

# THE DIRECTOR ACCOUNTABILITY TEST



RUSSELL PARROTT

# COPYRIGHT

Copyright © 2026 Russell Parrott

All rights reserved.

No part of this publication may be reproduced, stored or transmitted in any form or by any means without prior written permission from the copyright holder, except for brief quotations used in reviews, commentary or scholarly reference.

This publication is provided for educational and informational purposes only. It does not constitute legal advice, regulatory advice or professional services.

## Licensing

This publication may be shared and redistributed under the terms of the Creative Commons Attribution–NoDerivatives 4.0 International licence (CC BY-ND 4.0).

The work may be copied and circulated with attribution but may not be modified, adapted or republished in altered form.

## THE DIRECTOR ACCOUNTABILITY TEST

AI is changing the environment directors operate in, yet most boards are still relying on oversight habits built for a much slower world.

In the past, decisions moved through relatively stable systems. The people involved could usually be identified, the records survived, and the chain of events could still be followed years later if questions arose. Directors depended heavily on the organisational structure because underlying operations changed slowly enough for ordinary oversight to remain useful.

AI disrupts those assumptions. Systems now operate at enormous speed, depend heavily on external suppliers and can change continuously after deployment. The result is an environment where organisations may struggle to explain later how a specific outcome occurred or why a particular decision was made.

Most directors have never previously carried this kind of exposure.

Virtually all existing AI governance material concentrates on process, structure and visible oversight. Boards are shown policies, committees, reporting lines and assurance mechanisms designed to demonstrate organisational control.

Those things may help a company appear organised internally, but they do not necessarily protect individual directors once scrutiny begins.

Commercial liability insurers see this differently from most governance advisers because they operate in the world after failure. D&O, cyber and professional liability claims often begin years after the original decisions were made.

They are not interested in whether governance looks convincing internally. But whether events can still be reconstructed years later. Does the conduct still appear reasonable? Can the surviving evidence withstand hostile scrutiny from regulators, claimants or investigators?

That creates a much harsher perspective. An insurer is rarely impressed by a polished governance presentation or a large oversight structure on paper. The real question is whether the organisation can still explain and defend its decisions once events are pulled apart under hostile scrutiny years later.

The central issue is not if governance looked impressive at the time, but can the responsibilities, judgement and decision-making can still be proven later once scrutiny begins.

Many directors already sense this problem privately. They are being asked to approve systems whose behaviour they cannot fully see, while public pressure, regulatory expectations and liability exposure continue to increase.

Ultimately responsibility still rests with humans, even as those systems become harder to follow, explain and examine later.

This creates a dangerous imbalance. Directors remain legally and financially exposed, yet their ability to clearly reconstruct and explain what happened may steadily weaken over time.

## THE DIRECTOR ACCOUNTABILITY TEST

---

The purpose of this test is not to teach directors how to build AI systems. It is to help Directors understand what happens later when decisions are challenged, evidence is demanded and they are expected to account for what occurred.

## One system at a time

This test should be completed separately for every AI system and whenever that system is changed in a way that could affect how it behaves.

Every AI system should be considered separately. You may be able to answer 'Yes' for one system and 'No' for another. The questions should therefore be answered for the specific AI system being considered, not for the organisation as a whole

### Personal decision participation

Y

N

**Can you show that you personally took part in reviewing the proposed AI use before it was approved, rather than relying only on the fact that the organisation had followed a process?**

This checks whether you were personally involved in the judgement, not merely present in the room or copied into papers.

### Information actually reviewed

Y

N

**Can you show what information you actually read or considered before approving, accepting or relying on the AI system?**

This checks whether there is evidence of what you personally reviewed, rather than simply proving that documents existed somewhere inside the organisation.

### Understanding of known risks

Y

N

**Can you show that you understood the known limits, weaknesses and risks of the AI system at the time you allowed it to be used?**

This checks whether you knew enough about the system's risks to make a responsible decision, rather than accepting a general assurance that the system was safe or compliant.

### Challenge and questioning

Y

N

**Can you show that you asked questions, challenged assumptions or requested clearer answers before accepting what you were told about the AI system?**

This checks whether your involvement showed active judgement, rather than passive acceptance of management, technical or vendor reassurance.

### Reliance on others

Y

N

**Can you prove and explain why it was reasonable for you to rely on the people, teams, advisers or suppliers who gave you information about the AI system?**

This checks whether your reliance on others can still be justified later, especially if the information they gave you proves incomplete, wrong or overly confident.

### Awareness of escalated concerns

Y

N

**Can you show whether any warnings, objections, unresolved concerns or internal doubts reached you before the AI system was approved or used?**

This checks what you knew before the event, not what the organisation later says in general terms after a problem has occurred.

### Authority to stop or delay

Y

N

**Did you have the real authority to stop, delay, limit or question the use of the AI system if the risks were not properly explained?**

This checks whether your oversight role gave you practical power, or whether your position only created the appearance of control.

### Use of independent judgement

- Y  N **Can you prove and show that you reached your own judgement about the AI system, rather than simply approving what management, advisers or suppliers recommended?**

This checks whether approval reflected your own judgement, not just your participation in a formal sign-off process.

### Record of personal concerns

- Y  N **If you had concerns, doubts or reservations about the AI system, is there a record showing what you raised and how those concerns were handled?**

This checks whether your concerns can still be proven later, rather than depending on memory after the organisation is already under scrutiny.

### Awareness of evidence gaps

- Y  N **Before approval, were you told about any missing records, incomplete testing, unclear responsibility or unresolved uncertainty about how the AI system would behave?**

This checks whether uncertainty was visible to you before the decision, and whether you knowingly accepted the risk of proceeding despite it.

### Personal accountability preservation

- Y  N **Could you still prove and show, five years later, what you knew, what you considered, what you questioned and why you accepted the decision at the time?**

This checks whether your reasoning can be reconstructed from evidence, rather than relying on memory once the decision is challenged.

### Insurance disclosure exposure

- Y  N **Can you prove and show that known AI risks, unresolved concerns or material uncertainties were not hidden from insurers when cover was requested or renewed?**

This checks whether insurers were given a fair picture of the risk, rather than being asked to cover exposure that was already known but not clearly disclosed.

### Personal conduct defensibility

- Y  N **If this AI decision were challenged today, would the surviving evidence make your own conduct appear careful, informed and reasonable?**

This checks the central issue: whether your personal conduct can still be defended once organisational language is stripped away.

## How to read your result

There is no completely safe score. Every "No" answer identifies somewhere your ability to later explain, defend or reconstruct your conduct is weaker than it should be. Some weaknesses may appear small today, but they may not appear small once regulators, insurers, claimants or investigators begin examining events years later.

It is tempting to focus on the total number of "No" answers. That is understandable, but it misses the more important question: What exactly can you still prove?

A director with only one evidential weakness may face greater difficulty than a director with several minor ones. The issue is not how many boxes were ticked. But whether a critical part of your involvement, judgement or decision-making can still be demonstrated from surviving evidence.

Every "No" answer should therefore be treated seriously.

A "No" means that something important may no longer be visible. It may be unclear what information you reviewed, impossible to show what concerns reached you or difficult to demonstrate what questions you asked, what risks you understood, or why you accepted the decision at the time.

The purpose of this test is not to determine whether the organisation appears well governed. It is to examine if your own conduct could still withstand scrutiny after the protection of hindsight, memory and organisational language has disappeared.

The central question is always the same: If this decision were challenged years from now, could you still prove what you knew, what you did and why you did it?

## What fixes a "No"?

A "No" answer does not exist because somebody made a mistake. It exists because something important cannot be shown.

Perhaps there is no record of what information you reviewed, or no evidence of what concerns reached you before approval. Even no surviving record showing what questions you asked, what risks you understood or why you accepted the decision at the time.

Whatever the reason, the result is the same. A gap exists between what happened and what can now be proved.

Many directors assume that can be repaired later. They think they will explain what they knew, describe what they considered and reconstruct their reasoning if questions ever arise. Unfortunately, accountability rarely works that way. Explanations are easy to produce after an event. Evidence is much harder.

As time passes, memories become less reliable. People move on, systems change, records disappear and the organisation develops a new narrative around what happened. However the question facing insurers, regulators, claimants and investigators remains remarkably simple: what can actually be demonstrated from the surviving evidence?

## THE DIRECTOR ACCOUNTABILITY TEST

---

That is why every question in this test ultimately points to the same issue. Not whether something happened, but can it can still be shown to have happened.

A process cannot repair a missing record. Good intentions cannot repair a missing record. Confidence in your recollection cannot repair a missing record. If the evidence does not exist, the weakness remains.

The only thing that changes a "No" into a "Yes" is evidence capable of supporting the answer.

That is why the most important asset in any future investigation is not the governance framework, the policy library or the organisational chart. It is the evidence that allows your involvement, judgement and conduct to be reconstructed long after the original decision was made.

## EXPLANATORY NOTES

The questions in this test may appear simple. In practice, each examines a different way personal accountability can become difficult to prove later.

The issue is rarely whether something happened. The issue is whether it can still be shown to have happened years afterwards.

A director may genuinely have reviewed the information, understood the risks, questioned the proposal and exercised independent judgement. Yet if no evidence survives showing those things occurred, demonstrating them later may become extremely difficult.

The explanations that follow examine each question individually. They are not intended to provide legal advice or prescribe particular actions. Their purpose is to explain why the question matters, what type of evidential weakness it is testing and why that weakness may become important once a decision is challenged.

### **Question 1: Can you show that you personally took part in reviewing the proposed AI use before it was approved, rather than relying only on the fact that the organisation had followed a process?**

For many years, directors could largely demonstrate reasonable conduct by showing that the organisation had followed an established process. A proposal was reviewed, papers were circulated, a committee considered the issue, a vote was taken and the decision was recorded. If questions arose later, those organisational records often formed the basis of the defence.

That position is becoming increasingly difficult to rely upon where AI is concerned. Systems are more complex, decisions are harder to reconstruct and the gap between approval and outcome is often much larger than directors expect. As a result, scrutiny is gradually moving away from whether the organisation appeared well governed and towards whether individual directors can still demonstrate what they personally knew, what they personally considered and why they acted as they did.

This is where many directors become exposed without realising it. It is entirely possible to attend every meeting, receive every board pack and participate in every approval process without ever engaging with the substance of the AI use itself. Years later, memory may create the impression of involvement simply because you were present. Accountability is interested in something different. It examines whether there is evidence showing what you actually did with your attention.

That is why the wording of this question matters. Personal participation is not the same thing as formal approval. It means engaging with the proposed use before the decision is made. It means examining what the system will actually do, considering where human judgement ends and automated activity begins, and satisfying yourself that accepted risks have genuinely been understood rather than merely recorded.

The distinction becomes important when decisions are later challenged. Organisational processes may explain how approval occurred, but they do not necessarily demonstrate that you exercised your own judgement. If the surviving evidence shows only that you attended meetings and received information prepared by others, there may be little evidence that you independently reviewed the AI use at all.

Accountability does not require directors to become AI engineers. It requires them to demonstrate that they engaged with the decision in a meaningful way. The issue is not whether the organisation reviewed the proposal. The issue is whether there is evidence that you personally read, considered or questioned something material about the AI system before accepting the decision. Without that evidence, the distinction between exercising judgement and simply trusting the process becomes increasingly difficult to prove.

### **Question 2: Can you show what information you actually read or considered before approving, accepting or relying on the AI system?**

Most organisations can prove that information existed. They can produce board packs, reports, presentations and meeting papers. They can show when documents were circulated and who received them. What they often struggle to prove is something much more specific: what individual directors actually read before making a decision.

That distinction becomes important when an AI decision is challenged months or years later. At that point, the existence of information is no longer the central issue. The issue is whether there is evidence connecting that information to your judgement. The fact that a document sat in a board portal or appeared in a meeting pack does not demonstrate that it influenced your decision, or even that you read it at all.

This is where many directors become exposed without realising it. After a failure, it is natural to say that the relevant information was available at the time. Unfortunately, availability and consideration are not the same thing. An investigator, insurer or claimant is unlikely to be interested in the thousands of pages held by the organisation. They will focus on a much narrower question: what information can you show that you personally reviewed before accepting the risk?

That is why the word "actually" matters. It removes assumptions. It removes the belief that because a document was sent, it must have been read. It removes the comfort of general awareness. The question is not whether information existed somewhere inside the organisation. The question is whether there is evidence that it formed part of your decision-making.

The consequences can be surprisingly severe. A board pack may contain detailed discussion of the AI system's limitations, known weaknesses, testing gaps or failure modes. Yet if there is no evidence that you reviewed those materials, you may later appear to have approved the system without understanding those issues at all. The information existed, but your connection to it cannot be demonstrated.

Accountability does not require directors to read everything. It requires them to be able to show that they considered the information a reasonable person would have needed before accepting the risk. Without that evidence, informed judgement becomes difficult to distinguish from blind reliance on the views of others.

**Question 3: Can you show that you understood the known limits, weaknesses and risks of the AI system at the time you allowed it to be used?**

Many directors can show that risks were discussed. Risk registers were reviewed. Presentations were delivered. Assurances were provided by technical teams, suppliers and advisers. Yet none of those things demonstrate that the risks were actually understood.

That distinction matters because accountability is concerned with understanding, not exposure to information. An organisation may have identified dozens of risks associated with an AI system, but a list of risks is not the same thing as evidence that a director understood what those risks meant in practice.

The difference is often surprisingly simple. Being told that an AI system may produce errors is little more than a general warning. Every AI system produces errors. The statement is so broad that it provides almost no insight into the system itself. Understanding begins when the discussion moves beyond categories and into specific failures that could realistically occur.

A director who understands the limits of a system should be able to describe, in plain language, at least one situation where the system is likely to produce a wrong, misleading or harmful outcome. That description does not need to be technical. It does need to be specific. The issue is not whether you understand how the system works. The issue is whether you understand how it can fail.

This becomes important when decisions are later examined. Investigators rarely focus on the fact that risks were recorded somewhere within the organisation. Their attention is usually drawn towards the failure that actually occurred. They want to know whether that type of failure was known beforehand and whether those responsible for approving the system understood the possibility of it happening.

Directors sometimes worry about appearing technically uninformed. In practice, the greater danger is appearing unable to demonstrate any understanding of the system's known weaknesses. Accountability does not require technical expertise. It requires evidence that you understood enough about the system's behaviour to make a conscious decision about the risks being accepted.

That is why this question matters. If you cannot describe a specific failure that was known before approval, it becomes difficult to demonstrate that you understood the system's limits. If you cannot demonstrate that understanding, it becomes equally difficult to show that the risks were knowingly accepted.

**Question 4: Can you show that you asked questions, challenged assumptions or requested clearer answers before accepting what you were told about the AI system?**

Most directors believe they will remember the questions they asked. The difficulty is that accountability is not interested in memory. It is interested in evidence.

When decisions are challenged years later, investigators do not have access to what you were thinking at the time. They can only examine what survived. If there is no record of a question, a challenge, a request for clarification or an expression of concern, there may be no evidence that critical engagement occurred at all.

This is why the question matters. It is not testing whether you were difficult, sceptical or confrontational. It is examining whether your involvement left any visible trace of independent thought. A question about a risk. A challenge to an assumption. A request for additional information. Anything that demonstrates you did more than simply receive and accept what was presented.

The distinction may appear minor at the time. In practice, it can become highly significant. Most board discussions occur in an environment where trust is normal and disagreement is often unnecessary. Experienced presenters appear confident. Reports look complete. Assurances sound convincing. Under those conditions, passive acceptance can feel entirely reasonable.

The problem emerges later.

When an AI decision is challenged, investigators are rarely interested in how persuasive the presentation appeared at the time. They are interested in whether directors exercised judgement before accepting the recommendation. If there is no evidence that questions were asked or assumptions were examined, the conclusion becomes difficult to avoid. The recommendation was accepted without meaningful scrutiny.

Accountability does not require directors to oppose every proposal. Nor does it require technical expertise. It requires evidence that they engaged with the decision rather than simply observing it. A record of a question often reveals far more about the quality of a director's judgement than a record of their attendance.

That is why silence creates risk. Not because silence proves recklessness, but because silence leaves little evidence that independent judgement was exercised before the decision was made.

**Question 5: Can you prove and explain why it was reasonable for you to rely on the people, teams, advisers or suppliers who gave you information about the AI system?**

No director operates alone. Modern organisations depend on specialists, advisers, suppliers and technical teams. AI systems make that dependence even greater because many of the most important questions involve knowledge that sits outside the experience of most boards.

For that reason, reliance is not the problem. Every director relies on other people. The real issue is whether that reliance can still be justified later when decisions come under scrutiny.

The distinction matters because accountability does not stop at identifying who provided the information. It also examines why that information was trusted in the first place. The fact that somebody held a senior title, worked for a respected supplier or spoke with confidence does not automatically make reliance reasonable. Those facts may form part of the explanation, but they are not the explanation itself.

This becomes important after an AI-related failure. At that point, attention often turns to the information that was available before the decision was made. If that information proves incomplete, inaccurate or misleading, the question naturally follows. Why was it accepted? What basis existed for trusting the source? What evidence shows that the reliance was reasonable at the time?

The answer cannot simply be that the information came from an expert. Expertise is contextual. A supplier may understand its product but know little about the environment in which it is deployed. An adviser may be highly qualified yet lack experience with the particular technology being used. An internal team may be capable and well intentioned while still operating beyond its experience. Reasonable reliance depends upon what was known at the time, not what titles appeared on an organisational chart.

That is why this question focuses on justification rather than delegation. Delegating a task does not delegate accountability. Directors remain responsible for demonstrating why reliance made sense given the circumstances they faced.

When investigators, insurers or claimants examine a decision, they are rarely interested in whether somebody else provided the information. They are interested in whether a reasonable director, with the same information available at the time, would have relied on that source in the same way.

The evidential challenge is therefore straightforward. You need to be able to show who you relied upon and why that reliance was justified. If the first can be demonstrated but the second cannot, the decision begins to look less like informed judgement and more like acceptance without scrutiny.

### **Question 6: Can you show whether any warnings, objections, unresolved concerns or internal doubts reached you before the AI system was approved or used?**

Most investigations begin with a simple question: what was known before the decision was made?

The answer matters because accountability is judged using information that existed at the time, not information discovered afterwards. Once an AI system has failed, caused harm or attracted scrutiny, organisations naturally produce explanations. Reports are commissioned. Reviews are conducted. Lessons are identified. None of that changes what was known before the original decision was taken.

This is where warnings become important.

Many directors assume that warnings arrive in formal reports or clearly labelled risk assessments. Sometimes they do. More often they appear in less obvious forms. A concern raised during a meeting. An unanswered question in a paper. A qualification hidden in a presentation. A hesitant observation from a risk manager, technical specialist or junior employee. Individually these moments may appear insignificant. Collectively they often form the earliest signs that something is wrong.

The evidential problem is not simply whether those warnings existed. It is whether there is any record showing that they reached you.

That distinction becomes critical later. If a warning reached you but there is no evidence that it did, demonstrating what you knew at the time becomes much harder. Equally, if a warning clearly reached you but there is no evidence showing how you responded, others may be left to draw their own conclusions about what happened next.

This is why accountability depends upon more than receiving information. It depends upon preserving evidence that the information was received and considered. A note, a follow-up question, a request for clarification or a record that the concern was discussed may later become the only evidence connecting a known warning to the decision that followed.

Years later, memories will compete with each other. Explanations will change. Organisational narratives will evolve. The surviving evidence will remain. That evidence is often the only way to demonstrate what concerns were visible before approval and what was done once they became known.

For that reason, this question is not really about warnings. It is about whether there is evidence showing what reached you before the decision was made and whether your response to those concerns can still be reconstructed afterwards.

### **Question 7: Did you have the real authority to stop, delay, limit or question the use of the AI system if the risks were not properly explained?**

Most directors assume that accountability and authority travel together. If a person is expected to oversee a decision, it seems reasonable to assume they also possess the power necessary to influence that decision. In practice, the relationship is often less clear.

Organisation charts, committee structures and terms of reference can create the appearance of authority. A director may be asked to review proposals, approve recommendations and oversee implementation. Yet those formal powers do not always reflect reality. Projects gather momentum. Contracts are signed. Timelines become fixed. Expectations become embedded. By the time concerns are raised, the practical ability to alter the outcome may already have disappeared.

This question examines the difference between formal authority and real authority.

The distinction becomes important because accountability is ultimately concerned with influence over outcomes. A director who could genuinely stop, delay or limit a decision occupies a very different position from a director whose approval was expected but whose objection would have changed nothing.

Many directors discover this tension only when decisions are challenged later. The organisation points to the approval process and shows that oversight existed. The director points to concerns that were raised or reservations that were held. The difficulty arises when those two positions collide. If meaningful concerns existed, why did the decision proceed? If the decision would have proceeded regardless, what authority actually existed in the first place?

The issue is not whether a legal power appeared on paper. The issue is whether an objection could have altered the outcome in practice. Could additional information have been demanded? Could approval have been delayed? Could deployment have been restricted? Could the decision have been stopped entirely until concerns were addressed?

Those questions matter because authority forms part of the foundation of accountability. It is difficult to claim oversight of something that could not be influenced. Equally, it is difficult to explain responsibility for a decision if the practical ability to affect that decision never existed.

For that reason, this question is not really about governance structures or formal powers. It is about whether the authority associated with the role existed in reality and whether that authority could have been exercised when it mattered most.

### **Question 8: Can you prove and show that you reached your own judgement about the AI system, rather than simply approving what management, advisers or suppliers recommended?**

Every significant decision arrives accompanied by recommendations. Management presents a proposal. Advisers provide opinions. Suppliers explain benefits and risks. Technical teams offer assessments. By the time a matter reaches a director for approval, much of the analytical work has usually been completed by other people.

There is nothing unusual about that. Modern organisations could not function any other way, the difficulty arises later when decisions come under scrutiny.

At that point, the focus shifts away from the recommendation and towards the decision itself. Management may be asked why it supported a particular course of action. Advisers may be asked why they provided a particular opinion. Suppliers may be asked why they made particular claims. The director, however, faces a different question. Why did you accept the recommendation?

That distinction sits at the heart of this question.

Agreement is not the same thing as judgement. A director may ultimately reach exactly the same conclusion as management, advisers or suppliers. Accountability does not require disagreement. It does require evidence that the conclusion was reached through the exercise of independent judgement rather than simple acceptance.

This becomes important because recommendations are not a substitute for accountability. If a decision later proves controversial, harmful or mistaken, pointing back to the recommendation rarely resolves the issue. The recommendation explains what others thought should happen. It does not explain why you accepted it.

For that reason, investigators, insurers and claimants often look for evidence that a director formed a view of their own. That evidence may appear in questions that were asked, concerns that were raised, alternatives that were considered or conditions that were attached to approval. The form matters less than the underlying principle. There must be some indication that the recommendation was evaluated rather than merely received.

Many directors assume that attending meetings, reading papers and participating in discussions automatically demonstrates judgement. It does not. Those activities may show involvement. They do not necessarily show independent thought. A director can be present throughout the entire process and still contribute nothing beyond agreement.

That is why this question matters. If there is no evidence showing how you moved from another person's recommendation to your own conclusion, it becomes difficult to demonstrate that independent judgement was exercised at all. The decision may then appear to belong to those who made the recommendation rather than to the person who approved it.

### **Question 9: If you had concerns, doubts or reservations about the AI system, is there a record showing what you raised and how those concerns were handled?**

Not every concern begins as a formal objection. Many start as something much smaller. A feeling that a claim sounds too confident. A suspicion that a risk has been downplayed. A doubt about whether the system will perform as expected outside a controlled environment. These moments are common in decision-making. What makes them important is not whether they ultimately prove correct. What matters is whether there is any evidence that they existed.

This question examines the difference between a concern that was experienced and a concern that can be demonstrated.

That distinction becomes important when decisions are challenged months or years later. By that stage, memory is rarely enough. People remember events differently. Discussions become blurred. Recollections evolve over time. A director may genuinely remember having reservations, but accountability is concerned with what can be shown rather than what can be recalled.

For that reason, a concern only becomes visible when it leaves a trace. An email, a note, a question in a meeting, a request for clarification or a record in the minutes can all demonstrate that a concern existed before the decision was made. Without that evidence, the concern remains locked inside memory and may be impossible to distinguish from hindsight.

The second part of the question is equally important. Accountability does not end when a concern is raised. It also examines what happened afterwards. Was the concern investigated? Was additional evidence provided? Was the issue addressed, rejected or simply ignored? The answer forms part of the evidential record surrounding the decision.

This matters because concerns often become highly significant after an AI-related failure. The existence of a concern demonstrates that somebody recognised a potential problem before the event occurred. The way that concern was handled may then become one of the most important pieces of evidence available when reconstructing the decision later.

Many directors avoid creating such records because they do not wish to appear obstructive, pessimistic or difficult. Unfortunately, the absence of a record rarely produces that impression afterwards. Instead, it creates a different problem. If there is no evidence that concerns were raised, there may be little evidence that concerns existed at all.

That is why this question matters. It is not asking whether you had concerns. Most directors occasionally do. It is asking whether those concerns can still be demonstrated today and whether there is evidence showing how they were addressed before the decision was taken.

### **Question 10: Before approval, were you told about any missing records, incomplete testing, unclear responsibility or unresolved uncertainty about how the AI system would behave?**

Every significant decision is made with incomplete information. Records may be missing. Testing may still be underway. Responsibilities may not be fully defined. Questions may remain unanswered. AI systems are no different. In many cases, uncertainty is not a sign of failure. It is simply part of reality.

The difficulty begins when uncertainty becomes invisible.

Most approval processes are designed to present information that appears complete, coherent and ready for decision. Risks are summarised. Findings are presented. Recommendations are made. What receives far less attention are the limits of current knowledge. What has not been tested. What cannot yet be demonstrated. What remains uncertain.

This question examines whether those limitations were visible before approval.

The distinction matters because directors cannot evaluate information they never received. If important uncertainties were not disclosed, they could not form part of the decision-making process. A decision made without knowledge of those uncertainties may appear reasonable at the time, yet still rest upon an incomplete understanding of the situation.

The issue becomes more significant after an AI-related failure. Investigators rarely stop at identifying what went wrong. They also examine what was known before the decision was made. If the failure can be traced to an area of uncertainty that existed beforehand, attention naturally turns to whether that uncertainty was disclosed, discussed and understood.

This creates an uncomfortable position for directors. If an important gap was never disclosed, they may be unable to show that they knowingly accepted the associated risk. Yet the absence of disclosure does not entirely remove scrutiny. Questions may still arise about whether sufficient effort was made to identify what remained unknown before approval was granted.

That is why this question focuses on uncertainty rather than error. The issue is not whether every question had an answer. Few complex decisions offer that luxury. The issue is whether there is evidence showing where knowledge ended, where uncertainty began and whether those boundaries were visible before the decision was taken.

A director cannot make an informed decision about information that does not exist. They can, however, be expected to understand what information is missing. The difference between those two positions often becomes critical when decisions are reconstructed years later.

**Question 11: Could you still prove and show, five years later, what you knew, what you considered, what you questioned and why you accepted the decision at the time?**

Most decisions are made in the present and judged in the future. That gap creates a problem that many directors underestimate. The information is available when the decision is taken. The discussions are fresh. The reasoning feels obvious. The people involved understand why the decision was made. Years later, much of that context has disappeared.

Claims, investigations and regulatory actions rarely arrive immediately. They emerge after an event has occurred, after concerns have been raised or after harm has become visible. By then, people have moved on, systems have changed, records have been archived and memories have begun to merge recollection with hindsight.

This question examines whether the decision can still be reconstructed when that happens.

The distinction is important because accountability depends upon evidence rather than recollection. A director may sincerely remember what they knew, what information they considered, what concerns they raised and why they accepted a particular course of action. The difficulty is that memory cannot normally be tested, verified or examined by others. Evidence can.

That is why each part of the question matters. It is asking whether there is a surviving record of what information was available to you, what information influenced your thinking, what issues you challenged and why you ultimately accepted the decision. If those elements cannot be reconstructed from evidence that exists independently of memory, they become increasingly difficult to demonstrate as time passes.

Many directors assume that organisational records will solve this problem. Sometimes they help. More often they explain what happened rather than why it happened. Minutes record decisions. Papers record recommendations. Approval records show outcomes. None of those necessarily reveal what a particular director knew, considered or questioned before reaching a conclusion.

This is where personal accountability and organisational governance begin to separate. The organisation may be able to show that a decision was made. The director may still be unable to show how their own judgement was exercised.

The issue therefore is not whether you remember the decision today. The issue is whether somebody who knew nothing about the decision could reconstruct your involvement five years from now using only the evidence that survives.

If the answer is no, then a significant part of your defence depends upon memory. Memory fades, changes and can be challenged. Evidence survives. That is why this question sits at the centre of the entire test.

**Question 12: Can you prove and show that known AI risks, unresolved concerns or material uncertainties were not hidden from insurers when cover was requested or renewed?**

Insurance decisions are made using information. Insurers assess risk, set premiums and decide whether to provide cover based on what they are told at the time. The quality of those decisions therefore depends entirely upon the quality of the information available to them.

That creates a question many directors never consider. How closely does the insurer's understanding of the risk match the organisation's understanding of the risk?

The distinction matters because information does not always travel cleanly through an organisation. Concerns may be known within a project team but never reach the board. Issues discussed at board level may never appear in renewal discussions. Uncertainties identified during testing may remain trapped inside technical reports. As information moves between functions, some of it inevitably disappears.

This question examines whether anything important was lost along the way.

The issue is not whether every possible risk was disclosed. Few organisations possess complete knowledge of every future problem. The issue is whether known concerns, material uncertainties or significant unanswered questions were visible inside the organisation but absent from the information provided to insurers.

That becomes important when claims arise.

If an insurer later discovers that information material to its assessment was known within the organisation but never reached the renewal process, attention quickly turns to how that happened. What was known? Who knew it? Was it communicated? If not, why not?

Many directors assume this is somebody else's responsibility. Brokers manage renewals. Risk teams provide information. Legal teams review disclosures. Yet accountability rarely stops with the people who completed the paperwork. If material concerns were known before cover was requested or renewed, questions may arise about whether those responsible for oversight satisfied themselves that the insurer was receiving an accurate picture of the risk.

This is why the question focuses on evidence rather than process. It is not asking whether an insurance renewal took place. It is asking whether there is evidence showing that known risks, unresolved concerns and material uncertainties were reflected in the information provided to the insurer.

The protection offered by insurance depends upon the quality of the information supporting it. If the insurer was making decisions using a materially different understanding of the risk from the organisation itself, the value of that protection may become uncertain at precisely the moment it is needed most.

**Question 13: If this AI decision were challenged today, would the surviving evidence make your own conduct appear careful, informed and reasonable?**

Every question in this test ultimately leads to the same place. The earlier questions examine specific pieces of evidence. What information you reviewed. What risks you understood. What questions you asked. What concerns reached you. Whether you exercised your own judgement. Whether those things can still be reconstructed years later.

This question brings all of those issues together and asks a single, unavoidable question: what picture does the surviving evidence create?

That distinction matters because accountability is rarely judged using the information that once existed. It is judged using the information that still exists. Documents disappear. People leave. Memories fade. Explanations change. What remains is the evidential record available to those examining the decision long after the event itself.

At that point, the issue is no longer whether you believed you were acting responsibly. Most directors do. Nor is the issue whether you intended to act carefully. Most directors do that as well. The issue is whether an independent observer, looking only at the surviving evidence, would reach the same conclusion.

That is why this question focuses on conduct rather than outcome. AI systems can succeed for the wrong reasons and fail despite sensible decision-making. Accountability is therefore less concerned with whether the decision produced a good result and more concerned with whether the decision was reached in a careful, informed and reasonable manner.

The challenge is that personal confidence and evidential appearance are not the same thing. A director may genuinely remember reading the papers, considering the risks and asking difficult questions. Yet if there is no surviving evidence of those activities, they become increasingly difficult to demonstrate. The record may then tell a very different story from the one held in memory.

This is the point at which organisational governance and personal accountability finally separate. The organisation may be able to show that meetings were held, processes were followed and approvals were obtained. None of that automatically demonstrates how a particular director exercised judgement. That question can only be answered by evidence linked to the individual.

Viewed from that perspective, this is not simply the final question in the test. It is the reason for the test.

If an investigator, regulator, insurer or claimant examined the surviving evidence today, would they see the conduct of a careful, informed and reasonable director? Or would they see a person who approved a decision without leaving any meaningful evidence of their own involvement, understanding or judgement?

The answer to that question is the answer that matters most, because it is the question that all the others were designed to reach.

## About the Author

Not anti-AI, just honest about what can be proved later.

Russell Parrott writes independent explanatory material about AI accountability focused on one practical question: what can still be proven after an automated decision is challenged.

He treats accountability as an evidence problem rather than a policy statement or technical aspiration. It examines whether a specific outcome can be shown from records and evidence that already exist, whether that outcome can be linked to a named person or body with authority and whether the basis for the decision can still be demonstrated afterwards.

Because eventually somebody asks: who decided, why and where is the evidence?

Contacts:

- Website: <https://russellparrott.eu>
- LinkedIn: <https://www.linkedin.com/in/russell-parrott>
- Email: [parrott.russell@gmail.com](mailto:parrott.russell@gmail.com)