SAMPLE REPORT - FICTIONAL ORGANISATION

GOVERNANCE DOCS

AI Risk Assessment

Northfield Lending (sample)

Every AI system the company develops, provides or uses: the credit decision model, the customer service chatbot and the generative AI assistant used by staff.

ASSESSED	25 September 2026
METHOD	ISO/IEC 42001:2023, clauses 6.1.2 and 6.1.3, with the AI system impact assessment of clause 6.1.4 in mind
IN SCOPE	11 scope items, 12 risks
APPETITE	Medium and below (level 9 or less) accepted without treatment
PROFILE	Banking and financial services
PREPARED BY	Governance Docs

CONTENTS

Where the organization's risks stand

PROCESS SCORE

91 / 100

Level 2 · Developing

The risks are identified, but gaps in rating or treatment would not survive an audit.

Held at Developing while 1 critical finding is open; the score alone would place it higher.

Initial | **Developing** | Defined | Managed

The score measures how complete and defensible the assessment is, not how risky the organization is.

RISKS	ABOVE APPETITE	CRITICAL	ABOVE AFTER TREATMENT	FINDINGS
12	5	0	1	4

WHAT THIS ASSESSMENT TELLS YOU

- 5 of 11 rated risks sit above your appetite line (Medium and below is acceptable).
- The highest risk is R-01, AI system treats groups of people unfairly, at level 16 (High), owned by Head of Data Science.
- 1 risk above the line has no treatment decision yet (R-05). Until it has one, the treatment plan is incomplete.
- Planned treatment brings the number of risks above the line from 5 to 1.
- 1 risk is not rated yet, so it is missing from the heat map and the priority order.

DO THESE FIRST

	SEVERITY	ACTION	RISKS	OWNER	DOCUMENT THAT CLOSES IT
1	CRITICAL	**Above the appetite line with no treatment decision** Choose modify, avoid, share or retain for each one, with the risk owner.	R-05	Procurement Manager	Risk treatment plan
2	MAJOR	**Not yet rated for likelihood and impact** Rate each one against the likelihood and impact scales in your criteria.	R-12	Legal Counsel	Risk assessment and treatment methodology

Heat maps: today and after treatment

Each cell shows how many risks sit at that likelihood and impact, with their references. The heavy line is the appetite line: Medium and below (level 9 or less) is accepted without treatment. The second map shows where each risk is expected to be once its treatment is carried out; avoided risks are removed.

Current risk levels

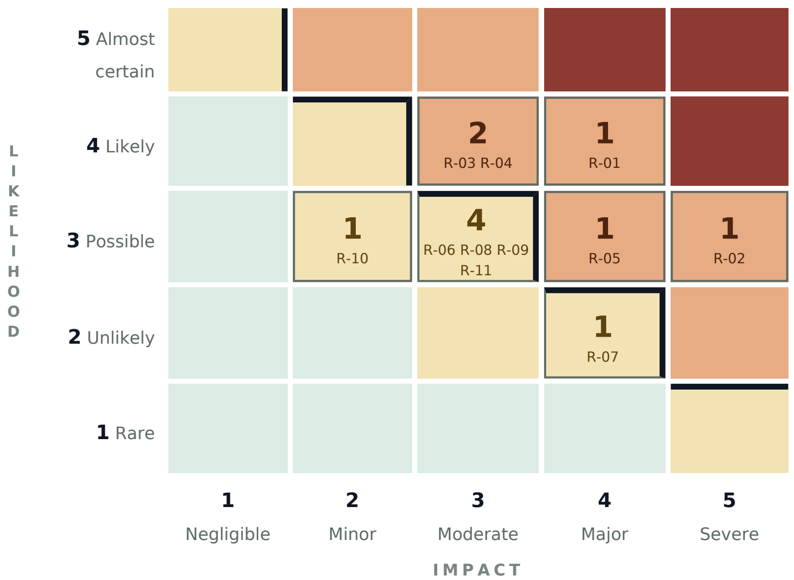

11 of 12 risks rated.

Target levels after treatment

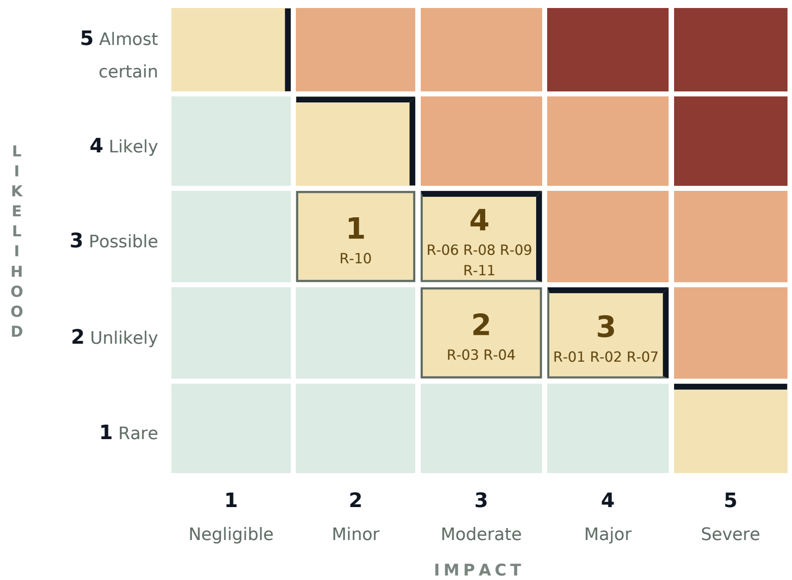

10 risks plotted; 0 removed by avoidance; risks above the line with no decision are not plotted.

Low 1–4 · Medium 5–9 · High 10–16 · Critical 17–25 · Appetite line

RISKS BY LEVEL

Analysis and priorities

IN ONE SENTENCE

Your AI risks are well mapped; decide how to treat the third-party model risk and settle the gaps in the register next.

AI-assisted analysis, generated from the information the organization provided and checked automatically against the scores and findings in this report. It supports planning and is not an audit opinion; see "Basis of this report and limitation of liability".

FOR SENIOR MANAGEMENT AND THE BOARD

Our AI risk assessment is largely complete: 12 risks identified, 5 above our appetite, and treatment planned that brings that number to 1. The largest exposure is the credit decision model, where applicants over 60 are approved less often with no credit reason found, and no impact assessment is on file. One decision is outstanding: how we govern the third-party model behind our chatbot and staff assistant.

WHERE YOU STAND	You have identified 12 AI risks across 11 scope items, and 11 are rated. Your process score is 91 out of 100, but the level is held at Developing because one risk above your appetite line, relying on a third-party AI model without due diligence, has no treatment decision. Every AI system, dataset and group of people in scope has at least one risk, and every treated risk has a target, an owner and a due date.
WHAT IT MEANS FOR THE ORGANIZATION	5 risks sit above your appetite line, all of them High, and the two highest concern the credit decision model: unfair treatment of groups of applicants and the missing impact assessment on the people it affects. These are the risks that fall hardest on individuals and draw the closest regulatory attention. Your planned treatment brings the number above the line from 5 to 1.
WHERE TO FOCUS FIRST	Decide how to treat the third-party model risk first, because it is what holds the level back and two of your AI systems depend on that model. Then close the gaps in the register: rate the legal and regulatory risk, give the event logging risk an owner, and record the owner's acceptance of the residual risk for AI data leakage.

PRIORITIES

	PRIORITY	WHY IT MATTERS	FIRST STEP	START FROM	WITHIN
1	**Decide how to treat the third-party model risk** R-05 Third-party AI model or service used without due diligence · Above the appetite line with no treatment decision	This High risk sits above your appetite with no decision, and two AI systems depend on the model, so the treatment plan is incomplete.	Agree a treatment with the Procurement Manager, such as a supplier due diligence review against A.10.3 Suppliers, and record it.	Risk treatment plan	30 DAYS
2	**Deliver the fairness testing and impact assessment** R-01 AI system treats groups of people unfairly · R-02 Impact on people not assessed before use	These are your two highest risks and both fall on loan applicants; the planned treatments bring each within appetite.	Set the fairness measures and the applicant groups the impact assessment will cover.	Risk treatment plan	60 DAYS
3	**Rate the legal and regulatory risk** R-12 An AI use breaches legal or regulatory obligations · Not yet rated for likelihood and impact	An unrated risk has no level, so it is missing from the heat map and cannot be compared with your appetite.	Rate likelihood and impact against your criteria, with a sentence of rationale.	Risk assessment and treatment methodology	30 DAYS
4	**Name an owner for the event logging risk** R-11 No record of what the AI system did · Risk with no owner	ISO/IEC 42001 expects every AI risk to have an owner, and chat transcripts are deleted before the complaint window closes.	Assign the risk to the Product Manager and review how long transcripts are kept.	Risk assessment and treatment methodology	30 DAYS

30 / 60 / 90-DAY ROADMAP

FIRST 30 DAYS

- Record a treatment decision for the third-party model risk.
- Rate the legal and regulatory risk.
- Assign an owner to the event logging risk.

DAYS 31 TO 60

- Collect residual-risk acceptance for AI data leakage.
- Complete the AI system impact assessment for the credit model.
- Ground chatbot answers in the approved knowledge base.

DAYS 61 TO 90

- Run fairness testing on the next credit model release.
- Send the selected controls to the Statement of Applicability.

What was in scope

The AI system, model, dataset, use case or process, AI provider or supplier and people affected the assessment covered, with their owners. How much harm this AI system, model or dataset could do to people or to the organization if it fails or behaves wrongly. H high, M medium, L low.

Type	Scope item	Owner	Criticality	Risks
AI system	**Credit decision model**	Head of Data Science	H	4
AI system	**Customer service chatbot**	Product Manager	M	3
AI system	**Staff generative AI assistant**	Chief Technology Officer	M	2
Model	**Third-party large language model**	Chief Technology Officer	M	1
Dataset	**Loan application training data**	Head of Data Science	H	3
Dataset	**Knowledge base used for answers**	Product Manager	M	2
Use case or process	**Loan applications**	Compliance Officer	H	3
Use case or process	**Customer support**	Product Manager	M	2
AI provider or supplier	**Foundation model provider**	Procurement Manager	M	1
People affected	**Loan applicants**	Compliance Officer	H	2
People affected	**Employees using AI tools**	Information Security Manager	L	1

Risk register

Every risk, highest level first. Level = likelihood x impact, rated with the controls already in place. The rationale is the organization's own reason for the rating. The full scenario for each risk (threat, weakness and consequence) is in the workbook.

#	Ref	Risk scenario	Scope items	Objectives at risk	Owner	Existing controls	L	I	Level	Rationale
1	R-01	**AI system treats groups of people unfairly** Weakness: No fairness testing before release or in operation	Credit decision model, Loan applicants, Loan application training data	Fairness, Accountability	Head of Data Science	Accuracy tested before each release	4	4	16 · High	Approval rates for applicants over 60 are 11 points lower, with no credit reason found so far.
2	R-02	**Impact on people not assessed before use** Weakness: No AI system impact assessment, or one that is not recorded	Credit decision model, Loan applicants, Loan applications	Fairness, Safety, Privacy	Compliance Officer	Model review by the credit committee	3	5	15 · High	Credit decisions shape access to finance; no impact assessment on applicants is on file.
3	R-05	**Third-party AI model or service used without due diligence** Weakness: No review of how the model was built, tested and governed	Foundation model provider, Third-party large language model	Robustness, Accountability	Procurement Manager	Standard software contract	3	4	12 · High	The chatbot and the staff assistant both run on one provider's model, whose training and testing are not documented to us.
4	R-03	**Generative AI produces false or made-up output** Weakness: Outputs used without checking or grounding in reliable sources	Customer service chatbot, Knowledge base used for answers, Customer support	Transparency, Robustness	Product Manager	Disclaimer at the start of each chat	4	3	12 · High	A review of 300 chats found 14 answers that misstated fees or repayment terms.
5	R-04	**Confidential data leaks through AI tools** Weakness: No rules on what may be entered, and no approved alternative	Staff generative AI assistant, Employees using AI tools	Privacy, Security	Information Security Manager	IT acceptable use policy	4	3	12 · High	Web logs show staff using three public AI tools; the IT policy does not mention them.

6	R-06	**Over-reliance on AI with weak human oversight** Weakness: No defined human review for decisions that affect people	Loan applications, Credit decision model	Fairness, Safety, Accountability	Compliance Officer	Underwriters review declined applications over 25,000	3	3	9 · Medium	Smaller declines are automatic; underwriters override fewer than 1% of cases.
7	R-08	**Live data drifts away from the training data** Weakness: Inputs not monitored against the data the model was trained on	Credit decision model, Loan application training data	Robustness	Head of Data Science	Monthly accuracy report	3	3	9 · Medium	The application mix changed after the broker channel opened; accuracy is holding so far.
8	R-09	**AI systems in use that nobody has recorded** Weakness: No inventory of AI systems, the models behind them and the data they use	Staff generative AI assistant	Privacy, Security, Accountability	AI Governance Lead	List of approved software	3	3	9 · Medium	AI features in the HR and CRM systems were found only during this assessment.
9	R-11	**No record of what the AI system did** Weakness: Event logging not built in, or logs not kept long enough	Customer service chatbot, Customer support	Transparency, Accountability	**None**	Chat transcripts kept for 30 days	3	3	9 · Medium	Transcripts are deleted before the complaint window closes.
10	R-07	**Personal data used for AI without a lawful basis** Weakness: No privacy review before data is used for AI	Loan application training data, Loan applications	Privacy	Data Protection Officer	Privacy notice covers credit scoring; DPIA completed for the model	2	4	8 · Medium	Training data is pseudonymised and the DPIA is current.
11	R-10	**Prompt injection or jailbreak takes control of an AI system** Weakness: Inputs and connected tools not restricted or tested against attack	Customer service chatbot, Knowledge base used for answers	Safety, Security	Information Security Manager	Chatbot has no access to customer accounts	3	2	6 · Medium	The chatbot answers from the knowledge base only and cannot take actions.
12	R-12	**An AI use breaches legal or regulatory obligations** Weakness: Obligations for each AI system not identified, assigned or tracked	**None linked**	Fairness, Transparency, Accountability	Legal Counsel		–	–	Not rated	

What will be done, by whom and by when

Every risk with a treatment decision, earliest due date first. The ISO 42001 Annex A controls listed are the ones the treatment relies on; they carry into the Statement of Applicability. Acceptance is the risk owner's approval of the plan and of the risk that remains.

Ref	Risk	Decision	Planned actions	ISO 42001 Annex A controls	Owner and due date	Current to target	Accepted
R-04	**Confidential data leaks through AI tools** Information Security Manager	**Modify**	Give every team the approved staff assistant, block unapproved AI tools and add AI rules to the acceptable use policy.	A.9.2 Processes for responsible use of AI systems A.2.2 AI policy A.10.3 Suppliers	Information Security Manager 15 Nov 2026	**12 · High** → **6 · Medium**	**Not recorded**
R-03	**Generative AI produces false or made-up output** Product Manager	**Modify**	Answer only from the approved knowledge base and hand fee and hardship questions to a person.	A.6.2.4 AI system verification and validation A.9.4 Intended use of the AI system A.8.2 System documentation and information for users	Product Manager 30 Nov 2026	**12 · High** → **6 · Medium**	Chief Executive 22 Sep 2026
R-01	**AI system treats groups of people unfairly** Head of Data Science	**Modify**	Add fairness testing across age, sex and disability to every release and quarterly in production; review features that act as proxies for age.	A.5.4 Assessing AI system impact on individuals or groups A.6.2.4 AI system verification and validation A.7.4 Quality of data for AI systems	Head of Data Science 31 Dec 2026	**16 · High** → **8 · Medium**	Chief Executive 22 Sep 2026

R-02	**Impact on people not assessed before use** Compliance Officer	**Modify**	Carry out and record an AI system impact assessment for the credit model, covering applicants and the groups most at risk of unfair outcomes.	A.5.2 AI system impact assessment process A.5.3 Documentation of AI system impact assessments A.5.4 Assessing AI system impact on individuals or groups	Compliance Officer 31 Jan 2027	**15 · High** → **8 · Medium**	Chief Executive 22 Sep 2026
R-06	**Over-reliance on AI with weak human oversight** Compliance Officer	**Retain**	Risk retained at its current level.		–	**9 · Medium** → **9 · Medium**	Chief Executive 22 Sep 2026

ABOVE THE LINE WITH NO DECISION YET

Ref	Risk	Owner	Level
R-05	**Third-party AI model or service used without due diligence**	Procurement Manager	**12 · High**

5 risks are within appetite (Medium and below) and need no treatment; 1 are not yet rated.

Where the assessment falls short

Each finding is a gap between the assessment as recorded and what ISO 42001 clauses 6.1.2 and 6.1.3 ask for, with the risks it applies to and the document that closes it.

Severity	Finding	Applies to	What it means	To fix	Document that closes it
CRITICAL	**Above the appetite line with no treatment decision**	R-05	These risks are rated above Medium, the highest level you accept without treatment, but nobody has decided what to do about them. ISO 42001 expects a treatment option to be chosen for every risk that fails the acceptance criteria.	Choose modify, avoid, share or retain for each one, with the risk owner.	Risk treatment plan
MAJOR	**Not yet rated for likelihood and impact**	R-12	Without both ratings a risk has no level, so it cannot be compared with the appetite line or placed in priority order.	Rate each one against the likelihood and impact scales in your criteria.	Risk assessment and treatment methodology
MAJOR	**Risk with no owner**	R-11	ISO/IEC 42001 expects every AI risk to have an owner: the person with the accountability and authority to manage it.	Name an owner, by role, for each risk.	Risk assessment and treatment methodology
MINOR	**Residual risk not yet accepted by the owner**	R-04	ISO/IEC 42001 asks risk owners to approve the AI risk treatment plan and accept the risk that remains. A name and a date is enough to show it happened.	Record the owner's acceptance once they have approved the plan.	Risk treatment plan

PROCESS SCORE IN DETAIL

Criteria defined Every likelihood and impact level has a definition, and the appetite line is set.	100%	11 of 11	weight 10
Scope covered Every scope item has at least one risk linked to it.	100%	11 of 11	weight 10
Risk owners named Every risk has an owner accountable for it.	92%	11 of 12	weight 15
Risks rated Every risk has a likelihood and an impact, so it has a level.	92%	11 of 12	weight 15
Ratings explained Existing controls and a rationale are recorded for each risk.	92%	22 of 24	weight 10
Treatment decided Every risk above the appetite line has a treatment option chosen.	80%	4 of 5	weight 15
Targets within appetite Treatments bring each risk to a target at or below the line.	100%	4 of 4	weight 10
Treatments scheduled Each treatment has an owner and a due date.	100%	4 of 4	weight 5
Residual risk accepted Risk owners have accepted what remains after treatment, by name and date.	75%	3 of 4	weight 10

Findings by severity

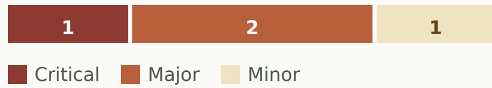

1 · Initial (0+)
A start on a register, not yet an assessment that decisions can rest on.

2 · Developing (40+)
The risks are identified, but gaps in rating or treatment would not survive an audit.

3 · Defined (60+)
A complete, repeatable assessment with a treatment plan behind it.

4 · Managed (80+)
Owned, justified and accepted: ready for management review and audit.

Criteria used and basis of this report

ISO/IEC 42001 asks for AI risk criteria that tell acceptable risks from unacceptable ones and give consistent, valid and comparable results, with consequences weighed for individuals and societies as well as for the organization. This is the record of the criteria every rating in this report was made against.

LIKELIHOOD SCALE

Value	Label	Definition
1	**Rare**	Not expected in the next 5 years; no known case here or in the sector
2	**Unlikely**	Could happen once in 2 to 5 years; has happened in the sector
3	**Possible**	Could happen about once a year
4	**Likely**	Expected several times a year
5	**Almost certain**	Expected monthly or more often, or already happening

IMPACT SCALE

Value	Label	Definition
1	**Negligible**	No noticeable effect on the people affected; under 10,000
2	**Minor**	Minor inconvenience to a limited group; corrected quickly; 10,000 to 50,000
3	**Moderate**	Wrong or unfair outcomes for many people; complaints; possible regulator query
4	**Major**	Significant harm to people or groups; regulator notified; media interest; 250,000 to 1 million
5	**Severe**	Serious, lasting or irreversible harm, or harm to society; enforcement or withdrawal; over 1 million

RISK LEVELS AND ACCEPTANCE

Band	Levels	Treatment
Low	1–4	Accepted without treatment
Medium	5–9	Accepted without treatment
High	10–16	Needs a treatment decision
Critical	17–25	Needs a treatment decision

Level = likelihood x impact. A risk above the appetite line may still be retained, but only with the risk owner's recorded acceptance.

Process score. Nine dimensions, each the share of risks (or of risks above the line) that meet it, multiplied by its weight; the weights add up to 100. A dimension with nothing in scope counts in full. Any critical finding holds the level at Developing, whatever the score.

Your workbook. The Excel workbook that comes with this report is the register as a working file. Change a likelihood or impact on the Risk register sheet and every level, the heat maps on the Dashboard and the process score recalculate. It also holds the treatment plan and a Statement of Applicability starter listing the ISO 42001 Annex A controls your treatments rely on.

BASIS OF THIS REPORT AND LIMITATION OF LIABILITY

This report is the output of a self-assessment. The risks, ratings, risk levels, treatment decisions, findings, analysis and recommendations in this report are derived solely from the information entered by the organization, and have not been independently reviewed, verified, tested or audited by Governance Docs. Governance Docs did not conduct an independent assessment of the organization, its controls or its operations.

The accuracy and completeness of the results depend entirely on the accuracy and completeness of the information provided. Governance Docs gives no assurance, certification or audit opinion, and accepts no liability for any inaccuracy, omission or deviation arising from that information, or for any decision taken, or not taken, in reliance on this report.

The report, including any AI-assisted analysis, is provided for guidance and planning only. It does not constitute legal, regulatory or other professional advice, and it does not replace an independent assessment or a certification audit.

AI risk assessment for Northfield Lending (sample), generated 25 September 2026 by Governance Docs. This is a structured self-assessment based on the organization's own criteria and ratings. It is not a certification, an audit opinion or a substitute for an independent risk assessment.