

ILII Public Case Study Series

PCS-001

Artificial Intelligence and the United Kingdom's Asylum Decision-Making Process

Applying the Institutional Legitimacy & Influence Index (ILII) Examination Methodology

Neil Sumaru
Institutional Governance Research Initiative

Published by Levuka Venture Lab

July 2026

This Public Case Study Examination applies the ILII methodology using publicly available information. It is intended to demonstrate the investigative methodology and does not constitute an institutional diagnosis, audit, or legal determination.

Public Case Study Examination

Current version: Version 1.1

Originally published: July 2026

Version 1.1: August 2026. Publication-format revision only; trademark markings removed. No substantive changes to the analysis, evidence, findings, or methodology.

Previous version: Version 1.0

© 2026 Neil Sumaru. All rights reserved.
Published by Levuka Venture Lab
This publication may be cited with appropriate attribution.

Abstract

Artificial intelligence is increasingly influencing institutional decision-making in legally and socially consequential domains. Existing analyses have largely examined these developments through the lenses of operational performance, legal compliance, technical assurance, explainability and implementation controls. These perspectives provide important insight into how AI systems are developed, deployed and governed.

This investigation examines a different question.

Using publicly available evidence relating to the United Kingdom’s asylum process, it applies the Institutional Legitimacy & Influence Index (ILII) Public Case Study Examination Methodology to ask whether the governance position from which institutional authority derives its legitimacy remains coherently reflected throughout decision formation as AI participates more directly in that process.

This examination does not seek to determine whether the Home Office’s use of AI is lawful or unlawful, nor does it evaluate the technical performance of the systems involved. It investigates whether publicly observable evidence suggests that governance coherence is becoming harder to demonstrate as institutional decision environments evolve.

The governance position is established first as the independent object of examination. The investigation then considers whether the historical assumptions that once enabled consequential decisions to be reconstructed, explained and defended continue to hold as AI participates in the formation of institutional judgment.

The examination identifies a cumulative pattern rather than a single point of failure. Publicly observable evidence suggests that direct engagement with evidence, human synthesis, identifiable judgment and reconstructability each become less secure under the documented conditions of AI-assisted decision formation. Human accountability formally remains in place, which weighs against any claim that accountability has disappeared or been transferred to AI. The concern is narrower: accountability may become harder to substantiate where the decision process beneath it becomes less reconstructable.

This case study demonstrates how a governance-focused investigative methodology can complement legal, operational and technical analyses by examining a structural question that has received comparatively little attention: whether institutions can continue to demonstrate the coherent exercise of authority as artificial intelligence increasingly shapes how consequential decisions are formed.

1. The Public Record

The use of artificial intelligence within the United Kingdom’s asylum decision-making process has become the subject of public, legal and institutional scrutiny.

In April 2025, the Home Office published an evaluation of two AI trials within the asylum decision-making process: the Asylum Case Summarisation tool and the Asylum Policy Search tool. The Asylum Case Summarisation tool was designed to summarise asylum interview transcripts. The Asylum Policy Search tool was designed to find and summarise country policy information. The Home Office described both tools as aids for decision-makers, not replacements for any part of the decision-making process.¹

The stated purpose of these tools was to reduce cognitive load, improve efficiency and shorten the time required for case review. The evaluation reported potential time savings in the pilot environment, including an average saving of 23 minutes per case for transcript review using the Asylum Case Summarisation tool and 37 minutes per case when searching for country policy information using the Asylum Policy Search tool.²

These pilots took place against a wider asylum system under operational pressure. Public reporting and inspection material describe large application backlogs, staff turnover, recruitment challenges and demands to accelerate decision-making.³ These conditions matter because asylum decisions depend on the careful interpretation of testimony, credibility, supporting documentation and country information. When AI-assisted summarisation is introduced into that environment, the governance question is not only whether the tools function as intended. It is also whether the institution can continue to show how judgment was formed.

¹ Home Office, *Evaluation of AI Trials in the Asylum Decision Making Process* (London: Home Office, 29 April 2025), <https://www.gov.uk/government/publications/evaluation-of-ai-trials-in-the-asylum-decision-making-process/evaluation-of-ai-trials-in-the-asylum-decision-making-process>

² Home Office, *Evaluation of AI Trials in the Asylum Decision Making Process*.

³ Independent Chief Inspector of Borders and Immigration, *An Inspection of Asylum Casework: June–December 2025* (London: Independent Chief Inspector of Borders and Immigration, 25 June 2026), <https://www.gov.uk/government/news/inspection-report-published-an-inspection-of-asylum-casework-june-december-2025>; Open Rights Group, “How AI in Asylum Decision-Making Drives Failure and Cost,” Open Rights Group, accessed July 1, 2026, <https://www.openrightsgroup.org/blog/how-ai-in-asylum-decision-making-drives-failure-and-cost>

The Open Rights Group commissioned an independent legal opinion examining the Home Office's use of AI within refugee status determination. The opinion raised questions concerning transparency, procedural fairness, meaningful human oversight, AI assurance and the adequacy of safeguards surrounding AI-generated summaries within asylum decision-making. It also considered the relationship between these practices and existing public law, human rights obligations and the UK Government's AI Playbook.⁴

The public record is not one-sided. The Home Office evaluation reported that the pilot tools did not appear to negatively or positively affect decision quality during the pilot. It also stated that the tools were designed to preserve human decision-maker autonomy. At the same time, the evaluation identified limitations. The pilots were small-scale, the decision-making units were not selected to be representative, and the methodologies relied in part on self-reporting.⁵

Alongside the legal analysis, broader public commentary has raised concerns that pressure to improve processing speed, combined with AI-assisted summarisation, staff turnover and performance targets, could affect decision quality while increasing downstream appeals, institutional costs and pressure on the wider asylum system.⁶

Taken together, this public record provides a suitable basis for examination. The case presents a visible institutional setting in which AI-assisted synthesis, human judgment, operational pressure and public accountability intersect.


⁴ Open Rights Group, *Legal Opinion: The Use of Artificial Intelligence Tools in Asylum Cases* (London: Open Rights Group, 16 March 2026), <https://www.openrightsgroup.org/publications/legal-opinion-the-use-of-artificial-intelligence-tools-in-asylum-cases>; Government Digital Service, *AI Playbook for the UK Government* (London: Cabinet Office, 10 February 2025), <https://www.gov.uk/government/publications/ai-playbook-for-the-uk-government>.

⁵ Home Office, *Evaluation of AI Trials in the Asylum Decision Making Process*.

⁶ Open Rights Group, "How AI in Asylum Decision-Making Drives Failure and Cost," Open Rights Group, accessed July 1, 2026, <https://www.openrightsgroup.org/blog/how-ai-in-asylum-decision-making-drives-failure-and-cost/>.

Figure 1. Publicly Documented Operational Environment Surrounding UK Asylum Decision-Making

Redrawn from Open Rights Group, "How AI in Asylum Decision-Making Drives Failure and Cost" (2026).



Source: Redrawn from Open Rights Group, "How AI in Asylum Decision-Making Drives Failure and Cost" (2026), <https://www.openrightsgroup.org/blog/how-ai-in-asylum-decision-making-drives-failure-and-cost/>.

2. Methodological Foundation

Every governance system rests on a source of legitimacy. Operational rules, controls and procedures may structure how decisions are made, but they do not create the authority to decide. They derive their authority from the governance position they are intended to preserve.

For the purposes of this examination, the governance position refers to the substantive principles, values and sources of authority from which an institution derives the legitimacy to exercise consequential decision-making.

Historically, this distinction was easier to preserve. Consequential decisions were formed through slower human processes. The interpretation of evidence, the exercise of judgment and the justification of decisions remained closely coupled to those ultimately accountable for institutional authority.

Artificial intelligence changes that relationship. As AI participates in summarisation, synthesis and decision support, the assumptions that once sustained governance coherence can no longer be taken

for granted. The question is not simply whether operational controls exist. It is whether those controls still preserve the underlying governance position.

This examination therefore begins upstream.

Rather than evaluating operational controls, AI models or implementation decisions in isolation, this examination first reconstructs the governance position that historically legitimised the exercise of institutional authority. It then examines whether that position remains coherently reflected throughout the decision process as artificial intelligence participates in consequential decision formation.

Because this is a public case study examination rather than an internal institutional audit, the reconstructed governance position operates as a provisional public baseline. It is derived from publicly available evidence and used to examine whether governance coherence remains demonstrable under documented conditions of AI-assisted decision formation. It should not be read as a definitive finding about every historical asylum decision.

The relationship between the governance position, institutional authority, operational controls and consequential decision formation is illustrated conceptually in Figure 2.

This sequencing is not a matter of investigative preference. It is a methodological necessity.

Methodological Principle

Operational controls derive their authority from the governance position they are intended to preserve. If the determination of that position were delegated to the implementation being governed, the distinction between the source of legitimacy and its implementation would begin to collapse. Operational controls may help preserve legitimacy, but they cannot become the basis upon which institutional authority is justified.

This is not simply a technical concern. It is an epistemological one.

Institutional consequences continue to terminate in the human domain. The asylum seeker experiences the decision. The minister answers Parliament. The judge reviews the process. The institution remains accountable for the exercise of its authority regardless of how extensively artificial intelligence participates in forming the decision.

For this reason, the determination of governance coherence remains a human responsibility. The examination asks whether the substantive legitimacy from which institutional authority is derived remains coherently reflected throughout the decision process, and whether that coherence can still be reconstructed, explained and defended by those ultimately accountable for its consequences.

Governance Drift is not presumed from the presence of artificial intelligence. It may be identified only where the cumulative evidence shows that the governance position has become harder to locate, reconstruct or defend. Where a historical governance assumption remains demonstrably intact, that finding counts against drift on that dimension. A non-drift finding would require the public record to show that AI participation has not materially weakened the institution's ability to demonstrate the five reconstructed governance assumptions examined below.

3. Historical Governance Assumptions

The purpose of this investigation is not to determine whether the United Kingdom's asylum system historically produced perfect decisions. No consequential public decision-making system does.

The relevant question is narrower. Before examining the effect of AI-assisted decision support, the investigation must reconstruct the assumptions that historically allowed asylum decisions to be formed, explained and defended as exercises of institutional authority.

These assumptions do not describe how every decision was made in practice. They identify the governance conditions against which AI-assisted decision formation can be examined.

For an assumption to continue to hold, the institution must remain able to demonstrate that the relevant condition is still present in the decision process. Where that condition becomes harder to locate, evidence or reconstruct, the assumption becomes less secure as a basis for governance coherence.

Based on the public record surrounding this case, five historical governance assumptions appear to have underpinned the United Kingdom's asylum decision-making process.

These assumptions are not the proprietary scoring dimensions of the full ILII diagnostic architecture. They are case-specific governance assumptions reconstructed for purposes of this public examination. The purpose is to identify the conditions that would need to remain demonstrable for institutional authority to remain coherent in this decision environment.

3.1 Direct evidentiary engagement

Asylum caseworkers were expected to engage directly with applicant testimony, supporting documentation and country information before reaching a decision. The act of reviewing evidence formed part of the reasoning process itself rather than simply documenting an already completed judgment.⁷

3.2 Human evidentiary synthesis

The production of summaries historically formed part of the decision-maker's own reasoning. Producing a summary was not merely an administrative step. It helped the caseworker organise, interpret and understand the evidence before exercising institutional authority. The legal opinion characterises this function directly, observing that producing a summary is “not done for record keeping alone, but also as a means for the officer to order their thinking.”⁸

3.3 Identifiable human judgment

Although asylum decisions were governed by legislation, policy and procedural safeguards, consequential judgment remained attributable to identifiable human decision-makers. Responsibility for interpreting evidence, applying policy and reaching a determination could generally be located within the institutional decision process.

3.4 Reconstructable decision formation

Where decisions were challenged, the process through which judgment had been formed could ordinarily be reconstructed through the evidence considered, the reasoning applied and the individuals responsible for exercising institutional authority.

3.5 Human institutional accountability

Regardless of procedural complexity, the institution ultimately remained accountable for the exercise of its authority. Caseworkers, supervisors, ministers and courts occupied identifiable governance roles through which institutional decisions could be explained, reviewed and defended when challenged.

⁷ Open Rights Group, *Legal Opinion: The Use of Artificial Intelligence Tools in Asylum Cases*.

⁸ Open Rights Group, *Legal Opinion: The Use of Artificial Intelligence Tools in Asylum Cases*.

These assumptions do not establish that every historical asylum decision satisfied these standards in practice. Rather, they identify the governance architecture through which institutional authority was expected to remain coherent.

Figure 3. Reconstructing Historical Governance Assumptions

Historical governance assumptions that underpinned the UK asylum decision-making process prior to the introduction of AI-assisted decision support.

Historical Governance Assumption	Description	Why It Historically Mattered	What It Enabled Institutions to Do
1  Direct evidentiary engagement Decision-makers directly engage with applicant testimony, supporting documents, and country information.	Asylum decisions depend on first-hand understanding of evidence. Nuance, context, and credibility signals are captured through direct review.	Direct engagement allows decision-makers to observe, question, and interpret evidence that cannot be fully captured through intermediated summaries or third-party interpretations.	Form independent, grounded understanding of the evidence, rather than relying on indirect representations.
2  Human evidentiary synthesis Decision-makers personally interpret and synthesise evidence to form judgment.	Synthesis is a reasoning activity, not a clerical task. It shapes understanding, highlights key patterns, and informs judgment.	Human synthesis enables decision-makers to connect evidence, assess credibility, and apply institutional knowledge and experience.	Ensure that reasoning reflects human judgment, not merely the output of automated or system-generated summaries.
3  Identifiable human judgment Consequential judgment is exercised by accountable human decision-makers and makes the decision.	Institutional authority resides in human decision-makers who accept responsibility for the outcome.	Identifiable judgment ensures that authority is exercised by accountable humans across the decision chain.	Locate authority and ensure that decisions are attributable to identifiable individuals.
4  Reconstructable decision formation The reasoning process can be reconstructed and explained when reviewed and questioned.	Decisions must be capable of explanation. The path from evidence to conclusion must be transparent and reviewable.	Reconstructability enables institutions to demonstrate how conclusions were reached and why decisions were made in a particular way.	Explain and defend decisions consistently when challenged through appeals, reviews, or oversight.
5  Human institutional accountability Accountability for outcomes rests with the institution and the humans who exercise judgment.	Institutions remain responsible through human officeholders, supervisors, ministers, and courts.	Human accountability ensures that consequences remain with those who hold authority, not with systems, vendors, or automated processes.	Ensure that the institution can be held responsible, and that remediation and redress are possible when decisions are lawfully challenged.


These historical governance assumptions formed the foundation of institutional authority in asylum decision-making.
 The examination in this case study asks whether these assumptions continue to hold as AI increasingly participates in decision formation.

Note: These assumptions are reconstructed from public descriptions of the UK asylum system prior to the introduction of AI-assisted decision support.

Figure 3 summarises the historical assumptions reconstructed from publicly documented descriptions of the UK asylum decision-making process. These assumptions establish the baseline examined throughout the remainder of this case study.

4. Examination of Historical Governance Assumptions

The five historical governance assumptions reconstructed in Section 3 provide the baseline for this examination. Each assumption is now assessed against the public record to determine whether it remains demonstrably coherent under the documented conditions surrounding AI-assisted asylum decision-making.

Authority must be justified through legitimacy. Operational controls exist to preserve that legitimacy, not to justify it. Their relevance in this examination depends on whether they help accountable

decision-makers maintain and demonstrate the connection between institutional authority and the process through which judgment is formed.

Governance Drift is not inferred from the presence of artificial intelligence, nor from the adequacy or inadequacy of any individual control. The examination considers each assumption separately. Where publicly observable evidence shows that a historical governance assumption remains demonstrably intact, that finding weighs against drift on that dimension. Where the public record shows that several assumptions remain demonstrably intact despite AI participation, the examination may support a non-drift finding. Where several assumptions become harder to locate, evidence or reconstruct, the cumulative pattern may indicate emerging Governance Drift.

Human institutional accountability is assessed separately, but it is also structurally dependent on the reconstructability of the decision process beneath it. A finding that formal accountability remains intact weighs against any claim that accountability has disappeared. It does not, by itself, resolve whether accountable actors can substantively reconstruct, explain and defend the judgment for which they remain responsible.

Figure 4 illustrates the investigative sequence applied in the examination: historical assumption, public evidence, governance observation, coherence assessment and diagnostic implication.

Figure 4. ILII Investigative Logic: From Assumption to Coherence Assessment

The ILII Public Case Study Examination Methodology assesses each historical governance assumption without presuming a uniform outcome.

Historical Governance Assumption	Public Evidence	Investigative Question	Coherence Assessment (Per Dimension)	Diagnostic Implication
1  Direct evidentiary engagement Decision-makers directly engage with applicant testimony, supporting documents, and country information.	Home Office evaluation of AI trials; policy descriptions; inspection findings; legal opinion; public reporting on asylum caseload and casework pressures.	Can direct engagement with primary evidence still be demonstrated as the basis for consequential asylum judgment under the current operational conditions?	 Intact Less Secure Weakened Not Demonstrable <i>Place assessment on the continuum</i>	Stronger coherence weighs against Governance Drift on this dimension. Weaker coherence increases exposure.
2  Human evidentiary synthesis Decision-makers personally interpret and synthesise evidence to form judgment.	Home Office evaluation; policy descriptions; casework observations; legal opinion; inspectorate findings.	Does human synthesis remain a substantive component of decision formation, or has it been partially displaced by AI-generated synthesis?	 Intact Less Secure Weakened Not Demonstrable <i>Place assessment on the continuum</i>	Reduced human synthesis increases the risk that reasoning is shaped upstream of accountable judgment.
3  Identifiable human judgment Consequential judgment is exercised by identifiable human decision-makers.	Home Office evaluation; policy statements; legal opinion; inspectorate findings; public reporting.	Can consequential judgment still be meaningfully attributed to identifiable human decision-makers?	 Intact Less Secure Weakened Not Demonstrable <i>Place assessment on the continuum</i>	Lower attribution confidence weakens accountability and increases governance exposure.
4  Reconstructable decision formation The reasoning process can be reconstructed and explained when reviewed and questioned.	Home Office evaluation; policy descriptions; legal opinion; inspectorate findings; public reporting on source reliance.	Can the reasoning pathway from evidence to conclusion still be independently reconstructed and explained when challenged?	 Intact Less Secure Weakened Not Demonstrable <i>Place assessment on the continuum</i>	Lower reconstructability increases the risk that decisions cannot be explained or defended consistently.
5  Human institutional accountability Accountability for outcomes rests with the institution and the humans who exercise judgment.	Home Office evaluation; policy statements; legal opinion; inspectorate findings; public reporting.	Is formal accountability maintained, and is it supported by a decision process that accountable actors can understand, evidence and defend?	 Intact Less Secure Weakened Not Demonstrable <i>Place assessment on the continuum</i>	Accountability formally intact reduces governance drift risk on this dimension, but may still be substantively exposed if the underlying process is less reconstructable.



A finding of coherence (Intact) weighs against Governance Drift on that dimension.

A cumulative pattern of weakened or less demonstrable assumptions may indicate emerging Governance Drift.

Note: This figure illustrates the ILII investigative logic. Actual assessments for this case study are presented in Figure 5.

5. Examination

This section applies the investigative sequence set out in Figure 4. Each reconstructed governance assumption is examined against the public record to determine whether it remains demonstrably coherent under the documented conditions surrounding AI-assisted asylum decision-making.

The examination does not ask whether every asylum decision has been affected in the same way. It asks whether the public record shows conditions under which the relevant governance assumption remains intact, becomes less secure, or becomes harder to evidence.

5.1 Direct evidentiary engagement

Historically, consequential judgment depended upon direct engagement with applicant testimony, supporting documentation and country information.

The Home Office evaluation describes the Asylum Case Summarisation tool as using AI to summarise asylum interview transcripts, while the Asylum Policy Search tool finds and summarises country policy information. Both tools were designed as aids for decision-makers rather than replacements for the decision-making process.⁹

The governance issue is not simply whether the summaries are accurate. It is whether the caseworker's first meaningful encounter with testimony, country information and evidentiary material may occur through a system-generated account of that material. In an asylum process, the difference matters. Testimony, credibility, inconsistency and country risk are not merely data points to be condensed. They are the materials through which judgment is formed.

On the public record, direct evidentiary engagement is not shown to have disappeared. The Home Office expressly describes the tools as aids. But where AI-generated summaries mediate the caseworker's practical engagement with the underlying record, the institution may face greater difficulty showing that judgment remained grounded in direct evidentiary review rather than system-assisted synthesis.

5.2 Human evidentiary synthesis

The production of summaries historically formed part of the decision-maker's own reasoning process. It was not merely clerical. It helped the caseworker organise, interpret and understand the evidence before exercising institutional authority.

The legal opinion captures this point directly, observing that summary production is “not done for record keeping alone, but also as a means for the officer to order their thinking.”¹⁰

AI-assisted summarisation changes the location of that work. If the system performs part of the organising and synthesising function before the caseworker engages the record, then part of the reasoning environment has already been shaped upstream. The question is not whether the caseworker remains formally responsible. The question is whether a component of judgment that historically occurred through human synthesis has been partially relocated into a system-generated process.

⁹ Home Office, *Evaluation of AI Trials in the Asylum Decision Making Process*.

¹⁰ Open Rights Group, *Legal Opinion: The Use of Artificial Intelligence Tools in Asylum Cases*.

This assumption therefore becomes less secure. Human synthesis may still occur, but the public record supports a governance concern that AI-assisted summarisation can substitute for, or materially shape, part of the reasoning function that historically belonged to the decision-maker.

5.3 Identifiable human judgment

The Home Office position is clear: the tools were designed to support decision-makers, not replace them. The evaluation also states that users did not report that the tools reduced autonomy or influenced decision-making during the pilots.¹¹

That evidence matters. It weighs against a simple conclusion that judgment has been transferred from human decision-makers to AI systems.

The governance concern is narrower. Judgment may remain formally assigned to the caseworker while becoming less clearly attributable in substance. Where AI-generated summaries, time pressure, performance expectations and reduced review capacity interact, the final decision may still be signed by a human decision-maker even though the conditions shaping that judgment are distributed across system outputs and institutional pressures.

For that reason, the identifiable exercise of judgment remains partly coherent but more difficult to locate with confidence. The public record does not show that human judgment has been displaced. It does suggest that the formation of judgment may become less visibly attributable to the human decision-maker alone.

5.4 Reconstructable decision formation

Reconstructability depends on more than a final decision record. It requires the institution to show how judgment was formed: what evidence was considered, how it was interpreted and how the decision-maker moved from evidence to conclusion.

The Home Office evaluation itself identifies limitations relevant to this assumption. In the ACS pilot, some users noted that the summary output did not reference where information could be found in the transcript. The evaluation also reported that 9 percent of summaries were deemed inaccurate or had missing information and were removed from the pilot, and that 23 percent of users were not fully confident in the summary information.¹²

¹¹ Home Office, *Evaluation of AI Trials in the Asylum Decision Making Process*.

¹² Home Office, *Evaluation of AI Trials in the Asylum Decision Making Process*.

The legal opinion raises a related reconstructability concern. It states that decision-makers may be unable to cross-reference AI-generated summaries back to the original source material in a reliable way, and that the prompt structure used to produce the summary was not disclosed. If the evidentiary account presented to the decision-maker cannot be readily traced back to the underlying record, reconstructability is weakened even where a final decision record remains available.

Those findings do not prove that decisions were wrongly made. They do, however, matter for reconstructability. If an AI-generated summary becomes part of the practical route through which the caseworker understands the record, then missing information, lack of source references or reduced confidence in the summary can make it harder to reconstruct how the decision-maker engaged with the underlying evidence.

This is the dimension where the governance concern is most direct. Even where formal decision quality is not shown to decline, the process by which judgment was formed may become less securely reconstructable.

5.5 Human institutional accountability

Human accountability formally remains in place. The institution remains responsible for the decision. The caseworker remains the decision-maker. Supervisory structures, ministerial answerability and judicial review do not disappear because AI has participated in summarisation or policy search.

This finding weighs against any conclusion that accountability has disappeared or been transferred from human institutions to AI systems. The assumption therefore remains formally intact.

The remaining question is whether accountable human actors can substantively reconstruct, explain and defend the decision process for which they remain responsible. If direct evidentiary engagement, human synthesis, identifiable judgment and reconstructability all become less secure, then accountability may remain located while becoming thinner in substance.

The concern is not that AI has become the accountable actor. It has not. The concern is that accountable human actors may be asked to stand behind a decision process whose formation has become less transparent to them and harder for the institution to evidence after the fact.¹³

¹³ Open Rights Group, *Legal Opinion: The Use of Artificial Intelligence Tools in Asylum Cases*; Open Rights Group, “How AI in Asylum Decision-Making Drives Failure and Cost”.

6. Synthesis of Findings

The examination does not identify a single point of governance failure, nor does it suggest that the use of AI necessarily invalidates institutional authority.

It identifies a cumulative pattern.

Across the first four reconstructed governance assumptions, the public record suggests that the conditions historically supporting coherent decision formation become less secure under AI-assisted summarisation and policy search. Direct evidentiary engagement becomes harder to evidence. Human synthesis becomes partially relocated into system-generated outputs. The exercise of judgment remains formally human, but becomes less clearly attributable in substance. Reconstructability becomes more fragile where summaries mediate access to testimony and country information.

The fifth assumption, human institutional accountability, differs from the others. Formal accountability remains human. The institution remains responsible. The decision-maker remains the decision-maker. That finding weighs against any claim that accountability has disappeared or been transferred to AI. The concern is narrower. Formal accountability becomes substantively exposed where the reasoning process beneath it becomes harder to reconstruct, explain and defend.

Figure 5. Emergent Governance Drift Pattern Across Historical Assumptions

Synthesis of Examination Findings

#	Historical Governance Assumption	Examination Finding (Public Record)	Governance Implication	Coherence of Assumption Under AI-Assisted Conditions
1	 Direct evidentiary engagement Decision-makers engage directly with applicant testimony, supporting documents and country information as part of reasoning.	AI summaries mediate first exposure to testimony and country information. Direct engagement remains intended but becomes harder to demonstrate in practice.	Direct review of primary evidence may be partially displaced by mediated accounts, making it harder to show judgment was grounded in the original record.	 LESS SECURE Harder to evidence direct evidentiary engagement.
2	 Human evidentiary synthesis Caseworkers organise and interpret evidence themselves as part of forming judgment.	AI performs part of the organising and synthesising function upstream of the decision-maker.	Reasoning environment is shaped by system outputs, reducing the extent to which synthesis occurs through human reasoning alone.	 WEAKENED Human synthesis partially relocated to system outputs.
3	 Identifiable human judgment Consequential judgment is exercised by identifiable human decision-makers.	Judgment remains formally human, but its formation is influenced by AI-generated outputs, time pressures and reduced review capacity.	Judgment may be less clearly attributable to the human decision-maker in substance, even where formal responsibility remains.	 LESS SECURE Harder to locate judgment with confidence.
4	 Reconstructable decision formation The path from evidence to decision can be reconstructed and explained when reviewed.	Summaries lack source references, may contain inaccuracies or omissions, and are not fully trusted by all users.	The reasoning chain from evidence to conclusion becomes more fragile and harder to reconstruct after the fact.	 WEAKENED Harder to reconstruct decision formation reliably.
5	 Human institutional accountability The institution remains accountable through human officeholders, supervisory roles and review processes.	Accountability structures remain in place. Decisions remain attributable to human decision-makers.	Formal accountability persists, but its substance may thin if the process beneath it becomes less transparent and harder to reconstruct.	 FORMALLY INTACT / SUBSTANTIVELY EXPOSED Accountability remains human, but may thin in substance.



Governance Drift may arise not because accountability disappears, but because the decision process beneath accountability becomes harder to locate, reconstruct and defend.

Key:



LESS SECURE



WEAKENED



FORMALLY INTACT / SUBSTANTIVELY EXPOSED

Note: Findings reflect the public record available in ORG's Legal Opinion and the Home Office Evaluation of AI Trials in the Asylum Decision-Making Process.

The investigation has not assessed the technical performance of artificial intelligence, nor has it determined whether individual operational controls are effective or ineffective. Rather, it has examined whether those controls appear to assist accountable decision-makers in preserving the governance conditions that legitimise the exercise of institutional authority.

On the public record, the issue is not that authority has disappeared. It is that the coherence through which authority is exercised may become harder to demonstrate when AI-assisted tools mediate evidence, synthesis and case orientation before final judgment is reached.

Within the ILII Public Case Study Examination Methodology, this cumulative pattern is consistent with emerging Governance Drift, while the intact finding on formal human accountability limits the conclusion to governance coherence rather than accountability displacement.

Governance Drift does not describe the presence of artificial intelligence. It does not describe the failure of any single control. It describes a structural condition in which decision influence changes

faster than the institution can demonstrate that authority remains coherently exercised through the decision process.

7. Implications for Ongoing Governance Coherence Examination

The findings of this Public Case Study Examination warrant continued governance coherence examination of the Home Office's asylum decision-making process as the operational role of artificial intelligence evolves.

The question is not whether AI-assisted tools should be treated as inherently illegitimate. Nor is it whether formal human accountability has disappeared. The public record suggests a narrower and more important governance concern: accountability may remain formally human while the decision process beneath it becomes harder to locate, reconstruct and defend.

Continued examination should therefore determine whether the historical governance assumptions reconstructed in this investigation remain demonstrably coherent over time. Particular attention should be given to whether:

- direct engagement with primary evidence continues to form the basis of consequential judgment;
- human synthesis remains a substantive component of decision formation rather than becoming displaced by systematised AI synthesis;
- the exercise of institutional judgment remains meaningfully attributable to identifiable human decision-makers;
- consequential reasoning can be independently reconstructed, explained and defended when challenged;
- operational controls continue to preserve the underlying governance position rather than becoming substitutes for demonstrating legitimacy; and
- formal human accountability remains supported by a decision process that accountable actors can understand, evidence and defend.

These questions extend beyond operational performance. They concern whether institutions can continue to demonstrate that consequential decisions remain grounded in the authority they claim to exercise.

Governance coherence examination is therefore not the conclusion of institutional governance analysis. It is a prerequisite to remediation. Before an institution can determine which legal, operational, technical or procedural controls should be strengthened, it must first identify where the connection between authority, judgment and decision formation has become less demonstrable.

Although these observations arise from the United Kingdom's asylum system, the underlying governance question is unlikely to be unique to this case. As artificial intelligence participates in the interpretation, synthesis and formation of consequential decisions across sectors, institutions may require governance diagnostics capable of examining whether historical assumptions about authority, judgment and accountability continue to hold.

Within the ILII Public Case Study Examination Methodology, persistent difficulty demonstrating this coherence may indicate the emergence of Governance Drift. Where Governance Drift remains unresolved over time, institutions may begin accumulating Governance Debt: the progressive accumulation of structural misalignment between the governance position that legitimises institutional authority and the operational processes through which that authority is exercised.

Governance coherence examination should therefore complement existing legal, technical and operational governance activities. Its purpose is not to replace legal oversight, technical assurance or implementation controls. It provides accountable institutions with an independent means of examining whether consequential decision formation remains connected to the authority that legitimises it.

As artificial intelligence participates more directly in consequential institutional decision-making, the central governance question may no longer be whether AI can be controlled. It may be whether institutions can continue to demonstrate that the authority behind their decisions remains coherently exercised throughout the process by which those decisions are formed.

About the ILII Public Case Study Examination Series

The ILII Public Case Study Examination Series applies the Institutional Legitimacy & Influence Index (ILII) methodology to publicly documented institutional decision environments.

These examinations do not determine legal liability, regulatory compliance, or technical system performance. Rather, they investigate whether publicly observable evidence suggests that the governance position from which institutional authority derives its legitimacy remains coherently reflected throughout consequential decision formation.

Each examination is conducted using publicly available information and is intended to demonstrate the ILII investigative methodology while contributing to the broader study of governance coherence in AI-influenced decision systems.

About the Institutional Legitimacy & Influence Index (ILII)

The Institutional Legitimacy & Influence Index (ILII) is a governance diagnostic methodology developed by Neil Sumaru through the Institutional Governance Research Initiative and published by Levuka Venture Lab.

Rather than evaluating the technical performance of artificial intelligence or the adequacy of operational controls, ILII examines whether consequential decision formation remains coherently aligned with the governance position from which institutional authority derives its legitimacy.

The methodology investigates whether historical governance assumptions continue to hold as artificial intelligence increasingly participates in institutional decision environments. Where this coherence becomes progressively more difficult to independently demonstrate, the methodology identifies the potential emergence of Governance Drift.

ILII is intended to complement, not replace, existing legal, technical, operational and assurance activities by providing an upstream diagnostic examination of governance coherence.

Published by: Levuka Venture Lab
Institutional Governance Research Initiative
© 2026 Neil Sumaru. All rights reserved.

