking more often affect which high-risk licence holders are selected, leading to an increase in missed high-risk inspections. This effect is stronger in Scenario 3, where uncertainty is higher and more licence holders reach high-risk levels.

Although the underlying idea of expertise allocation across branches is strong, the model's response to this behavioural pattern is limited. As also observed in the sensitivity analysis, the effect of improved expertise through communication is not strong enough to compensate for the reduced targeting of high-risk licence holders. As a result, this policy does not lead to improved outcomes under the tested conditions.

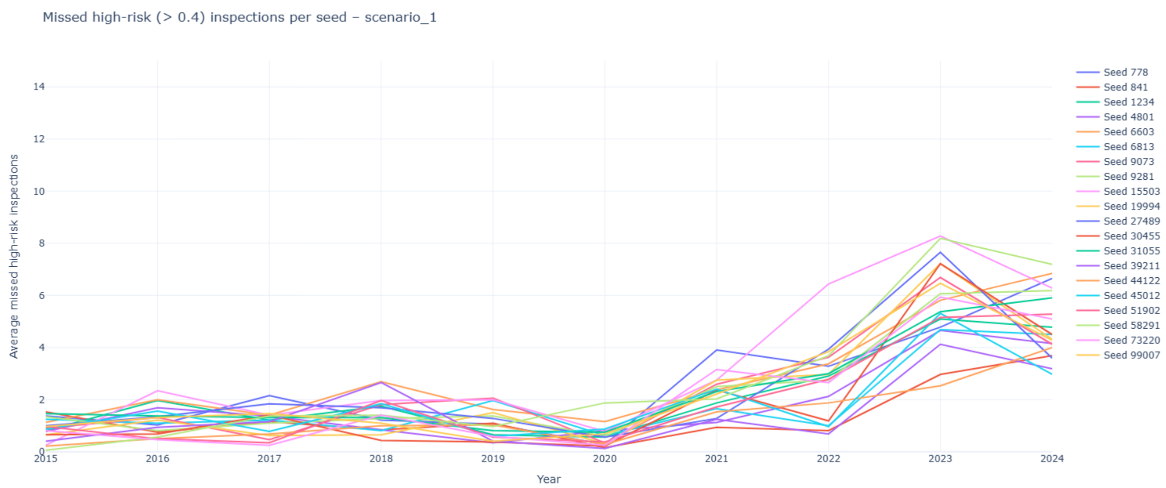

Figure 9.1: Missed high risk (>0.4) inspections - scenario 1 - partly random inspections

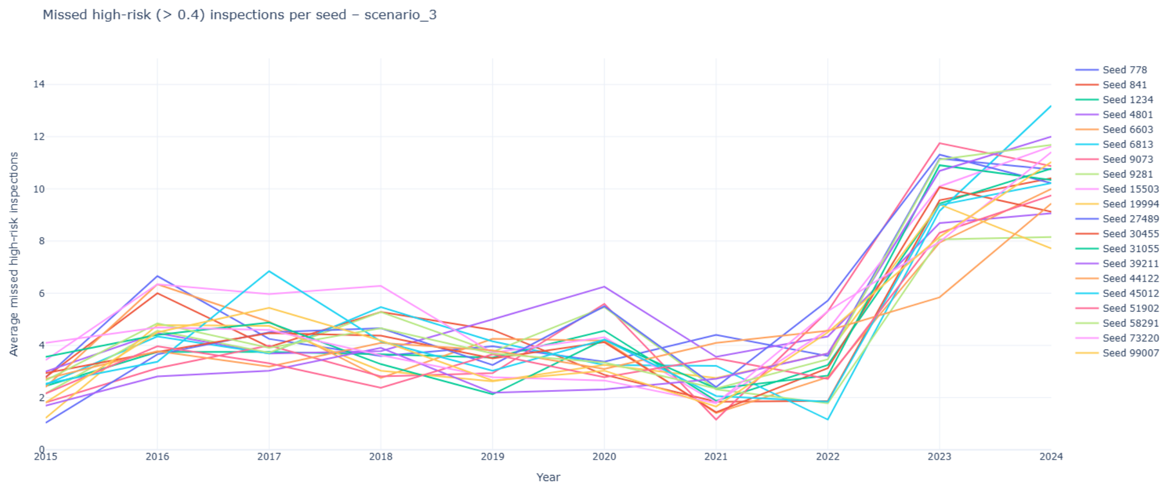

Figure 9.2: Missed high-risk (>0.4) inspections - scenario 3 - partly random inspections

Policy 2 - improved data quality

The data quality improvement policy directly affects the estimation of risk scores and, therefore, the inspection selection. By reducing uncertainty in the input data, the ranking of licence holders becomes more stable over time, which is especially beneficial when external conditions change.

Under this policy, the response to the unemployment trend disruption around 2020 remains visible. In both scenarios, the number of missed high-risk inspections decreases in 2020 (see figures 9.3 and 9.4). In the years after, missed high-risk inspections increase again, but this increase is less significant than in the baseline. Yearly variation is smaller compared to policy 0, especially in Scenario 3, indicating a more stable inspection selection in more extreme circumstances.

The worst-case risk paths also show a different development compared to the baseline (see figures E.3 and E.4). After 2020, risk levels slowly decrease, and the spread between licence holders' risk patterns becomes smaller. Ongoing high-risk patterns occur less frequently, which is especially visible in Scenario 3. This indicates that fewer licence holders remain in a high-risk state for long periods, as fewer high-risk licence holders are missed for inspection under this policy.

Compared to the baseline, improving data quality limits the strong reaction of the model to external changes. In the baseline model, fluctuations in unemployment lead to strong shifts in the inspection ranking when many licence holders approach high-risk levels. With improved data quality, these shifts are easier to stabilise. In scenario 3, external circumstances are worse, and the increase in missed high-risk inspections after 2020 is significantly smaller than in the baseline. Although this policy does not fully eliminate the impact of extreme external conditions, it reduces their long-term effects and results in a lower and more stable number of missed high-risk inspections overall.

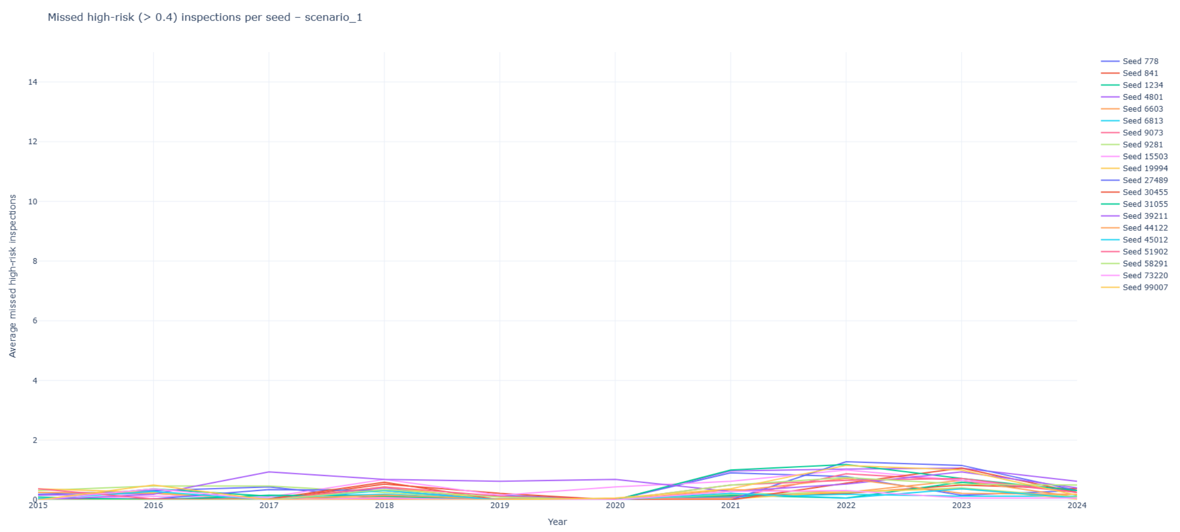

Figure 9.3: Missed high-risk (>0.4) inspections - scenario 1 - improved data quality

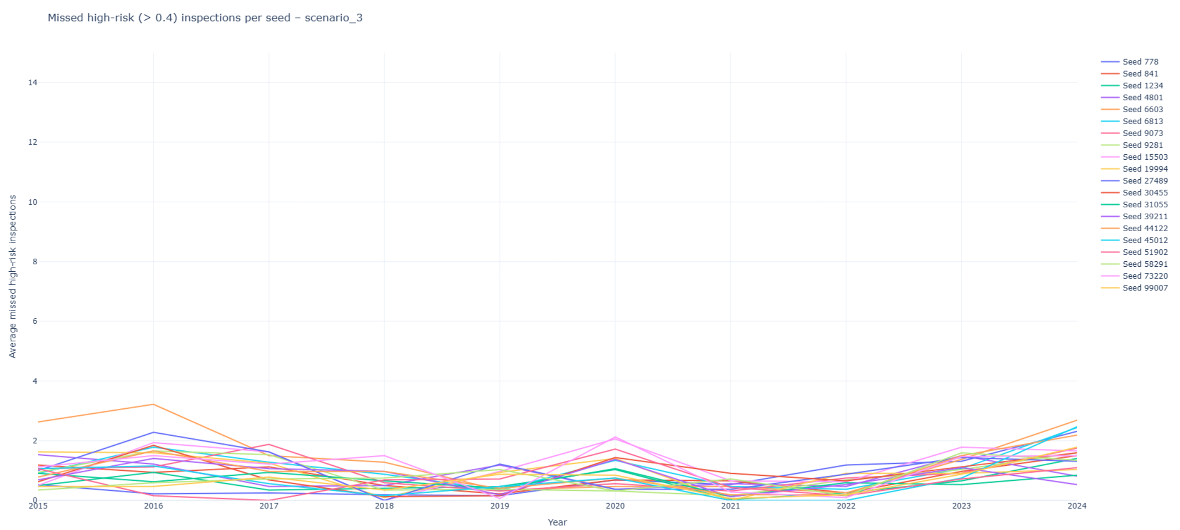

Figure 9.4: Missed high-risk (>0.4) inspections - scenario 3 - improved data quality

Policy 3 - increased inspection capacity

Increasing inspection capacity directly affects the coverage of licence holders that can be inspected each year. By inspecting more licence holders, the probability that high-risk licence holders are missed decreases, even under uncertainty in risk rankings.

With increased inspection capacity, the disruption around 2020 remains visible, but its impact on inspection outcomes is limited. In both scenarios, the number of missed high-risk inspections decreases in 2020 and remains low in the years after. Unlike the baseline and the random inspection policy, no strong increase in missed high-risk inspections can be observed in the years following 2020. Yearly variation is smaller, indicating a more stable inspection selection in different circumstances.

The worst-case risk paths show a short increase around 2020, followed by a relatively fast decline. Risk levels stabilise earlier compared to the other policies, and high-risk licence holders occur less frequently. This indicates that high-risk licence holders are inspected before risk levels can increase further.

This behaviour can be explained by the higher inspection coverage. When more inspections are available, fluctuations in the risk ranking have less influence on which licence holders are selected. Even when uncertainty increases, or more licence holders approach high-risk levels, sufficient inspection capacity remains to inspect the highest-risk licence holders. As a result, missed high-risk inspections are limited, and risk accumulation is prevented.

Compared to the policy 0, increased inspection capacity strongly reduces the reaction strength of the model to external changes, especially in Scenario 3. The increase in missed high-risk inspections observed after 2020 in the zero-policy is almost nonexistent. Among the tested policies, increased inspection capacity leads to a reduced number of missed high-risk inspections and provides a strong stabilisation of both inspection outcomes and worst-case risk paths under extreme external conditions. However, it is not significantly stronger than the improved data quality policy (2).

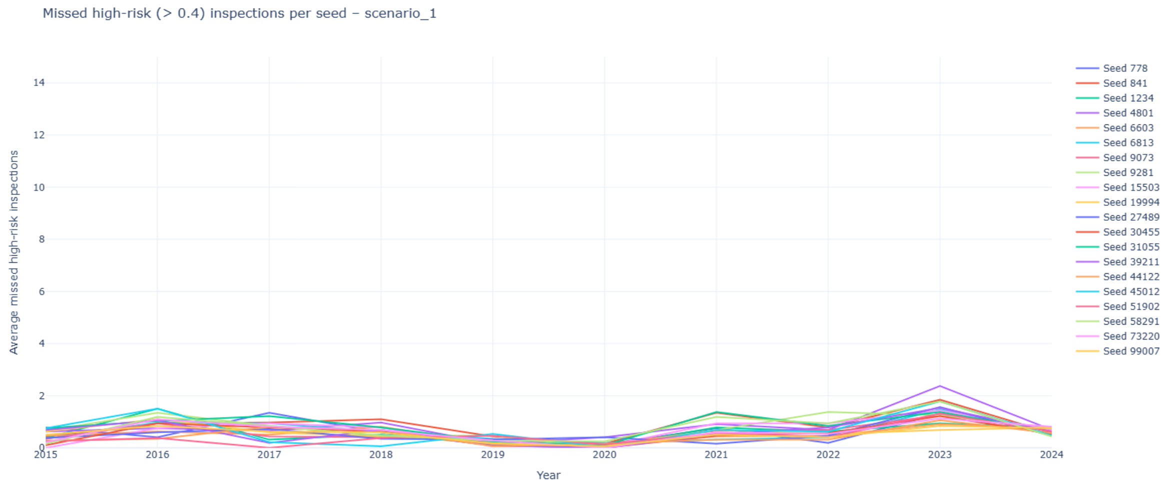

Figure 9.5: Missed high-risk (>0.4) inspections - scenario 1 - additional inspectors

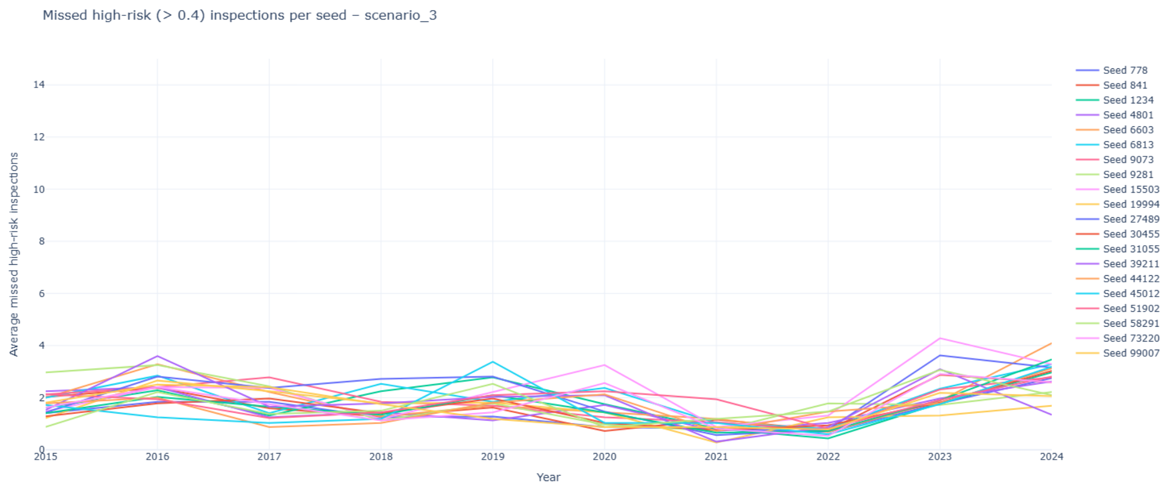

Figure 9.6: Missed high risk (>0.4) inspections - scenario 3 - additional inspectors

9.3. Conclusion - sub-question 5

In this chapter, three policy options were tested to reduce the negative effects of data-driven RBI under uncertainty. The policies differ in how they affect missed high-risk inspections, model stability, and feasibility.

The random inspection policy (1) is the least effective at reducing the number of missed high-risk inspections. Although spreading inspections across branches is ethically fair and easy to implement, it has an even more negative effect on inspection outcomes than if this policy were not implemented. This policy does not sufficiently prevent high-risk licence holders from being overlooked. However, it is straightforward to implement and does not require changes to current inspection practices. This policy could become more relevant if communication and expertise sharing improve in the future, both between inspectors and licence holders and among licence holders.

The data quality improvement policy (2) stimulates stability. Compared to the baseline, this policy reduces the strongest reactions to external changes and stabilises the development of continuous high-risk patterns, especially in more extreme scenarios. In addition, it supports more transparent and explainable inspection decisions because inspection prioritisation can be justified using a larger, more reliable dataset. However, the effects of this policy will be visible after a longer time period, since improving data quality requires consistent data collection over multiple years.

The policy of increasing inspection capacity (3) has a strong short-term impact within the system. Even a relatively small increase in capacity results in fewer missed high-risk inspections and more stable risk patterns per licence holder, even under extreme conditions. However, this policy requires additional inspection capacity and is therefore not in line with the goal of improving efficiency.

Based on these results, the advised approach is to implement a combination of policies. In the short term, a small increase in inspection capacity can support the safe implementation of data-driven RBI and reduce the number of missed high-risk inspections. At the same time, this provides opportunities to collect more and better data on inspections over multiple years. This is valuable for improving insight into the role of licence holders' characteristics in determining risk scores. However, simply increasing inspection capacity does not align with improving efficiency through risk-based inspection. It should therefore be considered a short-term measure. In this context, more junior inspectors can be deployed to conduct a larger number of relatively simple inspections, while more experienced inspectors are assigned to fewer but more complex and high-risk cases.

In the longer term, improved data quality can reduce uncertainty and improve the accuracy of inspection prioritisation. This is a long-term effect, since data gathering is a process that happens over a longer period of time. As the quality and reliability of the data increase, inspection capacity could slowly be reduced again, potentially to a level lower than the initial starting point while maintaining inspection accuracy. This implementation is preferable, as a 10% increase in data quality has an even stronger effect on inspection accuracy than a 10% increase in inspection capacity.

Random inspections remain an ethically fair policy, as they ensure that every licence holder has a probability of being inspected. This broadens inspection coverage and may contribute to maintaining compliance, as licence holders remain aware that inspection selection is always possible. However, under the current model assumptions, random inspections should not replace risk-based selection.

Overall, the results suggest that data-driven RBI can perform well under uncertainty, but that supporting policies are necessary. A phased approach, combining short-term stability with long-term improvement, is most appropriate within the scope of this study.

10

Discussion

This chapter discusses the limitations of the methods used in this study and reflects on the results and their interpretation. By reflecting on these limitations, this chapter clarifies what conclusions can and cannot be drawn from the thesis outcomes. This provides a realistic reflection of the current potential data-driven RBI approach, highlights aspects that require careful interpretation, and advises directions for further research. The aim of this research was to explore how uncertainty affects data-driven RBI selection and to assess the accuracy of inspection outcomes under different external scenarios and policy configurations. First, the approaches used to answer the sub-questions and the resulting findings are discussed. Secondly, the limitations of the study are reviewed. Thirdly, based on this discussion, the considerations for data-driven RBI in practice are described. Finally, the contributions of this research are discussed, and recommendations for further research are given.

10.1. Discussion of methods and results interpretation

10.1.1. Translating interviews into quantitative input variables

Translating interviews into quantitative input variables requires making tacit inspection knowledge explicit. Inspectors rely on experience-based judgement when evaluating risk, which is not directly measurable. Through interviews, this tacit knowledge is translated into risk factors by following the knowledge creation spiral described by Nonaka and Takeuchi (1995). The interviews are part of the externalisation phase of the knowledge creation spiral (see figure 4.2). As the concept of a spiral already suggests, knowledge creation is an ongoing and long-term process. In this thesis, externalisation was based on six interviews conducted within a relatively short time span. However, knowledge that needs to be externalised is not static, instead, it is expected to evolve as inspectors gain experience and share new insights over time. Therefore, iterations in the spiral are required to come closer to the underlying tacit knowledge, as interpretation involves bias. Given the limited duration of this thesis, a single iteration is explored. This limits the ability to use more of the potential of the knowledge spiral.

Despite these limitations, the factors that were carefully gained during the externalisation phase as saturation of interviews reached. These were grouped into latent variables to reduce model complexity and to represent qualitative aspects of risk, including the interconnectivity between quantitative risk factors. This interconnectivity was derived from the interviews and, therefore, required interpretation and simplification. This provides the base for the structured approach (see figure B.1.1) for using inspection knowledge as input in a quantitative modelling framework, however, interpretation can be done in different ways. Due to the saturation of the interviews, this subjectivity was reduced as much as possible, although some uncertainty remains that certain latent variables were not fully captured or clearly defined.

10.1.2. Use of agent-based modelling for risk-based inspection

The use of an agent-based model is suitable for studying risk-based inspection under uncertainty, as it allows for the representation of heterogeneous licence holders, dynamic risk evolution, and feedback effects between inspections and future risk profiles (Macal & North, 2010). Unlike static ranking or optimisation approaches, the ABM captures how inspection decisions influence system behaviour over multiple years, including changes in expertise, compliance, and inspection history.

The model was not intended as a predictive tool, but as an exploratory framework to analyse how inspection selection responds to uncertainty in a dynamic environment. The purpose of the model is to advise potential policy directions on the implementation of data-driven risk-based inspection within the ANVS. Rather than identifying an optimal inspection strategy, the model is used to test how different policy choices affect inspection outcomes under different assumptions and external conditions. This approach aligns with policy analysis under uncertainty, where models are used to assess the influence of model dynamics rather than to produce precise predictions (Bankes, 1993).

By modelling inspection and licence holders' behaviour, the ABM captures how uncertainty causes variability in inspection outcomes. Differences between simulation runs reveal how small changes in these behaviour assumptions affect inspection rankings and lead to different risk level paths of individual licence holders over time. This allowed inspection strategies and policy choices to be tested under uncertainty.

10.1.3. Representation of risk and uncertainty

Uncertainty is represented by applying a variation of −20% to +20% to the latent risk variable. This range was selected to introduce significant variability in risk estimates while maintaining the intended structure of the risk model. The chosen uncertainty range of −20% to +20% is a modelling assumption rather than an observed value. Because the data are limited and several risk variables are qualitative or latent, it is not possible to estimate the exact uncertainty in a statistically robust way. A more detailed approach could vary the uncertainty across different latent variables or licence holders to explore more specific patterns. However, the purpose of including uncertainty is not to define exact confidence intervals, but to test how inspection outcomes respond to variation in risk assessment. As long as uncertainty is applied in a systematic way, the main findings depend on the presence of uncertainty rather than on the exact variation of the chosen range.

10.2. Discussion of limitations

10.2.1. Data quality and behavioural assumptions

This study relies on simplified data due to the limited availability of detailed inspection and licence holder information over time. Some input data had to be categorised, which reduces detail and generalises differences between licence holders. The limited amount of historical data also limits the accuracy of regression-based risk predictions. In addition, there are no data available on the strength of behavioural patterns, meaning that their effects on yearly risk development are assumed and calibrated rather than directly measured. Although their influence was tested through a sensitivity analysis, the interpretation of absolute outcomes is limited, so results are mostly considered in relative terms.

In addition, the model assumes that inspections have a uniform effect among licence holders. Differences in inspection duration, inspection quality, severity of findings, and follow-up actions are not explicitly represented. While the model includes basic feedback effects such as inspection history and non-compliance accumulation, it does not capture adaptive or strategic behaviour by licence holders in response to inspection patterns. These assumptions simplify complex inspection dynamics to support modelling, but limit the level of behavioural detail that can be represented.

Finally, it should be noted that the uncertainty is defined per run for each latent variable, but not for each licence holder individually. This means that, realistically, there may be even more variations of uncertainty present, which are not captured in this model.

10.2.2. Model scope and capacity assumptions

First, the model includes only licence holders who have been inspected within the past five years. Licence holders that have not been inspected during this period are therefore excluded, even though they may still be relevant for inspection planning. This limitation may affect the regression analysis, as the effects of the characteristics of excluded licence holders are not represented and may therefore be interpreted differently than in the full inspection population.

Second, due to the limited amount of inspection data over time, there is little information on how risk scores change over time. In the model, risk development is therefore based partly on predicted values rather than observed changes. In practice, inspections generate new risk assessments each year, allowing regression models to be updated regularly. This would lead to a more accurate and up-to-date model over time, as the relative importance of licence holder characteristics and external variables can change over the years.

Third, in this model, external influences are included in a simplified way. Unemployment rates are used to represent changes in the external environment and to test how the model responds to such changes. The strong influence of unemployment in the results therefore mainly represents the model's reaction to environmental dynamics, rather than indicating that unemployment alone is the dominant driver or risk in practice. In reality, changes in unemployment are part of broader and more complex external developments that influence the risk scores of licence holders, this is not modelled in detail in this study. In addition, a linear regression approach is used, even though the relationship may be non-linear. Both low and high unemployment may result in higher risk, while medium levels reflect a more stable situation. Therefore, the role of external factors should be interpreted carefully, and further use of data-driven RBI should represent more detailed dynamics instead of single external indicators.

Finally, inspection capacity is modelled as a fixed number of inspections per year, assuming that all inspections require a similar amount of time and effort. In practice, inspections differ in duration and complexity. If, in a given year, mainly large or complex licence holders are prioritised, the actual number of inspections that can be carried out may be lower than assumed in the model. This simplification may result in an overestimation of feasible inspection capacity in certain scenarios.

10.3. Practical considerations for implementing data-driven RBI at the ANVS

This section discusses the practical considerations for implementing data-driven RBI within the ANVS, with the focus on data use, integration of inspector knowledge, and the organisation of the analytical support in the inspection process.

- The findings show that data-driven RBI can be initialised with information that is already available in ANVS's current systems. Qualitative risk assessments, including the scoring elements already used in inspection reports, provide a structured basis for quantification. As shown in this thesis, this existing structure can be translated into quantitative risk indicators without the requirement of fundamentally changing the inspection practice. Starting with this framework allows data-driven RBI to be introduced easily, and it reduces the risk that analytical support is separated from day-to-day inspection work.
- Due to the limited timespan of this thesis, the contribution of the spiral of knowledge creation is limited. The assumed uncertainty resulting from imperfect externalisation of knowledge and limited data availability does not directly lead to extreme biases or blind spots, as shown in Chapter 8 and Chapter 9. However, this imperfection should be managed by annually updating both the set of risk indicators and the relative importance of their historical values. Inspection outcomes and licence holder behaviour provide new information that can be used to refine the set of indicators taken into the analysis, this allows the risk assessment to strengthen over time.
- For data-driven RBI to be usable in the annual planning, inspection outcomes and external developments need to be captured in a consistent way across years. When updating risk scores, the new information should be combined with historical data, considering the relative importance

of data from each year.
- The output of every year's analysis should support the interpretation of the licence holders' prioritisation. In practice, this means presenting relative risk scores, including key contributing indicators and external developments. In this way, inspectors' professional judgement and analytical outputs can support each other in inspection planning. This can be facilitated through tools such as dashboards, which can be used during annual planning sessions.
- Based on the way data-driven RBI is applied in this thesis, the development and maintenance of analytical tools can be organised by analysts, as this mainly involves data processing, modelling, and updating indicators, and does not require in-depth knowledge of inspection practice. Inspectors' expertise remains important for interpreting analytical outputs and assessing whether the results align with observed risks in practice, while allowing inspectors to keep their focus on inspection activities and expertise. This way, analysts can be responsible for the data-driven RBI tool, while inspectors use the outputs as structured input in the annual planning process. This can be organised so that data-driven RBI supports annual inspection planning and is taken into account in decision-making, rather than functioning as a stand-alone tool.
- Given that inspection capacity is expensive and slow to expand, policy implementation should focus on using existing expertise more efficiently. More junior inspectors can be assigned to less complex inspections to gain experience, while more experienced inspectors focus on higher-risk or more complex licence holders. Currently, this is challenging to implement. An extension of data-driven RBI can support this by advising not only where to inspect but also what to inspect, to support this allocation strategy.
- Finally, the analysis shows that within the current model, the random inspection policy has a limited direct effect on inspection outcomes. This is partly due to the assumed strength of information-sharing behaviour in the model. If this behaviour were stronger, a broader range of risk-related information could be incorporated into inspection planning. This communication behaviour therefore has the potential to become more meaningful, in which case such a policy could have a stronger effect, even though this is not reflected in the current model results.

10.4. Contributions of this research

10.4.1. Scientific contributions

This research contributes to the scientific literature on risk-based inspection in three main ways. First, it contributes to the literature on tacit knowledge in regulatory oversight in combination with data-driven approaches in a risk-based inspection approach. Literature highlights that inspectors' tacit knowledge is essential for identifying risks in complex and high-impact sectors (Collins, 2010; Polanyi, 1966), which is also the case within the ANVS. At the same time, this literature highlights that externalisation of tacit knowledge is difficult, time-consuming, and holds cognitive bias. This research shows that externalising inspectors' experiential knowledge through multiple interviews, while acknowledging that this process should be iterative over more annual cycles and is therefore imperfect, still shows great potential for using this knowledge in a data-driven context.

The use of a data-driven approach in combination with the spiral of knowledge creation (Nonaka & Takeuchi, 1995) can be seen as a practical expansion of the spiral. This expansion is based on the idea of knowledge creation in risk-based inspection as two parallel and interacting paths. One path is based on socialisation and learning from interactions, and one is based on data-driven modelling. Both paths introduce bias, through cognitive judgement on the one hand and through data quality and modelling assumptions on the other, but their combination allows overlap in each other's bias to reduce blind spots in inspection prioritisation. This adds to existing work that treats expert judgement and data-driven RBI as a separate approach. By framing them this way, they become complementary within the iterative system.

Second, this research adds to the methods already described in the literature by integrating dynamic factors into an agent-based model that captures feedback loops, behavioural dynamics, and changing external conditions. While many RBI models rely on static predictions, this approach models inspection prioritisation with dynamics in several input values, such as inspection history, licence holder behaviour, and external changes that interact over time (Bankes, 1993; Macal & North, 2010). This shows the importance of considering feedback effects and interactions when assessing inspection accuracy.

Third, the policy analysis contributes to scientific understanding of how uncertainty affects risk-based inspection prioritisation and how different inspection strategies influence the effects. Rather than identifying a single optimal strategy, the results and system context show the importance of combining complementary policies when implementing data-driven inspection approaches. By using agent-based modelling, the research not only supports the identification of high-risk licence holders but also provides the first insights into the cause of the risks. When researched further, this has great potential to support better policy design.

10.4.2. Societal relevance

This research contributes to safer and more consistent regulatory supervision by supporting a more transparent inspection planning process. By making the assumptions and trade-offs behind inspection priorities explicit, data-driven RBI helps inspection authorities to better explain and justify their choices. This improves accountability and supports trust in regulatory decision-making. In addition, the ability to adjust inspection priorities to changing external conditions helps maintain effective supervision in years with increased uncertainty or limited inspection capacity, which is essential for safeguarding public safety and is in line with national and EU standards ((“Besluit basisveiligheidsnormen stralingsbescherming”, 2025; European Commission, 2025; International Atomic Energy Agency, 2006).

10.5. Recommendation for future research

Future research should focus on licence holders’ behavioural responses to inspections and on how these responses can be used to inform risk predictions. In practice, inspections can lead to different compliance responses, depending on inspection outcomes, enforcement actions, and previous inspection history. Modelling these differences more explicitly could improve the representation of risk development over time. For example, as mentioned during the interviews (Appendix A.4), it was suggested that the strength of ”expertise increases after inspections” differs between branches. This indicates that behavioural and learning effects may be sector-specific and could be important to include in more detailed inspection modelling in future research.

Future research should also focus on systematically internalising inspection outcomes into the annual inspection planning cycle. Every inspection results in an outcome that provides new information about the risk level of a licence holder. Future research should study how these inspection results can be incorporated into the risk score estimate after each inspection, so that predicted risk scores are continuously updated and remain accurate over time. Bayesian updating approaches provide a possible way to formalise this process by allowing risk estimates to be updated as new data become available while keeping information from earlier observations (Yuan et al., 2023). This includes studying how the regression model should be updated, how the relative importance of licence-holder characteristics changes, and how newly observed inspection data should be weighted relative to older data. With this additional data, more advanced methods for risk prediction can be explored as well.

In addition, external circumstances are included to represent different responses from the inspection system. In reality, inspection dynamics respond to a much wider range of influences than generalised economic conditions, innovation rates, or unemployment levels alone. Further research could extend the model by incorporating additional external variables, greater variation between sectors, and geographical differences. In addition, the external influences are not necessarily linear. The effect of every individual external factor could be more optimally used. These could be more closely aligned with their real influence by doing more research in previous literature, or this could be tested through more extensive statistical analyses than those used in this thesis.

These external factors, together with important licence holder characteristics, may not only be strong predictors of which cases are high risk, but also of what drives this high risk. This approach, therefore, has the potential to identify which latent variables contribute to high risk. If this is extended, it could result in a more targeted inspection strategy. However, further research is needed to again test accuracy, as focusing inspections on specific aspects of a licence holder also increases the risk of missing other relevant risks. In addition, it requires careful consideration of how such an approach can be incorporated into the current practice, since it would affect inspection outcomes differently from the current proposed approach, in which all licence holders within a branch receive the same inspection structure.

11

Conclusion

The potential of data-driven risk-based inspection (RBI) lies in the ability to support inspection planning by combining technical and social methods. In line with the national and European regulatory context, supervisory authorities are expected to prioritise inspections based on risk, where inspection intensity increases with the level of risk ("Besluit basisveiligheidsnormen stralingsbescherming", 2025; European Commission, 2025; International Atomic Energy Agency, 2006). For the ANVS, this expectation applies in a high-stakes regulatory context where mistakes have serious consequences. Within this context, data-driven RBI can support inspection planning by improving the accuracy and efficiency of inspection prioritisation. Currently, inspection planning at the ANVS is primarily based on professional judgement, with limited use of structured analytical support. The introduction of data-driven tools could take this to the next step in a socio-technical context, in which analytical methods are combined with expert judgement. The effectiveness of such tools, therefore, depends not only on their technical performance but also on how they are structurally used in yearly inspection planning.

Considering this context, this thesis examined how the current RBI approach at the ANVS can be improved using data-driven methods, and what risks arise when these methods are applied in a high-stakes and dynamic regulatory environment. A key challenge is translating inspectors' experience into quantitative factors that can be used in a data-driven inspection planning tool, while recognising that this experience is hard to fully capture in data. To address this, inspector interviews were conducted to externalise relevant factors and use them to predict licence holder risk scores. However, including these insights in a data-driven RBI approach introduces uncertainty, caused by current data availability and imperfect externalisation of tacit knowledge. In addition, external conditions can strongly influence licence holder risk scores, for example, through changes in economic conditions or radiation developments. The focus of this thesis was therefore on understanding how this uncertainty affects inspection selection, whether it leads to high-risk licence holders being missed in the annual planning under different dynamic circumstances, and how different policy choices influence the accuracy of data-driven inspection prioritisation.

Based on the methods used to implement quantitative approaches, this thesis shows that inspectors' experience can be externalised into a structured set of risk factors that can support inspection planning. Through interviews, relevant factors were identified that combine measurable characteristics, licence holder behaviour, and external developments. While the externalisation of these factors into quantitative input variables is imperfect, saturation in the interview responses indicates a shared view on a set of important risk factors. This provides a sufficient basis for their inclusion and shows that these factors give meaningful insight into differences in licence holder risk.

In addition, this indicates that data-driven RBI should be used in parallel with professional judgement, within the spiral of knowledge creation, rather than as a single decision-making instrument. Due to changes in the regulatory environment and licence holder behaviour, risk indicators should be redefined annually to ensure they continue to capture relevant risk factors. The value of data-driven RBI, therefore, lies in making risk considerations more explicit and comparable across licence holders by predicting quantitative risk scores for individual licence holders based on their characteristics and the system conditions at the time of prediction. At the same time, the approach structurally leaves room for inspectors' experience in the annual planning process.

By including these factors in an agent-based model, licence holder risk can change over time in response to inspections, inspection history, and external conditions. This allows inspection planning to be analysed as a dynamic process over time. At the same time, this dynamic setting introduces uncertainty. The results show that uncertainty leads to variation in inspection rankings. Under relatively stable external conditions, this variation has a limited effect on inspection outcomes and does not often result in missed high-risk licence holders. Under more challenging external conditions, uncertainty has a stronger effect and increases the likelihood of missed high-risk inspections. At the ANVS, such missed high-risk inspections mean an increased risk of radiation-related incidents. While uncertainty affects inspection accuracy, the worst-case analysis shows that risk remains bounded and does not develop into extreme risk levels. This can be explained by the inclusion of the licence holder's behavioural responses after inspections in the externalised risk factors.

The policy analysis shows that different policy measures influence inspection accuracy in different ways, and that no single measure fully reduces the effects of uncertainty. Increasing inspection capacity has a strong short-term effect, as it directly reduces missed high-risk inspections and stabilises risk patterns, but it requires additional resources and does not support improvements in inspection efficiency. Improvements in data quality reduce sensitivity to external changes and improve inspection performance, particularly in the longer term and have the strongest effect on missed high-risk inspections. However, this implementation requires a relatively longer period. Random inspections broaden inspection coverage and support fairness, but have a limited direct effect on reducing missed high-risk inspections within the current model. Their main added value lies in their potential indirect contribution to data quality by broadening the inspected sample.

Together, these findings indicate that a combination of policy measures is required to support the implementation of data-driven RBI. In the short term, relatively low-cost capacity measures, such as a strategic allocation of inspection expertise, can help reduce missed high-risk inspections while remaining feasible within the current organisational structure. Over time, increased inspection activity contributes to improved data quality, which strengthens the performance and stability of risk prediction models. Previous work already shows the importance of dynamic RBI with continuous adaptation of risk predictions (Bhatia et al., 2019). While high-frequency dynamic updates remain difficult to achieve within the ANVS context, yearly updates combined with a structured improvement cycle based on the spiral of knowledge creation (Nonaka & Takeuchi, 1995) can help the system to adapt to new insights over time. However, further research is needed on how this process should be structured in practice.

Most importantly, data-driven RBI can be introduced within existing inspection practices and can support annual inspection planning without changing the role of inspectors. By extending the information that is already available, data-driven RBI can contribute to accurate inspection prioritisation, provided that uncertainty is managed through iterative updates over time and analytical outputs are used as input for annual planning sessions rather than as stand-alone decision rules. Within the scope of this thesis, the results suggest that such an approach is feasible in practice, while remaining sensitive to the dynamics of radiation inspection.

Overall, the results indicate that data-driven RBI has clear potential, especially when combined with quality improvement policies, but its application remains highly context-specific. Beyond these policy results, this thesis contributes more generally by demonstrating a method for integrating data-driven models with expert knowledge, starting from the spiral of knowledge creation, which can be applied more generally. Although many risk factors identified in the literature are relevant, accurate application requires sector-specific externalisation of expert knowledge. Not all relevant risks can be captured or predicted using data-driven tools alone, as some factors, and especially combinations of factors, only become visible once they are explicitly identified. As a result, the performance of data-driven models depends on the relevance and continuous evaluation of their input variables. Within these limits, data-driven RBI supports inspection planning by structuring existing risk indicators more explicitly and comparably across licence holders. The approach does not replace professional judgement, but provides structured input that can be taken into account in the annual planning process. When updated over time using new inspection experiences, this allows inspection planning to improve accuracy.

Bibliography

Abbassi, R., Arzaghi, E., Yazdi, M., Aryai, V., Garaniya, V., & Rahnamayiezekavat, P. (2022). Risk-based and predictive maintenance planning of engineering infrastructure: Existing quantitative techniques and future directions. *Process Safety and Environmental Protection*, *165*, 776–790. https://doi.org/10.1016/j.psep.2022.07.046

ANVS. (2022). *Koersdocument*. Autoriteit Nucleaire Veiligheid en Stralingsbescherming. Den Haag.

ANVS. (2023). *Toezicht- en interventiestrategie 2023*. Autoriteit Nucleaire Veiligheid en Stralingsbescherming. https://www.anvs.nl

ANVS. (2025a). Internal trusted dataset for regression analysis [Confidential internal dataset; not publicly accessible].

ANVS. (2025b). *Jaarplan 2025: Voorbereiden, verbinden en verbeteren*. Autoriteit Nucleaire Veiligheid en Stralingsbescherming. Den Haag. https://www.anvs.nl

Awodi, N. J., Liu, Y.-K., Ayodeji, A., & Adibeli, J. O. (2021). Expert judgement-based risk factor identification and analysis for an effective nuclear decommissioning risk assessment modeling. *Progress in Nuclear Energy*, *136*, 103733. https://doi.org/10.1016/j.pnucene.2021.103733

Ayaydın, S. N. (2023, August). *Towards an assessment framework for inspection strategies: Combining risk-based and random approaches in food safety inspections* [Master's thesis]. Delft University of Technology.

Bankes, S. C. (1993). Exploratory modeling for policy analysis. *Operations Research*, *41*(3), 435–449.

Bayramova, A., Edwards, D. J., Roberts, C., & Rillie, I. (2024). Uncovering the genome of leading indicators from lagging indicators and normative documents: A proof-of-concept study. *Journal of Safety Research*, *91*, 230–244. https://doi.org/10.1016/j.jsr.2024.08.015

Besluit basisveiligheidsnormen stralingsbescherming [Regulation (ministeriële regeling), Netherlands. Geldend per 1 januari 2025]. (2025). *Rijksoverheid*.

Bhatia, K., Khan, F., Patel, H., & Abbassi, R. (2019). Dynamic risk-based inspection methodology. *Journal of Loss Prevention in the Process Industries*, *62*, 103974. https://doi.org/10.1016/j.jlp.2019.103974

Black, J., & Baldwin, R. (2010). Really responsive risk-based regulation. *Law & Policy*, *32*(2), 181–213. https://doi.org/10.1111/j.1467-9930.2010.00318.x

Campari, A., Vianello, C., Ustolin, F., Alvaro, A., & Paltrinieri, N. (2024). Machine learning-aided risk-based inspection strategy for hydrogen technologies. *Process Safety and Environmental Protection*. https://doi.org/10.1016/j.psep.2024.09.031

CBS. (2025a). *Economische groei (bbp)* [Interactief dashboard]. Retrieved January 3, 2026, from https://www.cbs.nl/nl-nl/visualisaties/dashboard-economie/bbp/economische-groei

CBS. (2025b). *Innovatie; kerncijfers* [CBS StatLine tabel 80066NED]. Retrieved January 3, 2026, from https://www.cbs.nl/nl-nl/cijfers/detail/80066ned

CBS. (2025c). *Werklozen* [Interactief dashboard arbeidsmarkt]. Retrieved January 3, 2026, from https://www.cbs.nl/nl-nl/visualisaties/dashboard-arbeidsmarkt/werklozen

Coffeng, T. (2022). *Bias in supervision: A social psychological perspective on regulatory decision-making* [Doctoral dissertation, Utrecht University] [Doctoral dissertation, Utrecht University]. https://doi.org/10.33540/1062

Collins, H. (2010). *Tacit and explicit knowledge*. University of Chicago Press.

Company.info. (2025). *Company.info – bedrijfsinformatie en jaarcijfers*. Company.info. Retrieved January 3, 2026, from https://www.company.info

Dyer, A. S., DeLaughter, J., Thomas, J., Krishna, M., Long, J., & Landon, M. D. (2022). Applied machine learning model comparison: Predicting offshore platform integrity with gradient boosting algorithms and neural networks. *Marine Structures*, *83*, 103152. https://doi.org/10.1016/j.marstruc.2022.103152

Edmonds, B., & Meyer, R. (2017). *Simulating social complexity*. Springer.

European Commission. (2025). *Radiation protection legislation*. European Commission. Retrieved January 4, 2026, from https://energy.ec.europa.eu/topics/nuclear-energy/radiation-protection/radiation-protection-legislation_en

Heerkens, L. (2023). *Inspection strategies under risk and uncertainty* [Master's thesis]. Delft University of Technology.

Hobé, R., van Asselt, E., van den Heuvel, L., Hoek-van den Hil, E. F., & van der Fels-Klerx, H. (2023). Methodology for risk-based monitoring of contaminants in food – a case study in cereals and fish. *Food Research International*, *168*, 112791. https://doi.org/10.1016/j.foodres.2023.112791

Huang, D., Li, W., Chang, X., & Tan, Y. (2019). Key factors identification and risk assessment for the stability of deep surrounding rock in coal roadway. *International Journal of Environmental Research and Public Health*, *16*(15), 2802. https://doi.org/10.3390/ijerph16152802

International Atomic Energy Agency. (2006). *Fundamental safety principles: Safety fundamentals* [Jointly sponsored by Euratom, FAO, IAEA, ILO, IMO, OECD/NEA, PAHO, UNEP and WHO]. https://doi.org/10.61092/iaea.hmxn-vw0a

Johannesson, P., & Perjons, E. (1996). *An introduction to design science*. Springer. https://link.springer.com/book/10.1007/978-3-319-10632-8

Kiviniemi, K., Salmivaara, L., Vainio, A., & Lundén, J. (2025). Food control strategies to support and enforce food business operators that repeatedly violate food safety legislation. *Food Control*, 111042. https://doi.org/10.1016/j.foodcont.2024.111042

Li, T., Langforth, S., Langkabel, N., Sotiraki, S., Anastasiadou, S., Nesbakken, T., & Meemken, D. (2023). Implementation of harmonised epidemiological indicators (heis) for pigs – a europe-wide online survey. *Food Control*, *153*, 109954. https://doi.org/10.1016/j.foodcont.2023.109954

Macal, C. M., & North, M. J. (2010). Tutorial on agent-based modelling and simulation. *Journal of Simulation*, *4*(3), 151–162.

Masad, D., & Kazil, J. (2015). Mesa: An agent-based modeling framework. *Proceedings of the 14th Python in Science Conference (SciPy)*, 53–60.

May, I. V., & Nikiforova, N. V. (2019). Methodological approaches to optimization of the laboratory control over product safety within risk-based surveillance framework. *Hygiene and Sanitation*, *98*(2), 205–213. https://doi.org/10.18821/0016-9900-2019-98-2-205-213

Mukelabai, M. D., & Barbour, E. R. (2024). Modeling and optimization of renewable hydrogen systems: A systematic methodological review and machine learning integration. *Energy and AI*. https://doi.org/10.1016/j.egyai.2024.100455

Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How japanese companies create the dynamics of innovation*. Oxford University Press.

Papadopoulos, D., & Benardos, A. (2021). Enhancing machine learning algorithms to assess rock burst phenomena. *Geotechnical and Geological Engineering*, *39*, 5787–5809. https://doi.org/10.1007/s10706-021-01867-z

Pasman, H., & Rogers, W. (2020). How to treat expert judgment? with certainty it contains uncertainty! *Journal of Loss Prevention in the Process Industries*, *66*, 104200. https://doi.org/10.1016/j.jlp.2020.104200

Polanyi, M. (1966). *The tacit dimension*. Routledge & Kegan Paul.

Rachman, A., & Ratnayake, R. M. C. (2019). Machine learning approach for risk-based inspection screening assessment. *Reliability Engineering & System Safety*, *185*, 518–532. https://doi.org/10.1016/j.ress.2019.02.008

Racicot, M., Zanabria, R., Paradis, M.-È., Gaucher, M.-È., Arsenault, J., Leroux, A., & Quessy, S. (2019). Identification of risk factors to be considered for food establishments' risk assessment models. *Microbial Risk Analysis*, *11*, 1–10. https://doi.org/10.1016/j.mran.2018.01.004

Rijksdienst voor Ondernemend Nederland. (2025). *Tvl: Mkb en grote ondernemingen*. Rijksdienst voor Ondernemend Nederland. Retrieved January 3, 2026, from https://www.rvo.nl/subsidies-financiering/tvl/mkb-grote-onderneming

Spirtes, P. (2001). Latent structure and causal variables. In N. J. Smelser & P. B. Baltes (Eds.), *International encyclopedia of the social & behavioral sciences* (pp. 8395–8400). Elsevier. https://doi.org/10.1016/B0-08-043076-7/00429-0 van Asselt, E. D., Hoffmans, Y., Hoek-van den Hil, E. F., & van der Fels-Klerx, H. J. (2021). Methods to perform risk-based inspections of food companies. *Journal of Food Science*, *86*(12), 5078–5086. https://doi.org/10.1111/1750-3841.15978

Wooldridge, J. M. (2010). *Introductory econometrics: A modern approach* (5th ed.). South-Western Cengage Learning.

Woreta, T. M., Worku, A. F., Tenagashaw, M. W., Yemata, T. A., Mamo, F. T., & Damtie, D. G. (2025). Influence of knowledge and attitudes on food safety perceptions and behavioral intentions among food business operators in water bottling factories: A pls-sem model. *SN Applied Sciences*, *7*(2). https://doi.org/10.1007/s42452-025-06542-z

World Nuclear Association. (2024). *Radiation and health effects – world nuclear association*. https://world-nuclear.org/information-library/safety-and-security/radiation-and-health/radiation-and-health-effects

Yazdi, M., Yakhchali, S., & Zarei, E. (2019). A methodology for enhancing the reliability of expert system applications in probabilistic risk assessment. *Journal of Loss Prevention in the Process Industries*, *58*, 51–59. https://doi.org/10.1016/j.jlp.2019.01.005

Yuan, S., Reniers, G., & Yang, M. (2023). Dynamic-risk-informed safety barrier management: An application to cost-effective barrier optimization based on data from multiple sources. *Journal of Loss Prevention in the Process Industries*, *83*, 105034. https://doi.org/10.1016/j.jlp.2023.105034

Zanabria, R., Racicot, M., Cormier, M., Arsenault, J., Ferrouillet, C., Letellier, A., & Quessy, S. (2018). Selection of risk factors to be included in the canadian food inspection agency risk assessment inspection model for food establishments. *Food Microbiology*, *75*, 72–81. https://doi.org/10.1016/j.fm.2017.09.019

Zhen, X., Ning, Y., Du, W., Huang, Y., & Vinnem, J. E. (2023). An interpretable and augmented machine-learning approach for causation analysis of major accident risk indicators in the offshore petroleum industry. *Process Safety and Environmental Protection*, *173*, 922–933. https://doi.org/10.1016/j.psep.2023.03.063

A

Interviews chapter 5 - sub-question 1

A.1. Interview 1

The interviewee discussed the various internal and external factors that influence inspection outcomes and overall safety performance.

One of the internal factors highlighted was the importance of communication within the organisation. In some companies, the presence of multiple languages can hinder effective communication and lead to misunderstandings regarding safety procedures. The organisation's history of violation reports was also considered a crucial indicator, as repeated issues often reflect deeper structural weaknesses.

The interviewee noted that many larger organisations currently struggle to attract and maintain qualified personnel, resulting in a shortage of in-house expertise essential for maintaining radiation safety. Frequent personnel changes were also mentioned as a factor that can reduce consistency and the overall level of safety awareness. The type of organisation also influences safety culture: organisations operating in fast-paced, production-oriented environments tend to take more risks, whereas those with a more careful and precise working culture usually demonstrate higher safety standards.

Interestingly, the interviewee observed that companies with more complex activities usually have better organised safety systems, as complexity requires clearer procedures and stronger internal controls. Therefore, the inspector advised against using company size as a direct indicator of risk, since the relationship between size and safety performance is not straightforward. Instead, the focus should be on organisational complexity and how safety is managed within that structure.

The behaviour of employees during inspections was also identified as an important indicator. The attitude and openness of company representatives can reveal a great deal about an organisation's underlying safety culture. Furthermore, the nature of operations plays a role: continuous operations involving frequent shift changes present different safety challenges to non-continuous operations involving more stable teams.

Several patterns were also noted. For example, safety performance often worsened during holiday periods when responsibilities were harder to distribute, and temporary replacements were unfamiliar with safety routines. In addition, after an inspection, companies usually demonstrate improved compliance and awareness, but this renewed focus tends to fade over time.

Regarding external factors, the economic situation was seen as a notable influence. During economic recessions, companies may pay less attention to safety and invest fewer resources in it. The interviewee also pointed out that technological innovation can introduce uncertainty, as organisations need time to adapt to new processes and define how safety should be maintained.

Finally, the interviewee highlighted that companies that go long periods without inspection often become careless, which can lead to a subtle decline in safety standards. Regular inspections are therefore essential to sustain awareness and guarantee ongoing compliance.

A.2. Interview 2

The interviewee highlighted several internal and external factors influencing inspection outcomes and risk assessments within the current regulatory context.

Documentation was identified as a key aspect, with documentation such as radiation alarm reports, transport documents and Annexe 7 forms being particularly important. Additionally, the quality of administrative processes and the level of radiation expertise within an organisation were considered important indicators of compliance and overall safety performance.

When a change in management occurs, inspectors aim to assess the organisation with a fresh and unbiased perspective regarding any previous violations or safety notices. However, frequent staff turnover is generally considered an indication of lower safety levels. One potential indicator of this could be the number of publicly available job vacancies. The size of an organisation could also affect its risk profile: larger organisations tend to be more complex, which increases the likelihood of safety issues arising, whereas smaller organisations often rely on external radiation experts. This dependency on external expertise can hinder communication and reduce the effectiveness of internal safety management.

The interview also revealed certain patterns across sectors. Companies sometimes inform one another about inspection visits after they have taken place, which may influence preparedness for future inspections. In smaller organisations that share the same external radiation expert, who is often engaged for only a few hours per year, inspection information may be indirectly shared, reducing the element of surprise and potentially affecting inspection outcomes.

Among the external factors, digitalisation was identified as having a significant influence. Organisations with an older workforce often struggle to maintain digital documentation, resulting in incomplete or disorganised records.

Finally, the interviewee highlighted that, although the Risk-Based Inspection (RBI) approach prioritises high-risk cases appropriately, it is equally important to continue inspecting low-risk facilities. Maintaining this balance ensures a consistent safety culture and helps to detect early warning signals that might otherwise be overlooked in lower-risk organisations.

A.3. Interview 3

The interviewee discussed the various internal and external factors that influence inspection outcomes and organisational risk levels.

Regarding the internal factors, the availability and quality of documentation were identified as key indicators. The organisation's behaviour when providing documentation was also considered highly relevant, for example, delays in submitting documents or unwillingness to share them could suggest underlying compliance issues or a less transparent safety culture. Repetition of incident notifications was also mentioned as an important signal of risk, as recurring reports may suggest that structural problems remain unresolved.

The interviewee noted that organisations that rely on an external radiation protection officer often experience weaker internal communication, which can lead to higher safety risks. A company's sectoral background also plays an important role. In certain sectors, such as vets, employees often have a basic understanding of radiation safety due to formal training or education, whereas in the food industry, for example, there is typically limited prior knowledge about radiation.

Changes in turnover or revenue were seen as potential indicators of future risk, particularly if they suggest possible financial instability or bankruptcy. Similarly, organisational changes, such as renovations or site modifications, may influence safety conditions and could be monitored through public data sources, such as the Dutch Land Registry (Kadaster).

The interviewee also mentioned several patterns that are relevant to inspection planning. It is considered important to have visibility within a sector, as licence holders often communicate with each other through professional associations. While this exchange of information is not necessarily negative, it is important for inspectors to be aware of it, as it can influence behaviour and provide opportunities for more efficient inspection planning.

Regarding external factors, the interviewee mainly highlighted economic conditions at the company level, such as turnover trends and potential bankruptcy risks.

A.4. Interview 4

According to the interviewee, inspectors first assess the risk associated with the radiation sources themselves. This includes the type of application, whether the source is used at multiple or changing locations, and how many people are typically present near the source, since this increases the potential severity of an incident.

The organisational context also plays a significant role. Larger organisations generally have more well-trained staff and stronger in-house expertise in radiation safety because they can invest more resources in these areas. However, this factor has two sides, as larger companies may also operate with more radiation sources and more complex systems, which could introduce additional risks. The quality of documentation is another strong predictor of inspection outcomes. Careless or incomplete documentation often reflects a broader attitude towards safety, with inaccuracies in administrative records frequently mirroring careless handling of radiation sources. Furthermore, personnel changes are checked before inspection planning as they can indicate organisational instability. The interviewee noted that this factor has not yet been analysed over time, but it could be valuable for identifying why certain violations tend to occur at specific times. Therefore, incorporating staff turnover over time may improve future risk prediction.

The interviewee also discussed recurring patterns in compliance behaviour. For example, after an inspection, companies typically show temporary improvements in compliance, handling radiation sources more carefully. However, the strength of this effect varies across sectors, some stay motivated to comply, while others resume normal operations as quickly as possible.

Information-sharing within branches is identified as another important pattern. Organisations often communicate with each other after inspections, meaning that companies inspected later tend to have already corrected issues identified earlier in the process. This form of peer learning can positively influence compliance across the sector. The ANVS already makes use of this effect by distributing fact sheets containing key findings and lessons learned following inspections. These are shared through umbrella organisations or directly with the entire branch when relevant, helping to raise awareness and improve safety.

Considering the external factors, specific incidents were noted as important triggers for increased awareness and stricter compliance among similar companies. However, sectors or licence holders that have not been inspected for several years tend to become lax, resulting in fewer reported incidents or violations. This highlights the importance of maintaining a regular inspection schedule as a preventive measure in sustaining a safety culture.

A.5. Interview 5

The interviewee highlighted that a company's core business is an important predictor of their risk profile. Whether radiation use is part of the organisation's primary activities or just a supporting function within broader commercial operations is important. In the second case, radiation safety tends to receive less attention over time.

Another key factor is the continuity of the internal radiation protection expert. While this tends to be more stable in larger organisations, the interviewee clarified that this is not directly related to company size, but rather to the scope of the licence, including the number and strength of radiation sources, and the extent of radiation-related activities.

The quality of the documentation provided prior to an inspection was also identified as a reliable indicator of compliance. Incomplete, delayed or disorganised documentation often correlates with a lower level of commitment to safety, in contrast to detailed, well-structured documentation, which signals higher compliance and awareness.

In addition, several organisational and financial characteristics were identified as being relevant to risk assessment. Company profiles, such as financial stability, the type of safety services used and the extent of import activities, can all influence overall risk. Companies that rely heavily on imports may face increased risks due to inconsistent quality control of incoming materials, as demonstrated by past issues in certain retail sectors.

Finally, a company's revenue and profit focus was identified as an indicator of safety culture. Organisations that are highly profit-driven may invest fewer resources in radiation safety, while those with a stable financial position and less commercial pressure are more likely to keep consistent compliance.

Also, several patterns that influence safety performance and inspection results were discussed. Safety levels often decline during the summer due to staff shortages and holiday schedules, making it harder to maintain continuity and oversight.

The interviewee mentioned that the history of incident notifications is taken into account when prioritising inspections. During the preparation phase, inspectors typically review two years of reports, using their professional judgement to detect emerging risk trends.

When looking at the external factors, the introduction of new radiation-related technologies can also affect compliance. The early stages of new technologies or applications are often associated with greater uncertainty, as organisations may not yet have developed the necessary expertise or procedures to manage them safely.

A.6. Interview 6

According to the interviewee, the first step in preparing for an inspection is always to check the general license and the specific conditions included in it. The size of the organisation of the license holder and the number of employees are also taken into account, as these factors determine the structure of the responsibility distribution. Additional important aspects identified were the age of the license and how recently the company has had contact with the ANVS. Organisations that maintain regular contact with the authority tend to perform better, while long periods without communication make it harder to assess their current level of compliance or awareness.

The interviewee also considers whether radiation is used as a core or side activity within the company. This affects how much priority radiation safety receives within the organisation. Additionally, previous inspection reports, as well as any other notifications or incidents, are reviewed to identify recurring weaknesses.

The size of an organisation influences the risk score. In very large companies, radiation-related activities may become lost within broader operations, resulting in reduced oversight. However, in smaller companies, limited employee and backup capacity can lead to a lack of expertise or gaps in safety management.

The interviewee noted that organisational culture is difficult to quantify, however, it remains relevant. During an inspection, a culture of openness and transparency, where inspectors can speak freely with multiple employees, is generally associated with higher safety. In contrast, limited communication to management could indicate lower internal awareness or limited knowledge of daily radiation practices.

patterns: The interviewee observed that companies that have not been inspected for a long time often show more chaos or confusion. This suggests that regular inspections increase compliance. Communication quality also varies significantly between sectors, some show strong communication towards other organisations in the same branch, while others do not. In certain highly regulated sectors, organisations are used to complying with multiple regulatory frameworks, which reduces the additional risk when ANVS conducts inspections.

The interviewee also highlighted the importance of not neglecting smaller or less well-known license holders. As little is known about these, it is hard to know if they have a high risk or not.

External factors: In terms of external influences, the interviewee discussed technological innovation as a key factor. There is an ongoing shift from radiation sources to X-ray devices, which offer easier control and a lower long-term risk. Similarly, some applications are being replaced by non-radiation alternatives, such as ultrasound imaging. Although radiation sources remain commonly used in hospitals, this technological transition is particularly notable in the industrial sector, where innovation is slowly reducing reliance on radioactive materials.

B

Appendix chapter 6 - sub-question 2

B.1. Latent variables

B.1.1. Causal Relation Diagram

Figure B.1: Causal relationship diagram

B.2. ABM model design

B.2.1. Pseudocode of the agent-based inspection model

INPUT: license-holder data, yearly external factors, inspection capacity N, latent-scale settings, optional fixed inspection lists.

OUTPUT: yearly KPIs, inspection history, risk rankings, final agent states.

INITIALISE

- Create one agent per license holder.
- For each agent set:
  - `years_since_last_inspection = 0-3`
  - `inspection_count = 0`
  - `branch_nr_inspected = 0`

FOR each simulation year

- Read external factors for the year.
- **Risk assessment**
  - FOR each agent:
    - Compute latent risk components: `FS`, `CB`, `CT`, `EX`, `LS`.
    - Combine components into `total_risk`.
    - Apply corrections based on inspection history and branch activity.
- **Ranking and selection**
  - Rank agents by `total_risk`.
  - Select top N agents for inspection.
  - If a fixed selection exists, use it instead.
- **Inspection and feedback**
  - FOR inspected agents:
    - Increase `inspection_count`.
    - Reset `years_since_last_inspection`.
  - FOR non-inspected agents:
    - Increase `years_since_last_inspection`.
  - Update `branch_nr_inspected` for all agents.
- **Store outputs**
  - Save yearly statistics, inspection log, and risk ranking.

Full python code can be found on: https://github.com/annabelverspeek1/abm-data-driven-risk-based-inspection

B.2.2. Variables in inspection model

Variable	Influenced by	Description
Fixed variables of licence holders		
Licence holder characteristics	-	Static characteristics of the licence holder, such as sector and type of activities.
Yearly changing variables of licence holders		
`FS, CB, CT, EX, LS`	Inspections, incidents, external developments	Latent risk-related variables representing safety culture, compliance behaviour, complexity, experience, and learning effects.
`Years_since_last_inspection`	Inspection history	Number of years since the licence holder was last inspected.
`Branch_nr_inspected`	Inspection allocation in previous year	Number of inspections performed in the same branch in the previous year, representing indirect learning and expertise spill-over for the licence holder.
Fixed variables of the inspection model		
`Run years`	2015 - 2025	Total number of years simulated in a single model run.
`Inspection capacity`	20	Total number of inspectors available in a run year
Yearly changing variables of the inspection model		
`Economic growth`	External scenario assumptions	Annual economic growth rate affecting overall system dynamics.
`Unemployment rate`	External scenario assumptions	Annual unemployment rate influencing organisational behaviour and risk.
`Innovation rate`	External scenario assumptions	Rate of technological and organisational innovation affecting licence holders.

Table B.1: Overview of variables in the inspection model.

B.3. Regression coefficients

This appendix documents the datasets used in the regression analysis presented in Chapter 6. The analysis is based on a combination of internal data provided by ANVS (2025a) and external publicly available datasets. Due to confidentiality requirements, the full internal dataset cannot be disclosed. This appendix, therefore, provides an overview of the data sources and variables used in the regression analysis.

B.3.1. Financial stability

The graph and regression analysis are based on the number of employees within each organisation. On the basis of this information, organisations are classified into size categories following the definitions provided by Rijksdienst voor Ondernemend Nederland (2025). Data on the number of employees and whether a licence holder operates on a profit-driven basis were obtained from Company.info (2025). GDP growth was gathered from data provided by CBS (2025a).

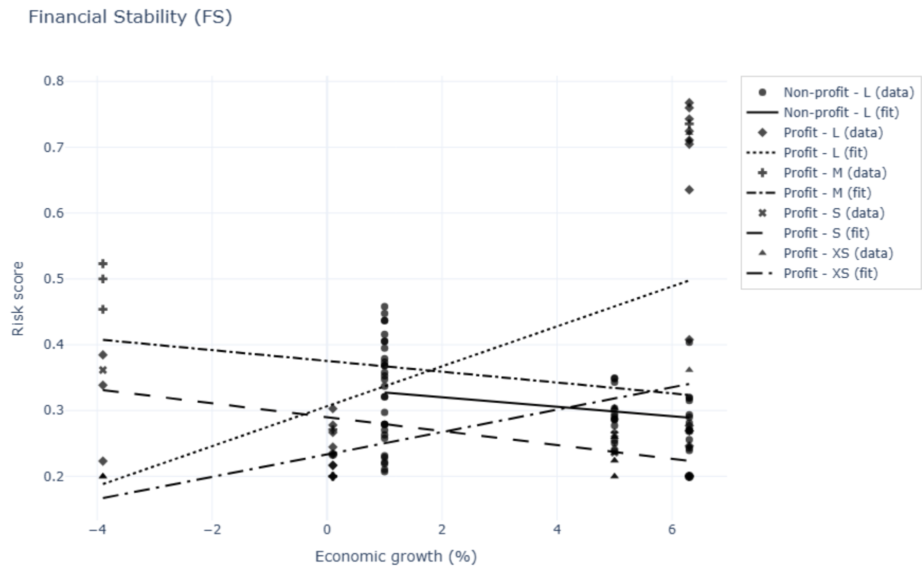

Figure B.2: Regression Financial Stability

B.3.1.1. Coefficients per group

Table B.2: Regression results by group

profit driven	organisation size	alpha	beta	n_obs
0	XS	-	-	0
0	S	-	-	0
0	M	-	-	1
0	L	0.303	-0.007	55
1	XS	0.233	0.017	18
1	S	0.290	-0.0106	7
1	M	0.375	0.008	9
1	L	0.307	0.030	25

B.3.2. compliance Behaviour

The regression of compliance behaviour is based on historical inspection data. Information on compliance outcomes and notice history was obtained from internal data provided by the ANVS (2025a).

B.3.2.1. Graph

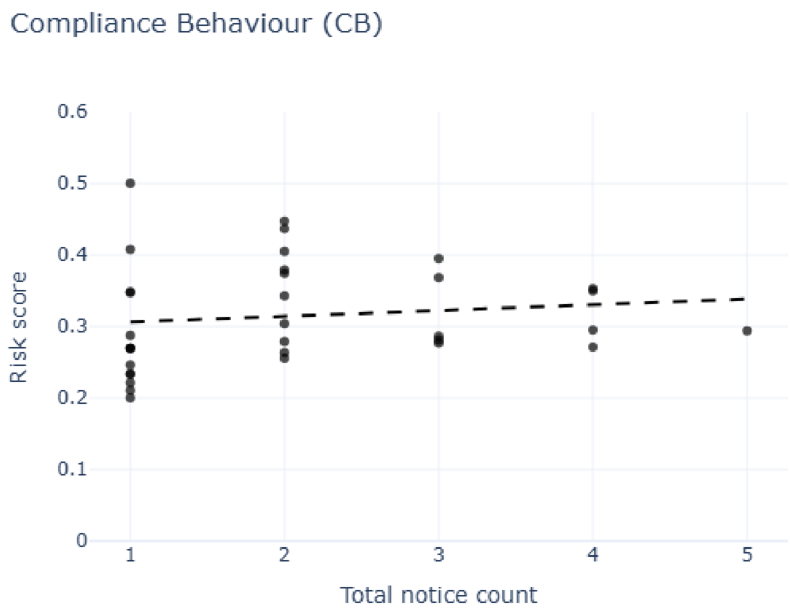

Figure B.3: Regression Compliance Behaviour

B.3.2.2. Coefficients per group

Table B.3: Regression notice history

alpa	beta_notice history	n_obs
0.298135	0.008	34

B.3.3. Continuity

The regression of continuity is based on internal supervision records of licence holders provided by the ANVS (2025a). Information on supervisory intensity and continuity was derived from these internal sources. Macroeconomic conditions were accounted for using national unemployment rate data obtained from publicly available statistics provided by the CBS (2025c).

B.3.3.1. Graph

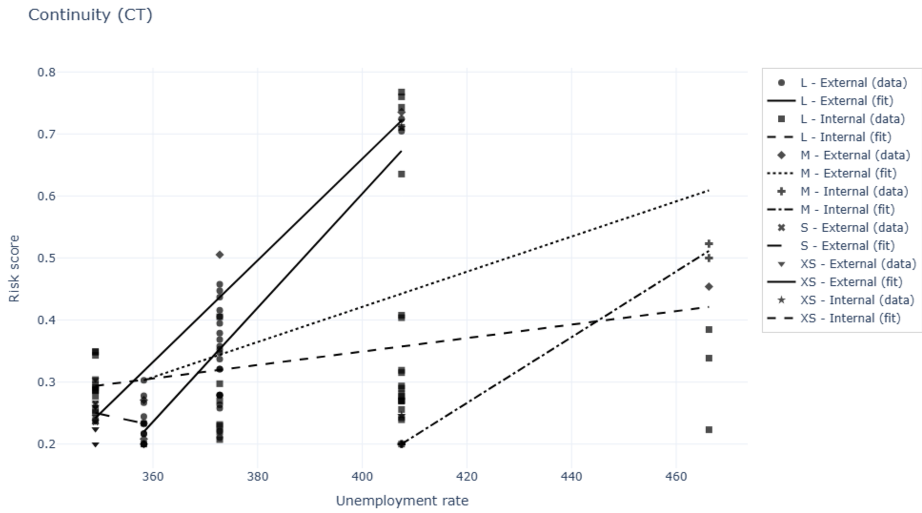

Figure B.4: Regression Continuity

B.3.3.2. Coefficients per group

Table B.4: Regression results by group

Radiation expert (I/E)	Organisation size	alpha	beta	n_obs
I	XS	0.000	0.001	2
I	S	-	-	1
I	M	-1.961	0.005	4
I	L	-0.085	0.001	44
E	XS	-2.619	0.008	11
E	S	0.879	-0.002	3
E	M	-0.713	0.005	4
E	L	-3.076	0.009	31

B.3.4. Expertise

The regression of expertise is based on internal supervision data provided by the ANVS (2025a). In addition, expertise reflects the extent to which the organisation of the licence holder is focused on radiation-related processes (ANVS, 2025a).

B.3.4.1. Graph

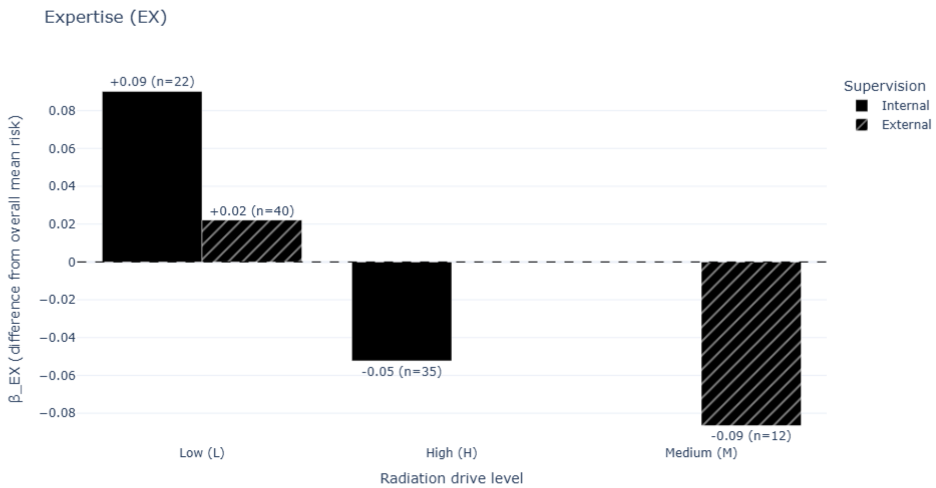

Figure B.5: Regression Expertise

B.3.4.2. Coefficients per group

Table B.5: Regression results by group

Radiation expert (I/E)	Radiation driven (L/M/H)	group_mean	beta	alpha	n_obs
I	L	0.420	0.090	0.330	22
I	M	-	-	-	-
I	H	0.278	-0.052	0.330	35
E	L	0.352	0.022	0.330	40
E	M	0.243	-0.087	0.330	12
E	H	-	-	-	-

B.3.5. License Scope

The regression of licence scope is based on the number of authorised activities, the number of different source types, and the innovation rate of licence holders from the internal sources of ANVS (2025a). Information on innovation rates was gathered from statistics provided by CBS (2025b).

B.3.5.1. Graph

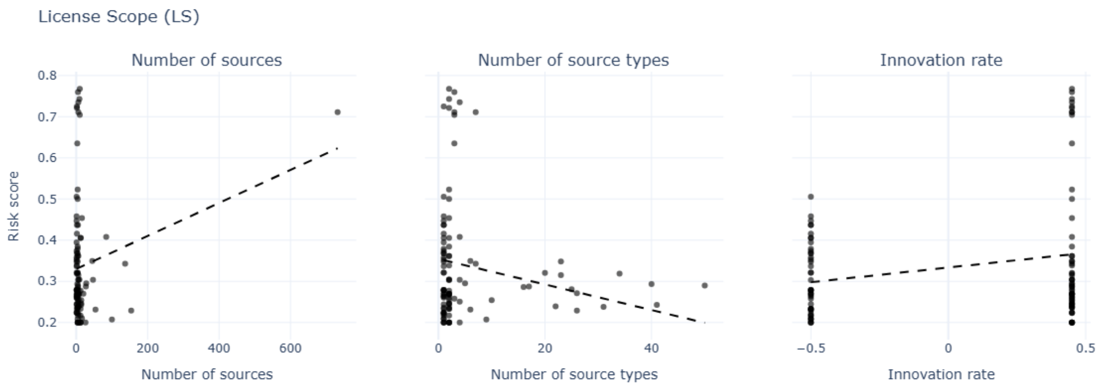

Figure B.6: Regression License Scope

B.3.5.2. Coefficients per group

Table B.6: Regression results by group

Alpha	beta source	beta nr types	beta innovation	n_obs
0.350	-0.001	-0.002	0.063	105

C

Appendix chapter 7 - sub-question 3

C.1. External factors

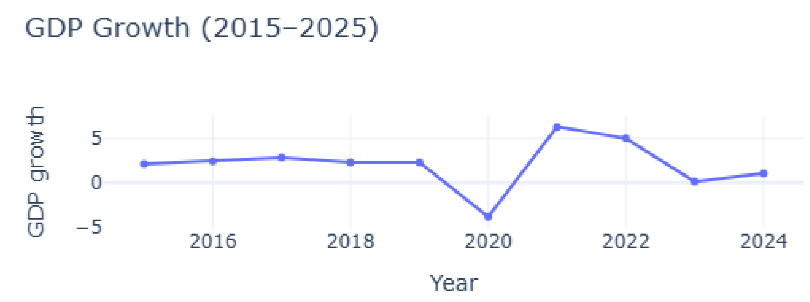

Figure C.1: GDP growth, from CBS (2025a)

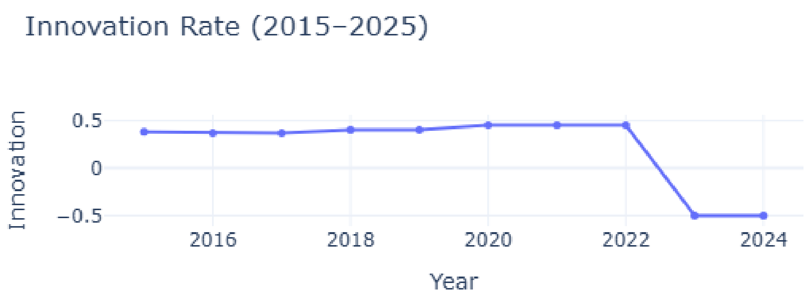

Figure C.2: Innovation rate, from CBS (2025b)

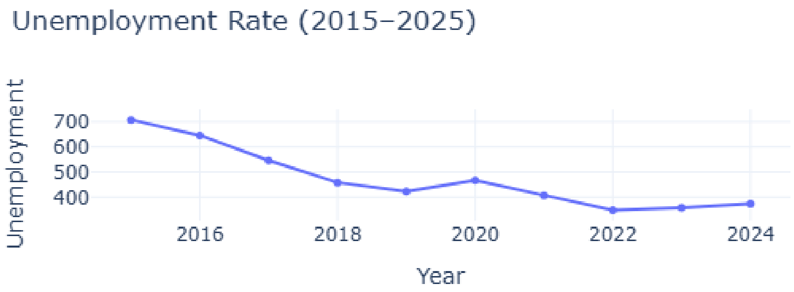

Figure C.3: Unemployment, from CBS (2025c)

C.2. Scenario analysis

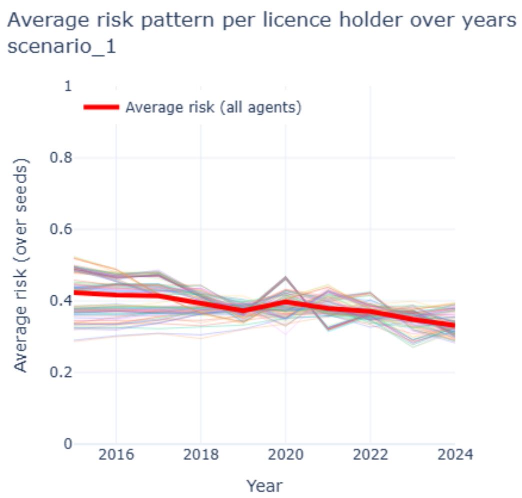

Figure C.4: Average risk pattern per licence holder over years, scenario 1

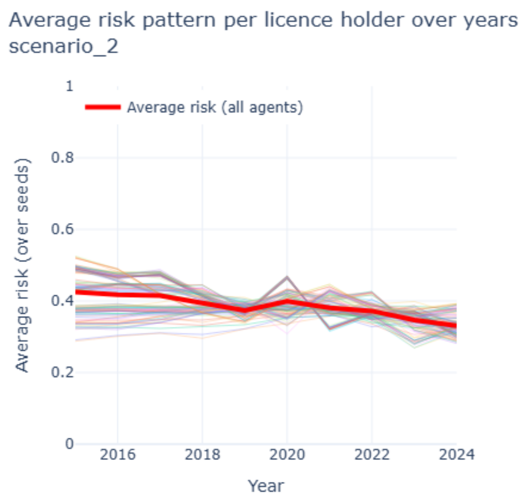

Figure C.5: Average risk pattern per licence holder over years, scenario 2

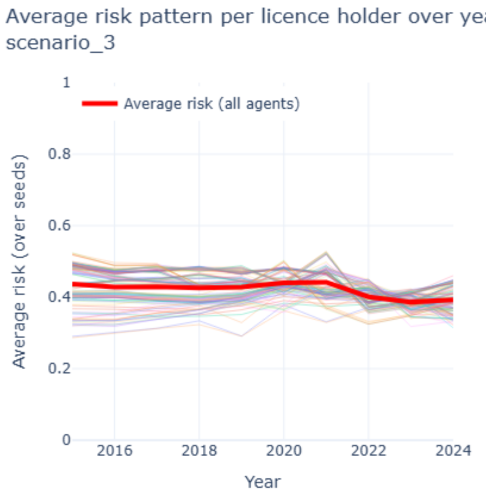

Figure C.6: Average risk pattern per licence holder over years, scenario 3

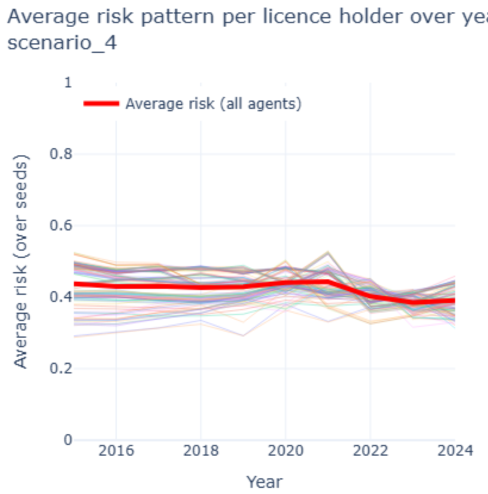

Figure C.7: Average risk pattern per licence holder over years, scenario 4

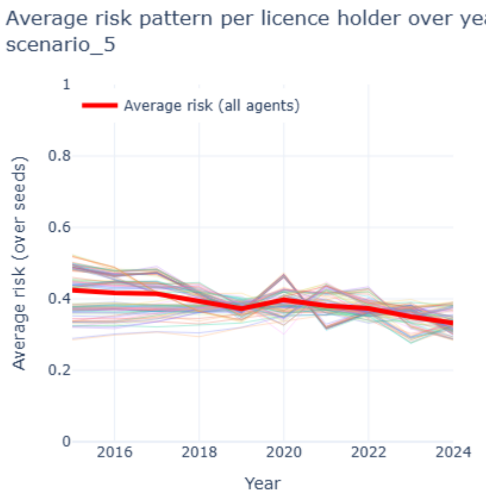

Figure C.8: Average risk pattern per licence holder over years, scenario 5

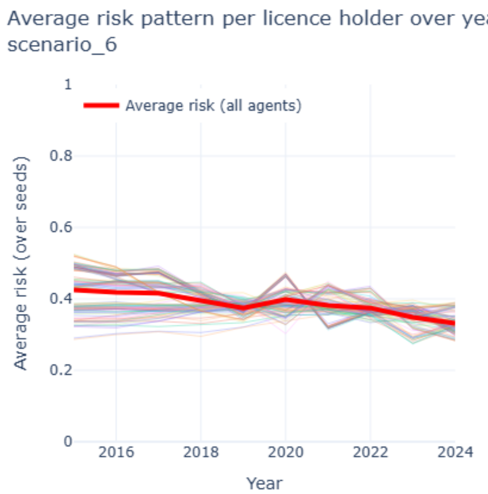

Figure C.9: Average risk pattern per licence holder over years, scenario 6

C.3. Sensitivity analysis

C.3.1. scenario 1 sensitivity

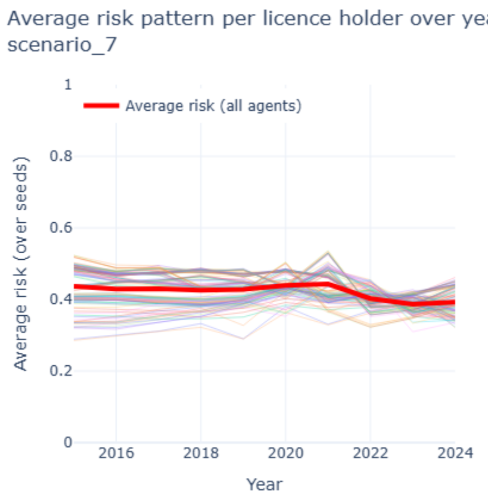

Figure C.10: Average risk pattern per licence holder over years, scenario 7

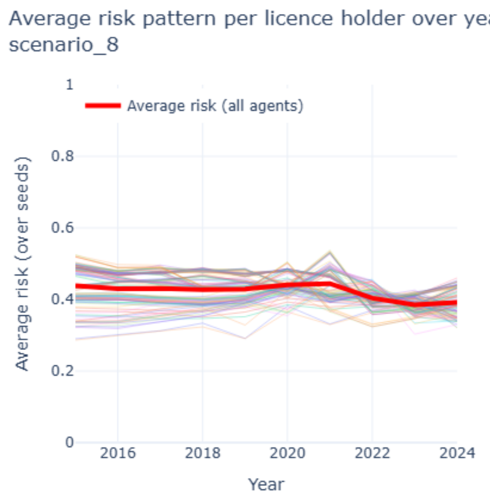

Figure C.11: Average risk pattern per licence holder over years, scenario 8

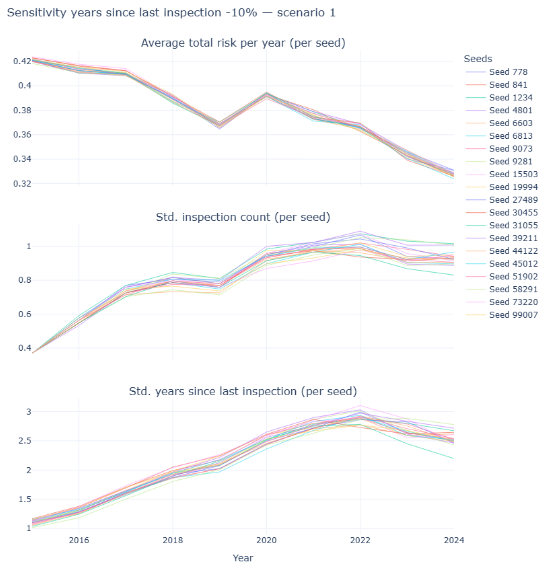

Figure C.12: Scenario 1 – years since last inspection −10%

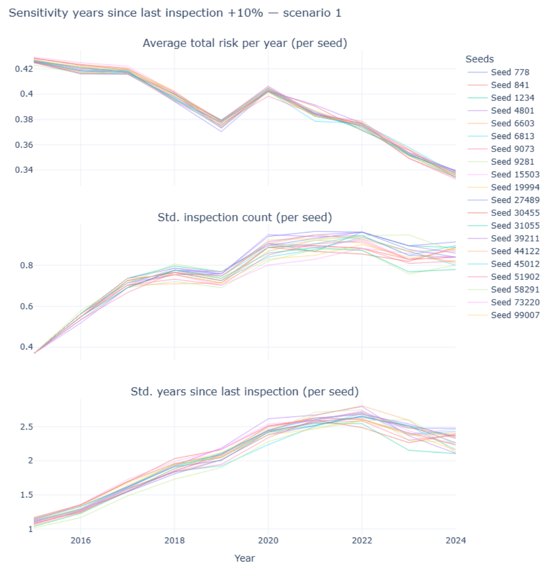

Figure C.13: Scenario 1 – years since last inspection +10%

Figure C.14: Scenario 1 – recent inspection bonus −10%

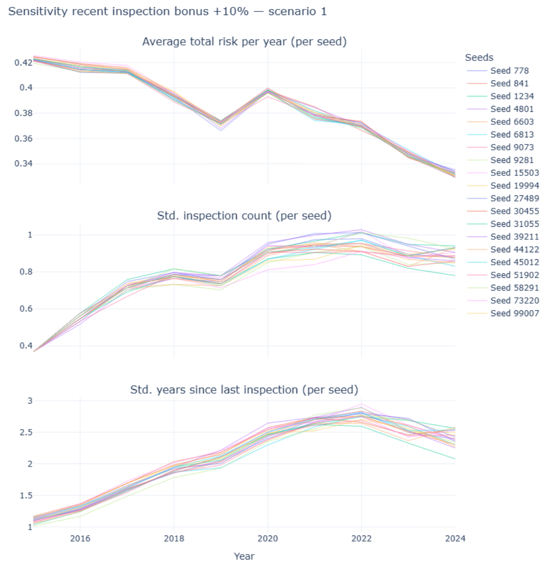

Figure C.15: Scenario 1 – recent inspection bonus +10%

Figure C.16: Scenario 1 – branch communication bonus −10%

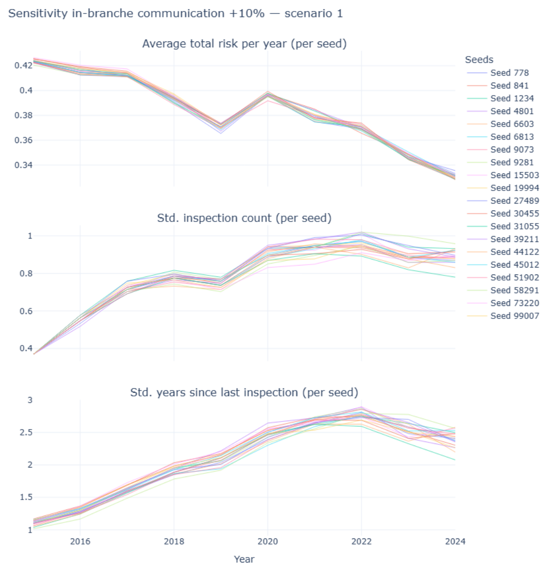

Figure C.17: Scenario 1 – branch communication bonus +10%

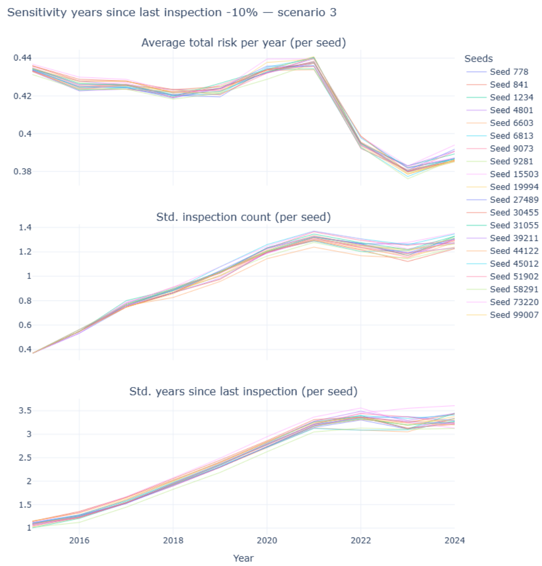

Figure C.18: Scenario 3 – years since last inspection −10%

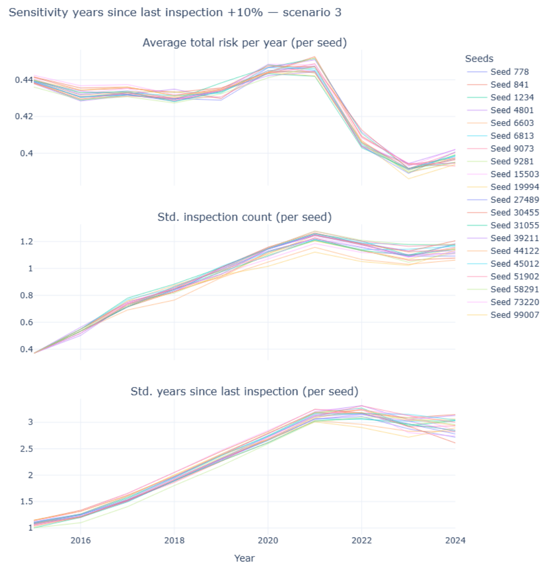

Figure C.19: Scenario 3 – years since last inspection +10%

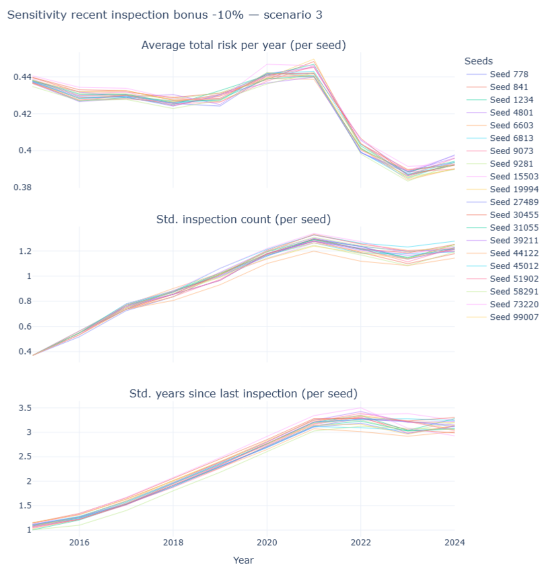

Figure C.20: Scenario 3 – recent inspection bonus −10%

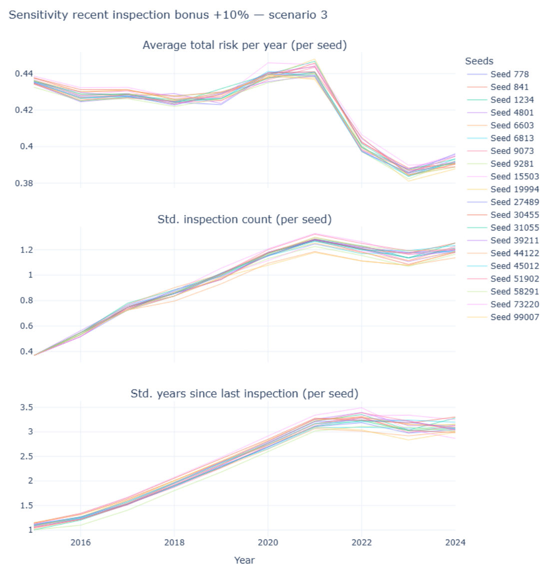

Figure C.21: Scenario 3 – recent inspection bonus +10%

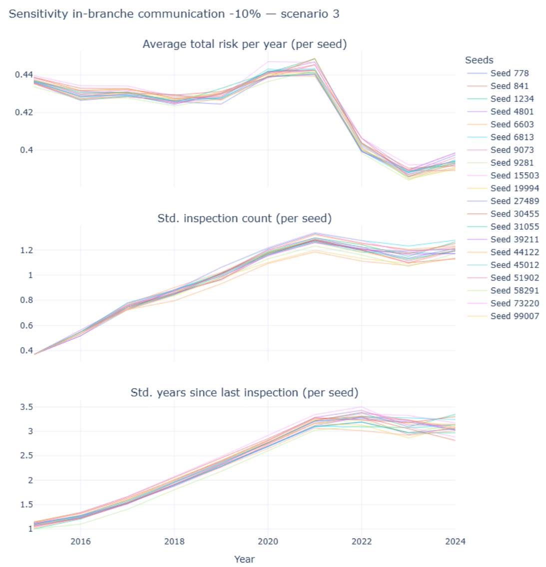

Figure C.22: Scenario 3 – branch communication bonus −10%

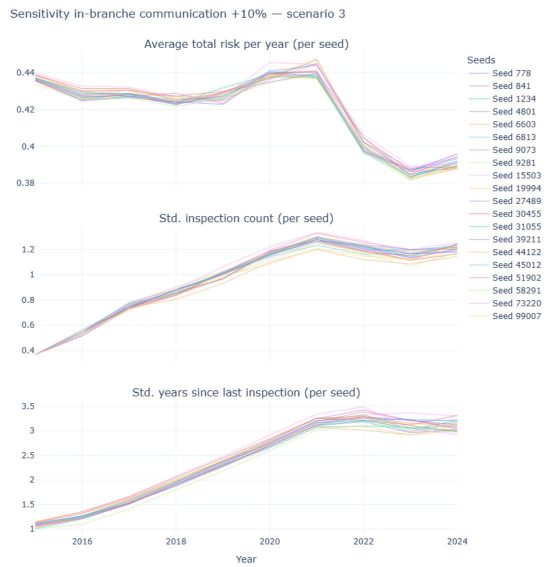

Figure C.23: Scenario 3 – branch communication bonus +10%

D

Appendix chapter 8 - sub-question 4

D.1. Threshold sensitivity

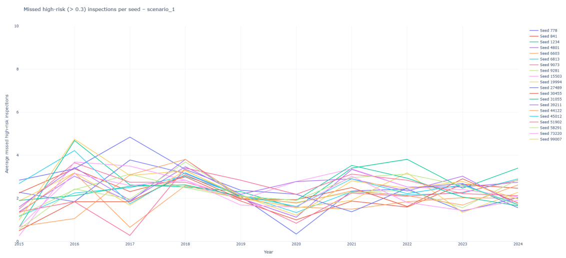

Figure D.1: Missed high-risk inspections (> 0.3), scenario 1

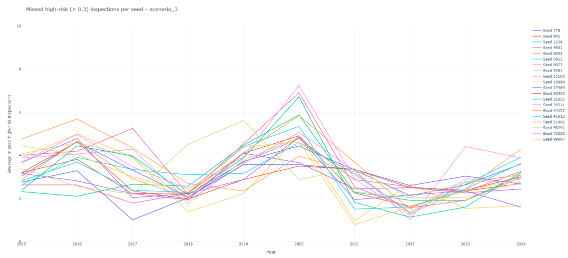

Figure D.2: Missed high-risk inspections (> 0.3), scenario 3

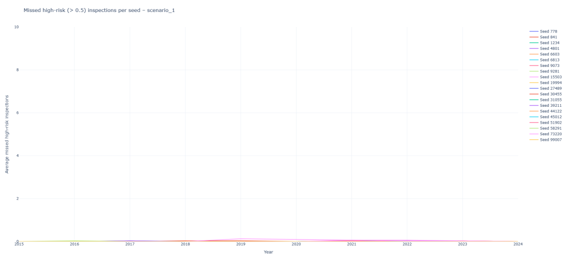

Figure D.3: Missed high-risk inspections (> 0.5), scenario 1

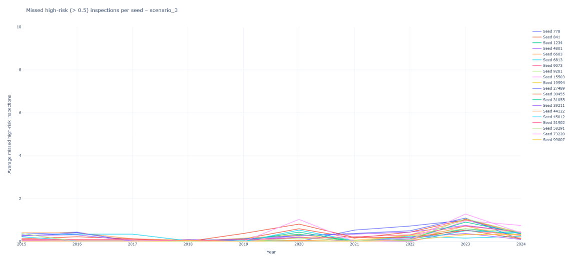

Figure D.4: Missed high-risk inspections (> 0.5), scenario 3

E

Appendix chapter 9 - sub-question 5

E.1. Policy 1 - random inspections

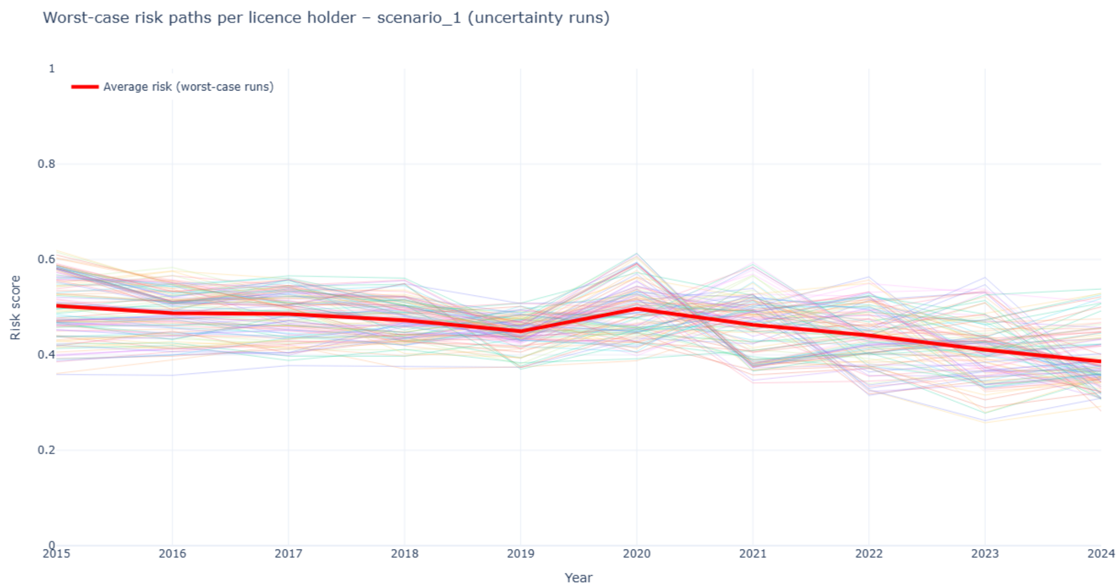

Figure E.1: Patterns inspections - scenario 1 - random

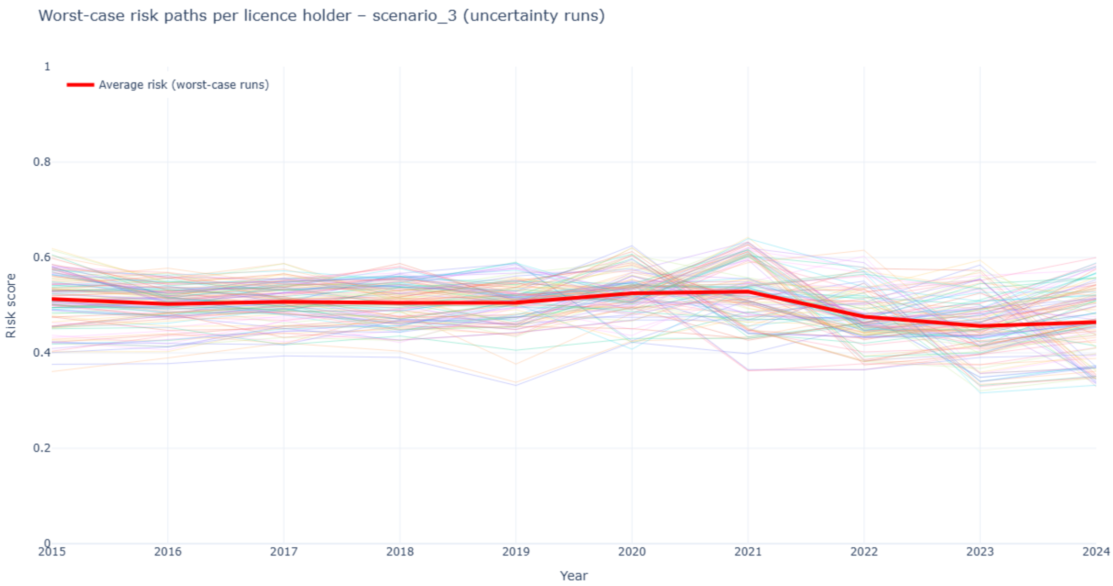

Figure E.2: Patterns inspections - scenario 3 - random

E.2. Policy 2 - Data quality

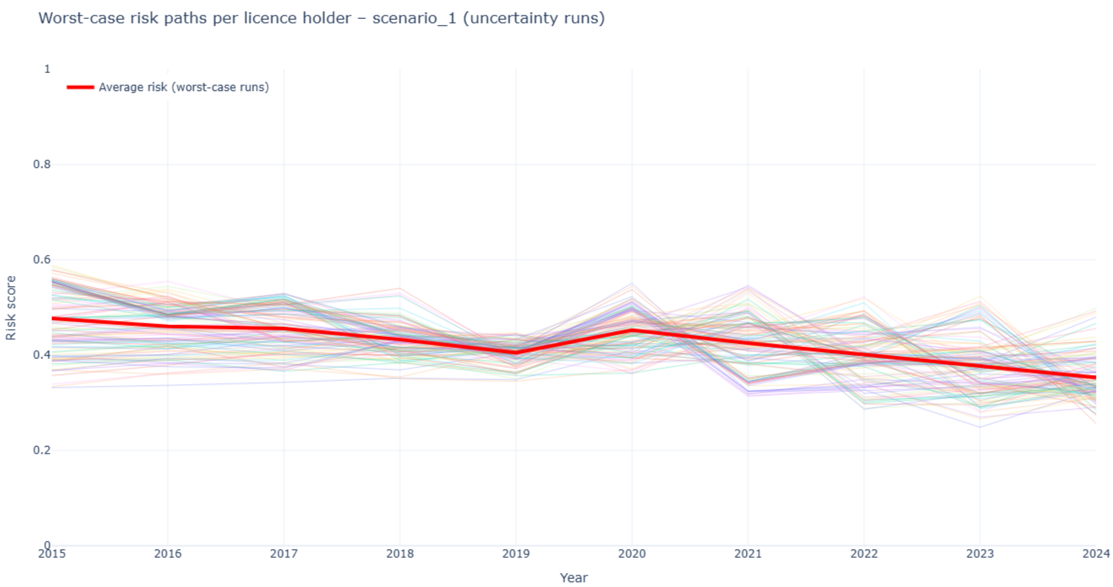

Figure E.3: Patterns inspections - scenario 1 - improved data quality

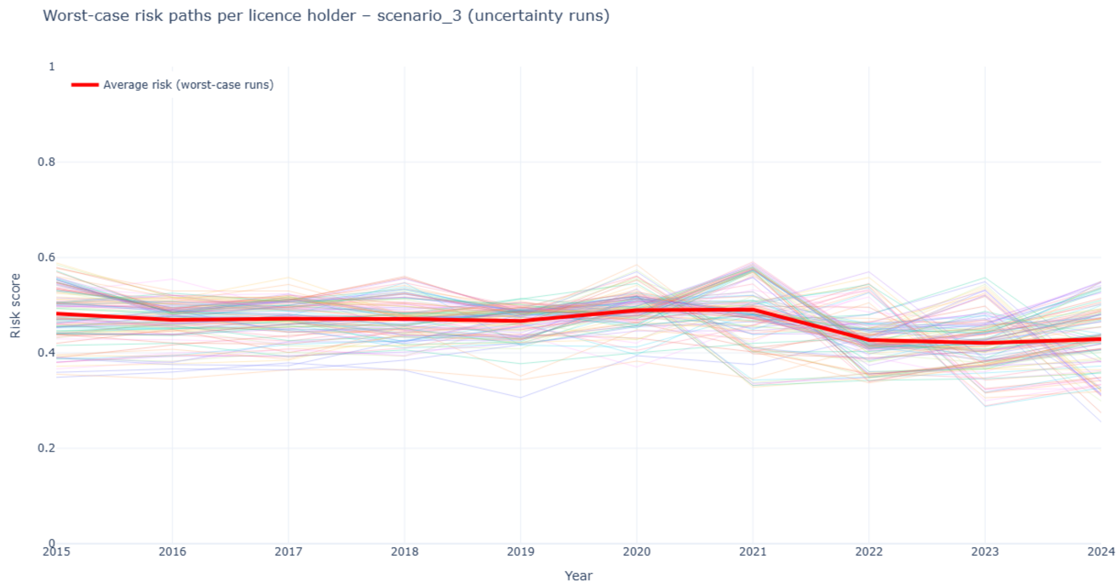

Figure E.4: Patterns inspections - scenario 3 - improved data quality

E.3. Policy 3 - Inspectors' capacity

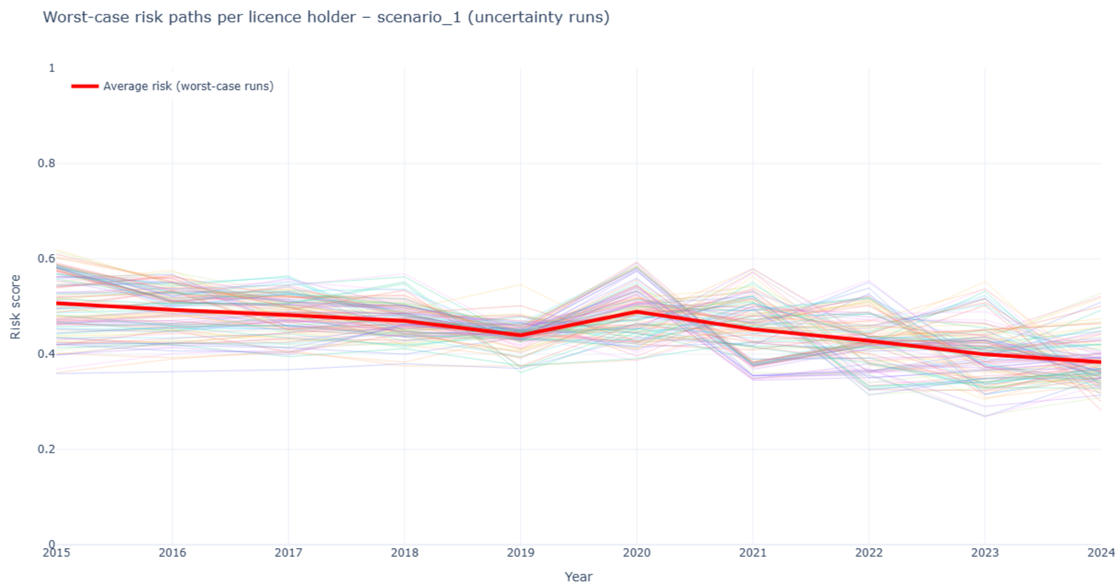

Figure E.5: Patterns inspections - scenario 1 - additional inspectors

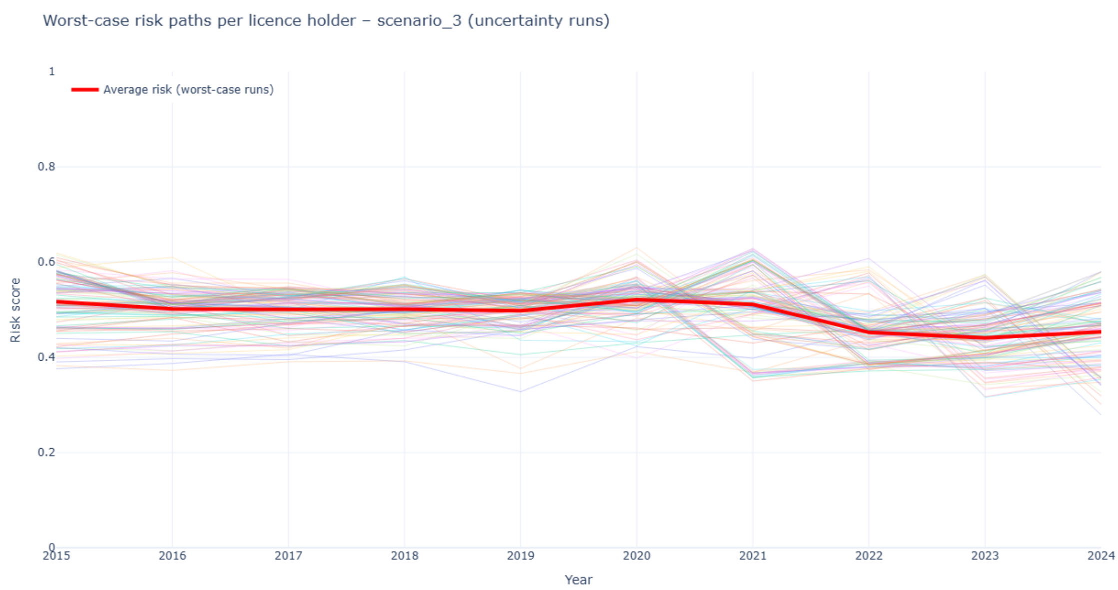

Figure E.6: Patterns inspections - scenario 3 - additional inspectors