SAFWAN ALMODHAYAN



OPINION AND ANALYSIS

Editorial date: October 29, 2026

The Name That Stopped a Probate Application After Its Owner Died

> In British probate records, an application may be stopped because the deceased person's name appears differently across documents. The story begins with a name field and reaches a larger question: can an institution serve people whose circumstances do not resemble the case around which its system was built?

On March 26, 2026, the UK Ministry of Justice published a statistical evaluation of the probate service in England and Wales. Among tables on processing speed and causes of delay appeared the short administrative label **Deceased Name**. It was one category used to record why some applications had been stopped, alongside problems with death certificates and document matching.1

Behind the phrase is a reality that cannot readily be reduced to a typing error. The evaluation documented cases in which the deceased person's name did not match across the probate application, will, and death certificate—sometimes because one document used a westernized form, or because a translation did not meet the application's requirements. The report tells no named individual's story and does not establish that every difference should have been accepted automatically. It shows how an apparently simple question on a screen can become a question about identity, evidence, and legal entitlement.1

The deceased person cannot correct the name. The heir does not need a computer to affirm their grief; they need a service capable of recognizing that a spelling difference may be a documentary matter that can be resolved, while another difference may reveal a real dispute requiring investigation. In either case, "the data do not match" is not enough.

The facts may be entirely true and still not fit the field.

This is where a test often absent from celebrations of digital transformation begins. Once a service succeeds at handling thousands of similar applications, unusual cases can resemble guests who arrived through a door missing from the building plan. Not every guest delayed at the entrance is wrong.

What Actually Stops When an Application Is Stopped?

Begin with the word “stop.” In British probate terminology, a **stop** means that a staff member encountered a matter requiring clarification, a document, or another action before processing could continue. An applicant may be asked to prove eligibility or supply an original document. A stop is **not the same as a final refusal**, nor is it a ruling that the applicant has no right.1

That distinction is not merely linguistic. Within its legal bounds, a refusal answers an application. A stop suspends the answer. It can provide necessary protection where a will is disputed, or be an extra loop created by an interface that failed to ask about alternative forms of a name at the outset.

Government guidance updated in July 2026 shows another side of the issue: the online **MyHMCTS** system does not accept every probate application. Some cases must be submitted on paper, including foreign wills, second grants for the same estate, certain applications to rectify a will, and cases requiring discretionary judicial powers. The official text says: “MyHMCTS cannot accept all types of probate applications.”2

Two entirely different kinds of case therefore sit beneath the label “non-standard.”

The first is a **genuine legal exception**: a file merits special examination because the law, the nature of the document, or a possible conflict of rights requires it. Erasing that exception in the name of a smoother user journey may remove the very protection for which the procedure exists.

The second is a **service-created exception**: a fact can readily be proved by a suitable means, but the application stalls because the form has no place for a previous name, instructions appear after rather than before submission, or the system asks a user to describe a real document using an option that does not apply.

The distinction determines where spending belongs. An institution should neither buy more automation to swallow legal judgment nor employ a judicial expert to resolve every difference that one extra field could explain.

Two Weeks for One, Fourteen for the Other

Official statistics published on September 24, 2026, show that probate applications resulting in grants during the second quarter of that year took an average of about **two weeks** if they had not been stopped and about **14 weeks** if they had. During the same period, **92 percent** of probate applications were submitted digitally and **94 percent** of grants were issued through the digital route.3

The twelve-week difference is not an estimate of harm caused by technology alone. Stopped cases already differ in complexity and documentary completeness, and some involve legitimate disputes. Comparing two different groups does not allow us to say that the stop itself produced the whole delay. It does reveal that the experience of the service does not have one duration.

The March 2026 evaluation, based on applications received between July 2019 and June 2025, found that **34 percent** were stopped at least once: **32 percent** of digital applications and **38 percent** of paper applications; **46 percent** for applicants without legal representation and **26 percent** for those using legal representatives.1

This is not the proportion of “rare cases” in society. Some stops resulted from a missing document in an ordinary file, and applicant populations differ across channels. Nor do the figures prove that paper causes problems or that using a lawyer causally reduces them. They do prevent us from treating stops as a negligible margin that need not appear in the design of the service.

The result looks strange beside an attractive presentation: an institution can draw an almost straight digital path, then discover in its operating report that roughly one-third of applications temporarily left it. The screen is not necessarily lying; it may simply have been photographed from a convenient angle.

Why Does a Different Name Become More Work?

It is not enough to place the word “bias” above the gap and consider it explained. The British evaluation also studied differences in the experiences of applicants who voluntarily answered questions about personal characteristics.

Among respondents to the ethnicity question, **56 percent** of applications from people in ethnic-minority groups were stopped at least once, compared with **44 percent** of applications from people identifying as white. For applications made in the first quarter of 2025, half of those in the white category were completed within **12 days**, compared with **25 days** for half of those in the ethnic-minority category, using the report’s method of measurement.1

The comparison has important limits. The personal-characteristics questions were offered to applicants without legal representatives and responses were voluntary; about **40 percent** of eligible applicants responded. Of those disclosing ethnicity, **96 percent** identified in the white category and **4 percent** in ethnic-minority categories. Individual circumstances were not all measured, and this sample is not a guaranteed portrait of all heirs. Types of wills, translation, legal representation, and estate complexity affect outcomes. The differences therefore describe **an association requiring explanation**; alone, they do not prove that the system treated someone differently because of ethnicity.1

What matters is that the agency itself looked for mechanisms closer to the ground. Its evaluation pointed to differences involving the status of wills, deceased names across documents, translations, and assets abroad. In December 2024 it said it had changed questions describing a will's condition from January 2023 and was redesigning how it collected a deceased person's name and its variants. It described early signs of improvement cautiously and said more completed files were needed for greater confidence.4

The scientific value lies not in attaching a grand name to the problem, but in discovering which part can be repaired by a better question, which requires new evidence, and which should remain under specialist review.

No One Lives in the Average

Consider a hypothetical service that completes ninety applications in two days and takes sixty days for ten others. Its average completion time is **7.8 days**. If the ninety are accelerated to one day while the ten remain unchanged, the average falls to **6.9 days**. The chart improves; nothing changes for the ten people still waiting.

This is an entirely hypothetical calculation, not a description of British data. Its purpose is to expose a weakness in how quality is seen: an improved average can coexist with an unchanged burden for people with difficult cases. The investment may help most people without answering the other ethical and administrative question: what did we do for those whose position did not improve?

In risk-management research, expected loss can be represented by combining each case's probability with the scale of the resulting harm. If case type is g, its probability p_g, and its average social cost c_g, then:

$$\mathbb{E}[L] = \sum_{g} p_g \, c_g$$

Here, c_g is not limited to the cost of an employee hour. Depending on the service and what can safely be measured, it might include an inability to deal with estate assets, delayed financial benefits, time lost resubmitting evidence, or avoidable legal-representation costs. This does not

give every effect a certain price. It explains why throughput alone cannot speak for everyone.

Longer durations, such as the 90th or 95th percentile of resolution time, can also be examined, provided the denominator and still-open cases are made clear. A tail measure exposes delay but does not determine whether the delay was necessary to protect a legal right. The institution must know **why the person waited**, not merely count the days.

That matters especially for cases not yet resolved. Measuring completed applications alone can brighten the picture because the longest-running cases remain outside the “completed” table. The probate report itself noted limits in recent-quarter duration data when too small a proportion of cases had been completed.1

When Should a Machine Say, “I Don’t Know”?

Machine learning offers a more modest answer than the promise to automate everything: **selective prediction**. A system need not answer every case. It can predict when conditions it trusts are met and abstain from judgment when the likelihood of error rises. In 2019, Yonatan Geifman and Ran El-Yaniv introduced **SelectiveNet**, which learns prediction and rejection together rather than adding only a confidence threshold after training.5

In technical language, the share of cases the system answers is its **coverage**, and its error is examined within the covered cases. It is relatively easy to reduce error if the system answers only obvious files. But a public service cannot leave everyone else in a void. “When should the model abstain?” must be paired with “what service begins the moment it does?”

In 2020, Hussein Mozannar and David Sontag developed the idea of **learning to defer**: balancing a model’s decision against referral to a human who may be better at a particular class of cases.6 The practical lesson is not that humans are infallible. Sometimes employee and model both err; an expert may excel at reading an unusual will while the system is more accurate at detecting an empty standard field.

A useful scientific objection followed. In 2022, Rajeev Verma and Eric Nalisnick showed that one earlier learning-to-defer formulation does not guarantee calibration of the probability that an expert is correct in the way a designer might expect, and they proposed a more disciplined probabilistic alternative.7 Writing “refer to a human when uncertain” in a product document is not enough. Estimates of model confidence and reviewer expertise must be tested against independent data and new cases.

The British report contains no evidence that the probate system used these AI models. **What is documented is a digital service and human procedures; the design described here is an inferential proposal for future automation.** That boundary matters. Without it, “AI” becomes a generic name for every field that frustrates its user.

Three Answers Better Than “Accept” and “Reject”

Suppose, as a **hypothetical example**, that an organization is designing a digital service for an application containing two documents with two forms of a name. It may not assume the names belong to one person; the difference could be serious. Nor may it assume two people merely because the strings differ.

A better service has three clear moves.

First, automatically verify the familiar case. The system checks document and field completeness, and conditions testable without legal judgment. A case satisfying the rules is processed, with a record of what the system actually verified.

Second, request specific clarification when information is missing. Instead of “the file is invalid,” the system identifies the document in which the name differs, permits an alternative name and evidence connecting it, and states the next review date. Privacy may prevent display of some matching details; the service can then explain what it is permitted to explain and provide a secure route for questions.

Third, refer decisions requiring human judgment or special authority to the right specialist. The employee is not asked to click “approve” on a recommendation they do not understand. They receive the evidence, amendment history, and unresolved questions. The applicant receives as much explanation as the law permits, with an opportunity to supply more information or challenge the decision under the applicable system.

These are not three new buttons. They are three different responsibilities: **cleaning data, completing evidence, and exercising judgment**. Mixing them makes a spelling error look like a legal dispute, or a legal dispute look like a text-formatting problem.

Fairness Has a Procedural Form

The legal psychologist Tom R. Tyler studied **procedural justice** in research published in 1988, based on interviews with **652 people** who had recent experiences with police or courts. People’s assessments of fairness were related not only to the outcome, but also to their opportunity to present their view, the quality of the decision, the availability of error correction, and other elements of the procedure.[8]

It would be unsound to project an old American sample onto every Saudi or British user today. The research does offer a testable institutional hypothesis: explaining why an application stopped and allowing an intelligible response may affect perceived fairness and trust even when a person cannot be given the decision they want.

Administrative oversight bodies give this idea practical language. A March 2025 guide on automated decisions from the Australian **Commonwealth Ombudsman** emphasizes reviewability, explanation, and accountability in proportion to a decision's impact.9 Guidance from the **NSW Ombudsman**, within its local legal context, explains that statutory discretion cannot be reduced to automated adoption that removes genuine consideration of a case, and that fair procedure requires a real opportunity for the affected person to be heard.10

These documents do not mean every country has the same legal rule, or that every probate application requires a hearing. They give the designer a question more important than "did the application reach the final stage?": **could the applicant submit evidence capable of changing the judgment?**

A digital form that respects this right may be more efficient than one repeating an obscure error message. A good explanation reduces correspondence that adds no information. A response that returns a citizen to the beginning is like an employee asking them to bring the original copy of the question.

The Strong Case for Digital Services

After all this, it would be easy to conclude that sensitive transactions should return to offices filled with paper. That would be a weak inference, and it would ignore the real beneficiaries of digital services.

Recent British statistics show that the overwhelming majority of probate applications are submitted electronically. In the second quarter of 2026, grants in unstopped cases took about two weeks on average—valuable speed for someone resolving an estate with clear documents.3 Government evaluations also documented changes to questions about wills and how name variants were recorded as part of addressing delays. The solution may itself be **digitalization that better understands variation**, not a retreat from it.4

A competing explanation for the delay differences must be taken seriously: more complex files take longer because they are more complex, whatever interface receives them. A foreign document or real dispute demands scrutiny that cannot be removed to satisfy a performance chart. Some applicants use legal advisers; others encounter the procedure for the first time. The descriptive comparison does not isolate every factor's effect.

That argument does not erase the service's responsibility. It changes the success test from "make every case equally fast" to "make every extra step explicable, necessary, and consistent with available review capacity."

There is another constraint: human specialists are themselves scarce. If every name difference is referred to the most senior expert, work will accumulate and automated abstention will become human waiting. Referral is safe only when designers know which expertise is required, how priority works, each team's caseload, and what support staff can complete before a file reaches the specialist.

In October 2026, the Institute of Chartered Accountants in England and Wales reported, in a summary of probate user-group meetings with HMCTS, that some missing-will and foreign-domicile files require specialist guidance, and that certain applications need both operational and registrar consideration.11 This is a professional account of meeting statements, not an independent experiment establishing the cause of delay. It shows that expertise does not disappear when the entrance moves online.

The Exception With No Name in the Database

There is an irony more recent than much of the theoretical discussion. On **September 9, 2026**, the UK Ministry of Justice answered a parliamentary question about specialist files. It said HMCTS does not centrally classify probate applications as "specialist" and "non-specialist," and does not hold comprehensive data permitting a precise figure for the time particular files spend under specialist review, including some cases involving missing wills or affidavits.12

The answer does not mean all operating records are absent or that the agency cannot process these cases. The ministry notes training, increased staffing, and procedural improvements during 2026. But it exposes a limit to knowledge: total waiting time may be known while it cannot be separated by the kind of expertise a file required.

Before the problem is algorithmic, it is definitional. How can we know whether using AI to triage exceptions succeeded if we cannot connect the type of complexity to review time and outcome? How can the effect of changing the name interface be tested if case categories change from one version to another?

The answer is not to collect everything about people. It is a purpose-limited record preserving **the type of stop reason; whether it is a remediable omission or a matter requiring judgment; dates of clarification, response, and resolution; the channel used; and the review outcome**. Privacy must be protected, access controlled, and sensitive categories kept from becoming unnecessary labels circulating among operating teams.

Imagine an institution saying, "We measure every click," then admitting it cannot tell you how long an application takes to reach the person with the right authority. It has mastered the science of movement before defining the destination.

A Proposal: Exception Containment Capability

The analytical standard proposed here can be called **exception containment capability**. It is **a term proposed here for administrative analysis, not an accepted academic theory or an existing government standard**.

It means that a service can recognize an unfamiliar case, identify why it left the ordinary process, bring in the right person or evidence, and produce an explained and correctable resolution within an intelligible time, without losing the file between channels.

The proposal can be tested without reducing a right to one number. First create documented categories of unusual cases, then observe four things together:

1. **Routing integrity:** how many cases reached the right kind of expertise on the first referral? An independent sample of files should be reviewed.
2. **Resolution time:** what are the median and 90th percentile from discovery of the difficulty to resolution, correctly accounting for files still open?
3. **Explanation quality:** did applicants understand what was missing, could they submit it through a channel available to them, and did they retain a right of challenge where the law provides one?
4. **Outcome correctness:** how many decisions changed after a document was completed or an independent review, and what type of error caused the change?

If resolution time falls while justified challenges remain high, the wrong path may have been accelerated. If stops fall because staff ceased recording them, visibility has been lost. The “exception” classification should therefore be tested on a sample of ordinary applications too, and paper-referred cases must remain in the denominator so that figures do not improve merely by removing the hardest cases from the screen.

Falsifiable prediction: if collection of alternative names and documentary guidance are redesigned while legal evidence requirements remain unchanged, avoidable name-matching stops should fall and their resolution should become faster among comparable cases, without a rise in erroneous decisions discovered later. Comparable service units or periods can be examined, adjusting for case mix and changes over time. If outcomes do not improve, the “field” hypothesis may be weaker than explanations based on missing evidence or case complexity.

The March 2026 figures alone cannot directly test this prediction; they are descriptive data combining operating changes and multiple years. Acknowledging that deficiency is part of the idea’s value.

Who Benefits When a File Deserves a Different Route?

In Saudi Arabia and the Gulf, the **idea can be applied as a proposal** to services dealing with names written in multiple forms, documents issued in different languages, or rights crossing jurisdictions, subject in each case to applicable local rules. No existing government program need be invented to see a research possibility.

At a financial institution, an exception might arise when verifying an heir's relationship to a deceased customer's account or reviewing a foreign document. Text matching is insufficient, and a customer-service employee cannot infer authority over an account from similar names. The value of exception design is to carry evidence to an authority entitled to examine it, while telling the family precisely what must be completed.

At an insurer, the occurrence of an incident may be clear while the beneficiary's status or supporting document is unfamiliar. An automated decision can check whether some data are complete and flag inconsistencies; the legal or contractual decision may remain a specialist's responsibility. The beneficiary wants a known time and an intelligible path. The institution wants to prevent improper payment without turning every atypical application into a suspicion.

Before investing in a new automated service, a board or regulator can test three questions: which cases is the system authorized to decide? How can an applicant add evidence capable of changing the outcome? What is the actual capacity of the team receiving files the system declines to decide?

These questions accord with the UK government service standard requiring that everyone be able to use a service, including support for people limited by digital skills or internet access, and with guidance updated in September 2026 calling for AI risks to be assessed against user need and service protection.[13][14] They are **service-design references**, not statements of Saudi legal obligations.

The Same Name, After the Field Changes Meaning

Return to "Deceased Name" in the March 2026 table. It may look like an internal code useful for counting applications. But it reveals a deeper problem: one person may leave valid documents whose letters do not match, or conflicting documents that demand serious judgment. The service fails when it treats both as one error.

Good digital transformation does not promise that exceptions will disappear; some must remain because justice requires deliberation. It makes a more exact and harder-to-prove promise: an unusual case will not become an additional time penalty merely because nobody imagined it while drawing the form.

The test is not the number of files that crossed a screen without stopping. It is whether a stopped file reaches the right place, whether its owner knows what is happening and why, and whether there is a route to correct the record when the record is wrong.

In the end, the field should be the smallest thing in the service—not the smallest thing in the life of the person with the right.

References

1. **UK Ministry of Justice / HM Courts & Tribunals Service.** *Probate Reform Evaluation Data Summary*, published **March 26, 2026**, an 82-page original government report analyzing July 2019–June 2025 data. Primary source for stop rates, duration differences, and stop-reason labels. **Limit:** descriptive study; personal-characteristics data were voluntary and mainly available for unrepresented applicants, and do not establish causality.
2. **HM Courts & Tribunals Service.** *Apply for probate on paper as a practitioner*, first published October 29, 2020, updated **July 13, 2026**. Official operating guidance listing applications excluded from MyHMCTS. **Limit:** applies to England and Wales; the cases stated are specific examples from the text.
3. **UK Ministry of Justice.** *Family Court Statistics Quarterly: April to June 2026*, published **September 24, 2026**, accredited official statistics; probate section. **Supports:** two-week/14-week figures among applications resulting in grants, and digital-channel rates. **Limit:** descriptive averages for different groups; administrative-database figures may be revised.
4. **HM Courts & Tribunals Service.** *Assessing Access to Justice in HMCTS Services – December 2024*, **December 2024**, original government evaluation; "Probate" section. **Supports:** redesign of will-condition questions from January 2023, work on name variants, and caution about the changes' effectiveness.
5. **Yonatan Geifman and Ran El-Yaniv.** *SelectiveNet: A Deep Neural Network with an Integrated Reject Option*, **ICML 2019**, peer-reviewed conference paper. **Supports:** integrating prediction and rejection in one training objective; does not establish success in a judicial field setting.
6. **Hussein Mozannar and David Sontag.** *Consistent Estimators for Learning to Defer to an Expert*, **ICML 2020**, peer-reviewed original research. **Supports:** framing the choice between a model answer and expert referral; its theoretical and empirical analysis concerns the paper's tasks, not the probate service.
7. **Rajeev Verma and Eric Nalisnick.** *Calibrated Learning to Defer with One-vs-All Classifiers*, **ICML 2022**, peer-reviewed research critiquing calibration in an earlier formulation. **Supports:** testing the probability of expert correctness and confidence calibration; does not establish that every deferral model fails.
8. **Tom R. Tyler.** *What Is Procedural Justice? Criteria Used by Citizens to Assess the Fairness of Legal Procedures*, *Law & Society Review*, **1988**, empirical research with **652** participants in a US context. **Supports:** the value of voice, error correction, and neutrality in perceived procedural fairness; generalization across countries and services is limited.
9. **Commonwealth Ombudsman, Australia.** *Automated Decision-Making – Better Practice Guide*, **March 2025**, a 61-page official oversight guide. **Supports:** designing review, explanation, and accountability in automated decisions; guidance within its Australian scope.

10. **NSW Ombudsman, Australia.** *Administrative law and automated decision-making*, official administrative-law guidance, accessed **October 2026**; publication date not established here. **Supports:** limits on automating discretion and the need for genuine consideration and a hearing in the New South Wales context.
11. **Institute of Chartered Accountants in England and Wales.** *HMCTS update for probate firms: October 2026*, **October 22, 2026**, professional coverage of user-group meetings with HMCTS. **Supports:** examples of complex files and registrar consideration; a **secondary account** of the agency's presentation, not an independent test.
12. **UK Ministry of Justice, parliamentary answer.** *Written Question 26374*, answered **September 9, 2026**, official primary record. **Supports:** absence of a comprehensive central specialist/non-specialist classification or complete duration data for some files; does not imply absence of every working record.
13. **UK Government Digital Service.** *Service Standard, Point 5: Make sure everyone can use the service*, official government design guidance, reviewed **October 2026**. **Supports:** accessibility, non-exclusion, and support for people without digital access or skills.
14. **UK Government Digital Service.** *How to apply the Service Standard*, first published **May 8, 2019**, updated **September 7, 2026**. **Supports:** understanding AI risks when it is used and ensuring it serves a user need; British guidance, not a Saudi rule.