The Eligibility Edge: Why Borderline Grant Cases Fail

A documentary review and constructed evidence model for mandatory requirements unresolved facts and review decisions

Authored by Chennakeshav Adya

Independent Researcher

November 2026

Abstract

Background. Grant screening can leave different questions unresolved about the applicant, project or governing terms. Published processes separate mandatory requirements from assessment and funding decisions; some borderline scope cases proceed to assessment.

Objective. This paper develops a source-linked evidence record for deciding which requirements are supported, failed or unresolved, and which additional review work is justified.

Approach. A dated review of selected European, UK, US and Gulf official materials is combined with two constructed exercises: exhaustive comparison of 729 vectors with six three-state positions and an evidence-resolution model using twelve authored cases, six checks and two fixed review orders. Work units and case outcomes are assumed; no applicants are observed.

Findings. A compensating score at five-sixths passes six vectors containing a failed mandatory predicate. Ten of twelve constructed cases have a certificate within the selected check set, including three initial certificates. Entity-first reaches ten certificates across the twelve cases with a per-case budget of ten work units; cheap-first reaches ten with a per-case budget of fifteen. Two cases remain unresolved. These are logical and constructed-case results, not applicant success rates.

Implications. A reviewable decision should preserve the exact applicant, rule version, facts, decisive evidence and unresolved action. Programme interpretation, competitive selection and actual funding require their own evidence. A prospective applicant-cohort protocol is proposed for empirical evaluation.

JEL Classification: O31, O38, D83, C61, G32

Keywords: grant eligibility, innovation finance, SME definition, ownership and control, evidence provenance, application screening, SBIR, EIC Accelerator, unresolved evidence, funding decisions

1. Classify the decision before explaining a failure

A grant application needs an identifiable basis for each mandatory requirement before its funding prospects can be discussed responsibly. We propose an evidence record that connects the applicant, the selected call, the applicable terms and the facts used in a screening decision. The central question is which uncertainty remains and what evidence would resolve it. This paper examines that question through a dated documentary review, a formal representation of mandatory requirements and explicitly constructed evidence cases. The recommendations are

designed for founders, finance leaders and advisers deciding whether to prepare, clarify, revise or discontinue a particular application.

The title supplies no finding that every borderline application fails. Innovate UK's published process allows on-edge scope cases into assessment when judged within scope. [1] A documented failure of a selected requirement, an unresolved evidence gap and an unsuccessful funding decision retain different meanings. Actual feedback is needed before attributing an outcome to eligibility.

Consider a hypothetical company whose project fits a programme's technology theme and whose ownership documents have not been reviewed. Its theme fit provides evidence about the project. The missing ownership review leaves a separate question open. A positive commercial assessment cannot settle that question. Equally, the absence of a reviewed ownership record provides no affirmative evidence of a prohibited ownership structure. A finance manager needs both distinctions when authorising additional preparation work. The appropriate next action may be a document request, a programme clarification or a reasoned decision that the selected route should no longer be pursued.

Our contribution is a bounded decision design. It specifies how to record a rule, attach evidence to a proposition, retain an unresolved state and produce a reviewable explanation. The numerical exercises examine the consequences of that design within fully disclosed artificial cases. We make no priority claim that these logical or evidential concepts are new. The present review also contains no representative applicant population, rejected-application dataset or consented company cases. It cannot establish how common eligibility failures are, predict a company's award probability or measure the financial return from using the proposed process.

The practical output is a decision certificate for a defined set of requirements at a stated time. A certificate identifies the decisive facts, the rules consulted, any outstanding issue and the person authorised to approve the next action. Its scope is deliberately limited to the questions actually checked. Eligibility, competitive merit, contracting and payment each require their own evidence. Management should carry those separate states into the funding plan so that an application under review is described consistently in the board pack, financing model and communications with potential investors.

2. Define the documentary study and its limits

We selected official materials that expose different sources of boundary uncertainty. The European materials cover applicant routes, enterprise size, group relationships and access to a later application stage. The United Kingdom materials distinguish application requirements, scope judgement and competitive assessment. The United States materials describe selected small-business eligibility conditions and a recent change in submission limits. Hub71 provides a Gulf example in which permission to apply, assessed selection and later participation obligations must be recorded separately. These sources form a purposive comparison of decision structures. Their inclusion does not establish that the selected programmes represent all innovation funding in any jurisdiction.

The documentary observation date is 10 November 2026. We preserve a document's own version, release date, applicability date and the precision of its stated timetable. These fields serve different purposes. An access date establishes when the research team consulted a page. A release

date records the source's publication statement. An applicability date identifies the events to which a rule says it applies. A month-only cohort deadline remains month-only until a more precise authoritative statement is available. The study does not reconstruct every historical page version or assume that the first observed wording is the first wording ever published.

For each claim, the research record contains the official URL, document role, relevant clause or section and a local statement of the claim's limits. The evidence may consist of an actual programme document, an agency restatement or a provider's explanatory page. Those roles remain visible when sources disagree. A FAQ can help identify an operational question while leaving an inconsistency with a work programme unresolved. The applicant's legal advisers or the responsible programme authority must address any interpretation material to a real case. This paper provides selected checks and a way to organise that work.

Direct access to the current US electronic regulations was unavailable during this review. We therefore identify the accessible NIH policy statement and the posted SBA directive by their actual versions and reading scope. The March 2026 NIH statement predates the April 2026 reauthorisation Act. The posted directive contains legacy expiry wording. Neither source is presented as a complete audit of all legislation, later regulations or current agency solicitations. These limitations matter particularly when a team is preparing to certify eligibility, because a selected summary cannot replace the applicable solicitation and its amendments. [8, 9, 11]

The analytical work is divided into documentary findings, proposed operating practices and constructed calculations. Programme facts carry source references. Proposed record structures and decision rights are identified as recommendations. Constructed cases state their assumed facts and units before any calculated result is discussed. No invented applicant is presented as a client, award recipient or rejected company. This division allows readers to reuse the method while checking the actual rules and evidence relevant to their own financing decision. It also makes clear which claims could be tested later using a properly assembled applicant dataset.

Table 1. Programme comparison and the boundary of this review

Source set	Selected decision issue	Further evidence for a real case
EIC Accelerator	Entity route and requested component	Full current programme and relevant country/call conditions
Commission SME definition	Measurement and aggregation boundary	Accounts, work units and ownership/control records
Innovate UK	Submission, scope and selection	Exact competition brief and actual feedback
NIH and SBA	Owner, affiliate and phase conditions	Current solicitation, amendments and applicable rules
Hub71	Application versus participation	Current cohort instructions and actual offer terms

Selected documentary comparison [1-14]; no programme-wide or applicant-specific eligibility certificate.

3. Separate the stages of the funding decision

We use admissibility to describe the required form and submission conditions, eligibility to describe the selected applicant or project requirements, and scope to describe the project's fit with the call. Merit concerns the assessment of the proposal against published award criteria. Selection records the funding body's decision after the relevant assessment and allocation process. Contracting concerns the conditions for executing the funding agreement. Payment records the actual release of funds. These are analytical labels for organising the review; a programme may

use different terminology or combine several checks in one stage. The underlying source wording must remain available.

The analyst should identify the exact question answered by a document before using it elsewhere. An incorporation certificate concerns the existence and registered details of an entity. A payroll return records employment information under its own reporting convention. A technical report supports specified claims about a technology and test environment. A portal receipt records a submission event. Each document has an author, period and purpose. Its usefulness for a grant requirement depends on the match between those characteristics and the requirement being assessed. A document's presence in the data room alone provides an incomplete basis for that match.

An important practical distinction is the difference between a rule and an evaluative judgement. A published numerical limit can be expressed as a predicate once its measurement basis and exceptions are understood. A panel's judgement about commercial potential needs a different record, including the criterion and the evidence considered. The analyst should preserve the judgement as a judgement. Assigning it an invented percentage or treating it as a deterministic legal test would add information that the underlying record does not supply. The proposed method permits structured reasons without manufacturing certainty about institutional decisions.

A hypothetical applicant can be legally established in a permitted jurisdiction, provide a complete application and still receive an unfavourable merit assessment. Another applicant can receive a strong technical assessment while remaining subject to later agreement conditions. These scenarios demonstrate why one overall traffic-light field is insufficient for the financing team. A compact stage record can state the requirement, current evidence and unresolved action for each decision. It should preserve the programme's actual status terms alongside the team's analytical description so that a reader can reconcile the internal assessment with official correspondence.

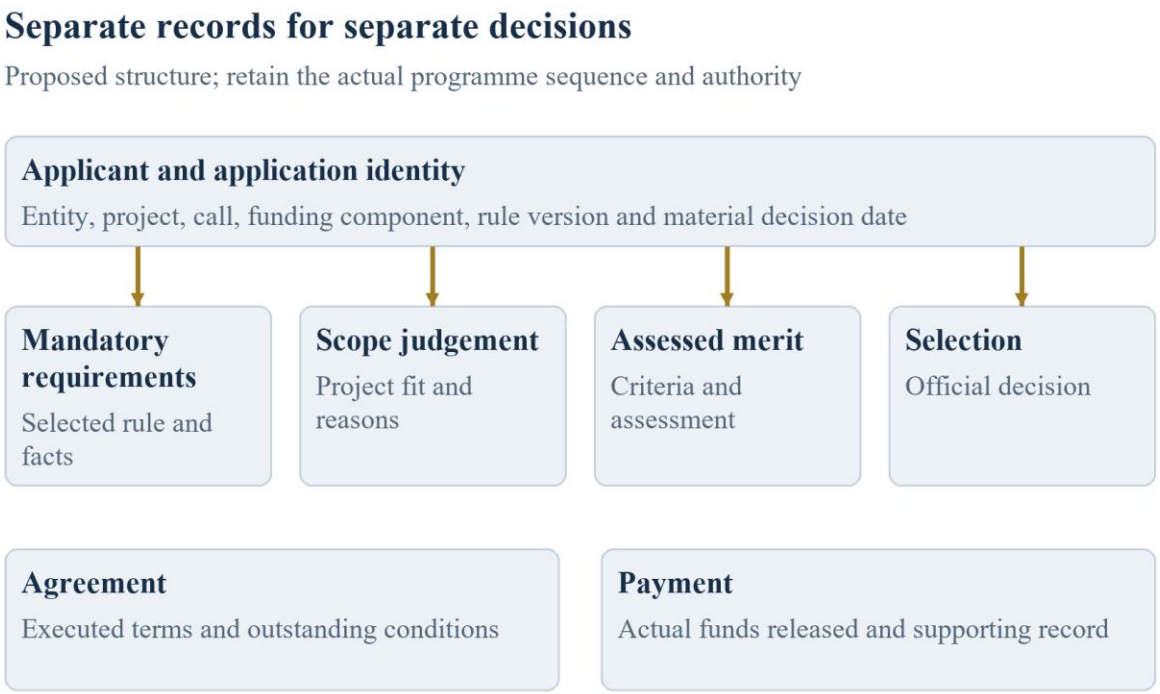
For management purposes, we recommend assigning authority at the stage where a decision is made. The company can approve its own preparation budget and verify its records. Its advisers can review identified rules within their professional remit. The funding body retains the decisions allocated to it under the programme. A board's willingness to proceed with an application does not change the funder's authority or establish an award. The funding plan should therefore identify which decisions remain external and which commitments the company can support using resources already available to it.

Table 2. Distinct decisions in an application record

Stage	Question to record	Evidence or authority
Admissibility	Are the stated submission conditions met?	Current brief and actual submitted package
Selected eligibility	Are the checked mandatory requirements supported?	Applicable clauses and applicant evidence
Scope	Does the project fit the call?	Submitted explanation and responsible decision
Merit	How is the proposal assessed?	Published criteria and actual assessment
Selection	Has support been offered?	Official decision and its conditions
Agreement	What has been executed?	Signed governing documents
Payment	What funds have been released?	Actual payment evidence

Proposed analytical record; preserve each programme's actual terminology and sequence.

Figure 1. Separate stage records for the financing decision



Authored decision architecture. An earlier stage does not establish the outcome of a later stage.

4. Bind each requirement to an applicant and a date

A reusable eligibility record needs a stable identity for the proposed applicant. We recommend recording its legal name, registration identifier, jurisdiction and relationship to any trading name used in the pitch. The selected call identifier, funding component and proposed project should be attached to that identity. This makes the entity and project covered by the review explicit. Where incorporation or relocation is planned, the record should distinguish the entity that currently exists from the entity expected to enter an agreement later.

Each requirement also needs a material decision date. Depending on the governing text, the relevant fact may be assessed at application, full proposal, award, agreement signature or another checkpoint. The analyst should transcribe that timing condition with its clause. A current fact can support a present assessment while leaving a future condition open. For a proposed corporate transaction, the team should retain both the current ownership structure and the approved transaction assumptions. The later assessment should be updated using the structure that actually becomes effective, with evidence of its effective date.

The record should separate the rule's applicability from the age of the evidence. An annual financial statement can be the relevant document under a specified accounting-period rule even when a newer management report exists. A recent payroll extract can be insufficient for a measure defined over a longer period. The date field should therefore identify what the document measures and why that period is appropriate. A simple recent-document filter would fail to capture that distinction. The reviewer needs to check the measurement convention against the selected rule before accepting the evidence.

We propose retaining a clause identifier, source URL, document version, local observation date, fact period and review owner for every material proposition. A source capture may be identified

by a file hash where the actual bytes have been retained. A hash supports later comparison with that captured file. It supplies no independent proof that the content is authoritative, complete or correctly interpreted. For a web observation retained as extracted text, the record should state that capture form. It should avoid implying possession of an original PDF or complete historical page when only a selected passage was read.

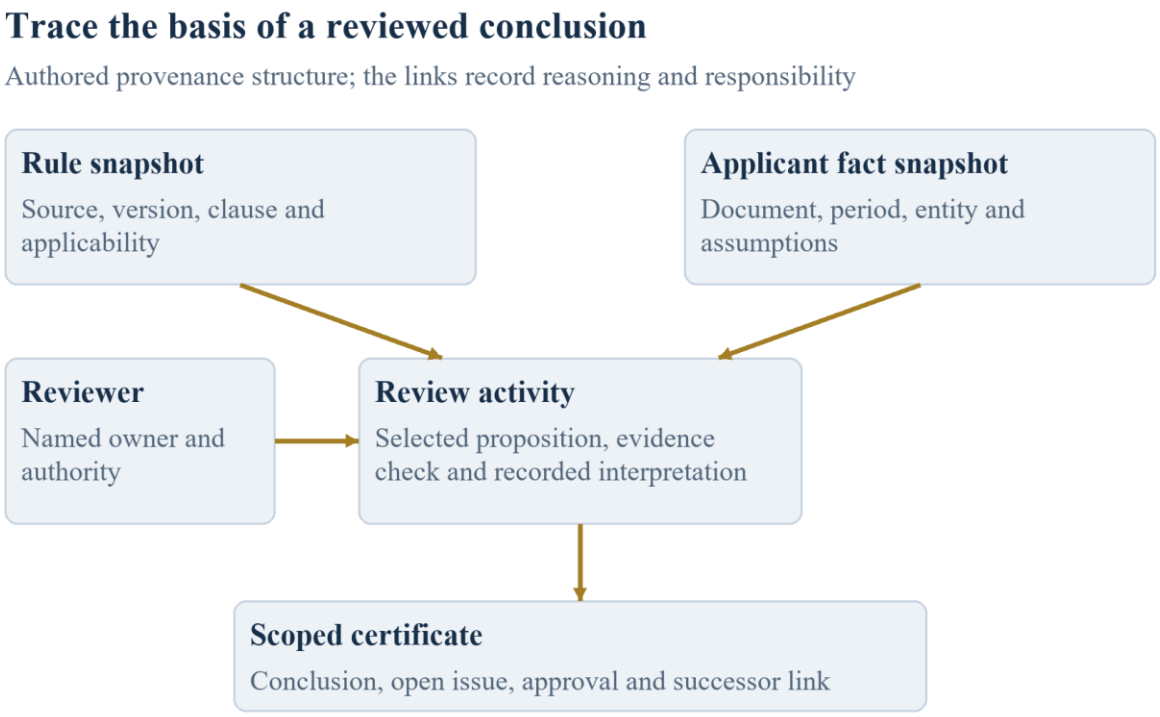
The practical acceptance question is whether another qualified reviewer can reproduce the reasoning from the recorded source and facts. If that reviewer cannot identify which entity was assessed, when a requirement applied or what a number measured, the certificate needs further work. That work can be precisely assigned: obtain the group chart, identify the applicable annual accounts, confirm the prior submission or seek a programme response. Naming the missing evidence allows management to authorise a bounded task with an explicit scope.

Table 3. Eligibility evidence matrix

Analytical category	Candidate evidence	Question remaining for the reviewer
Entity	Registry record and applicant identifier	Does this entity meet the selected route at the relevant checkpoint?
Control	Shares, voting and contractual rights	Which relationships are material under this rule?
Size	Accounts, staff basis and group schedule	Are perimeter, period and aggregation correct?
Activity	Defined project and technical evidence	Which specified requirement does the evidence support?
Cost	Budget and eligible-category reconciliation	What is included and on which stated basis?
Submission	Required materials and portal receipt	Was this exact package submitted under the applicable conditions?

Proposed review matrix; document names alone establish no conclusion. Multiple predicates may be needed in each category.

Figure 2. Connect a rule to the evidence and reviewed conclusion



Proposed record structure informed by PROV concepts [16]; no claim of programme authority or proof of document accuracy.

5. Read the European applicant route by component

The EIC Accelerator Work Programme 2027 describes established and prospective enterprises, with component-specific conditions. The route for an applicant from a non-associated third country requires effective establishment in a Member State or Associated Country by full-proposal submission. The small-mid-cap route has component-specific conditions. Annex 2 and any relevant call conditions remain part of the check. These provisions do not support a universal statement that every applicant must complete every incorporation step before any initial application. [4]

For a hypothetical cross-border applicant, we recommend a route-specific evidence schedule. It would identify who submits the initial proposal, which legal entity is intended to receive support and the document required to substantiate establishment at the relevant checkpoint. The schedule should also identify any country or funding-component exception requiring further review. The team's willingness to incorporate does not itself establish effective establishment. Its evidence plan needs an identifiable deliverable and the relevant decision date. If the available documentation cannot support the selected route, the appropriate assessment remains unresolved pending clarification or changed facts.

The same work programme describes completed TRL 5 validation and separately presents the full-stage award criteria. Table 8 states an overall threshold of 13 out of 15, with individual criterion thresholds. Its quality-seal procedure also contemplates proposals that meet the relevant assessment criteria while remaining unfunded within the available budget. [4] These are different decision records. An analyst should retain the technical evidence, the assessment result and the

funding outcome separately. An enterprise-size certificate cannot answer the technical question, and an assessment score cannot establish that a funding agreement has been executed.

A practical project file should identify the exact technological claim being supported. In an authored example, a battery-monitoring startup might have test results for a laboratory prototype while proposing commercial deployment in a harsher operating environment. The reviewer would document the tested configuration, the environment and the performance claim before assessing what the report supports. This example makes no judgement about an actual EIC application or the appropriate readiness level of a real technology. It illustrates why a technical label should be traceable to evidence and why the funding body's assessment must remain distinct from the applicant's description.

The component record also matters to financing discussions. Management should identify whether it is considering a grant component, investment support or a permitted combined route and retain the relevant conditions for each. The commercial model should describe the proposed instrument accurately. It should avoid carrying a headline programme amount into available cash before the supporting decisions and agreements exist. For a real transaction, advisers should verify the current programme documents and executed terms. This paper's selected provisions are a basis for organising that review, with no applicant-specific entitlement or legal opinion supplied.

6. Preserve an official source conflict without guessing

The current EIC materials provide a concrete documentary boundary. The Work Programme 2027 full-proposal text names successful short applications under the 2026 or 2027 programmes. Footnote 129 contains a conditional route for 2024 and 2025 GO outcomes where a full proposal has already been submitted and submission limits remain available. The live FAQ's first answer names 2025, 2026 and 2027. Another FAQ answer refers to 2026 for the described 2024 case. Those statements were observed together and are retained as an inconsistency. This review does not select a favourable interpretation for an affected applicant. [4, 5]

A team encountering that conflict should first identify whether it is material to its actual route. The necessary case record would include the short-proposal result, the programme year, any prior full submission, the legal entity concerned and the relevant submission history. A question framed only as whether an old GO is valid omits facts needed to locate the stated exception. We recommend submitting a concise clarification request that reproduces those facts and cites both documents. Any response should be retained with its author, date, scope and connection to the identified application.

Until the issue is resolved, the evidence status should state the open interpretation. The company can separately decide which reversible preparation tasks to continue. Management might authorise work that remains useful for another funding route, while holding a commitment that depends specifically on access to the disputed stage. These are proposed decision practices. The paper supplies no measured cost of waiting and no evidence that a particular clarification route will succeed. The point is to keep the unresolved rule question visible when the company chooses how much additional work to undertake.

The decision record should preserve both source observations after a clarification arrives. A later response can explain how the programme authority intends to handle the specified case. It does

not change what the earlier publications said when they were consulted. Retaining the earlier observations supports a reconstruction of why the team sought clarification and which facts it supplied. If a programme later publishes a corrected version, that version should receive its own observation and applicability fields. The reviewer should then determine which assessments need updating under the corrected text.

The wider lesson is a design recommendation about documentary uncertainty. A source hierarchy and a version log help organise review, while material inconsistencies still need a disposition. The analyst should record whether the issue concerns contradictory text, a missing fact, an undefined term or a requirement whose decision date has not arrived. These categories lead to different next actions. A search for another convenient summary cannot establish the governing interpretation. The certificate should identify the precise question that remains open and the authority or evidence required to answer it.

Table 4. A decision date and deadline tracker

Item	Published or proposed date	Required disposition
EIC2027 full-proposal batches	20 January, 17 March, 1 September and 4 November 2027	Check exact call conditions and Brussels-time instructions before submission [4]
Hub71 cohort deadline	February 2027; no day stated in sampled page	Obtain the exact operative cutoff [13]
Hub71 review timing	Access says by April; FAQ says April-June 2027	Retain the disagreement and obtain cohort-specific clarification [13,14]
HHS submission-limit notice	Released 10 July 2026; due-date applicability from 13 April 2026	Retain release/applicability separately [12]
Company clarification review	Unassigned in this research	Management sets a real date after obtaining the case timetable
Company final submission approval	Unassigned in this research	Use the actual brief and readiness evidence; do not invent a deadline

Official facts observed 10 November 2026; this is a tracker structure, not a complete application calendar. Company dates are deliberately unassigned.

7. Calculate enterprise size within the correct boundary

The Commission's current SME summary combines staff below 250 with either turnover no greater than EUR 50 million or a balance-sheet total no greater than EUR 43 million. It also warns that group information may need inclusion. [6] The underlying recommendation distinguishes autonomous, partner and linked enterprises. Its aggregation provisions use proportional partner data and full linked-enterprise data, subject to the specified rules and avoidance of amounts already consolidated. Staff are measured using annual work units, with the relevant accounts and reference-period provisions retained. [7] These selected conditions require a properly defined calculation before a company can rely on a size description.

A hypothetical calculation illustrates the logical structure only. Assume a correctly defined standalone enterprise has 240 annual work units, turnover of EUR 52 million and a balance-sheet total of EUR 40 million in the applicable period. Under just the stated numerical conjunction, the staff condition is satisfied and the balance-sheet alternative is satisfied. The turnover number alone does not settle that selected test. This illustration excludes group relationships, timing provisions, exceptions and other programme requirements by assumption. It provides no full SME determination and should not be reused for an actual applicant without the omitted checks.

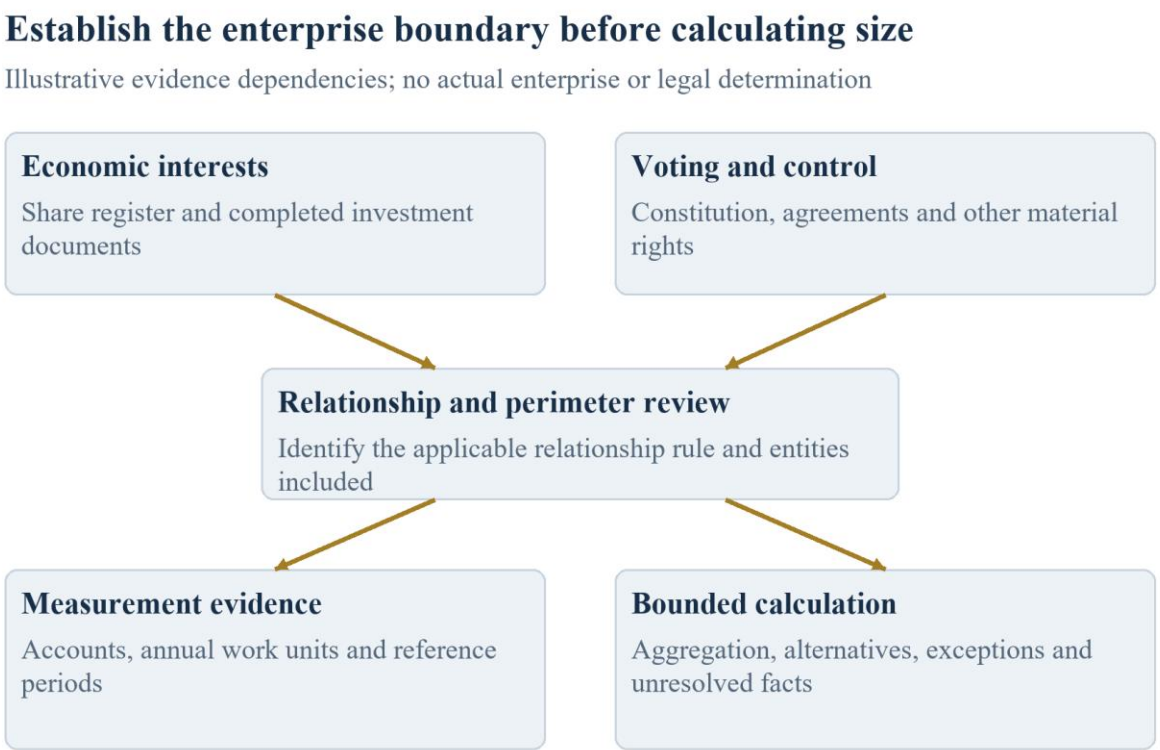
A second authored example concerns the calculation boundary. Suppose the same business has an investee or shareholder whose relationship may require aggregation. The reviewer must establish

the relationship under the applicable control and ownership provisions before deciding which figures enter the calculation. A cap table can identify economic interests while leaving voting arrangements or contractual influence unexamined. The evidence request should therefore state which rights and agreements are needed. Where the necessary relationship information is unavailable, the group-based size test remains unresolved even when the standalone accounts are complete.

We recommend documenting a reconciliation between the organisational perimeter and the numerical schedule. Each included entity should have a stated relationship, measurement period, aggregation treatment and source for the data used. Any adjustment for information already consolidated should be shown explicitly. A second reviewer should be able to follow the calculation without reconstructing the group from an informal pitch deck. The layout is an authored operating proposal. The proposed layout makes the analyst's choices visible for professional review.

Reference periods deserve the same attention as thresholds. The recommendation contains a two-consecutive-period threshold provision, and new enterprises have an estimate-based provision. [7] This selected reading does not establish an unconditional grace period after every acquisition or ownership change. An applicant planning a financing round or group transaction should obtain advice on the specific status-change facts and applicable rule. The proposed evidence record would preserve the current calculation, the contemplated change and the new assessment once the event occurs. It would avoid converting a general timing provision into an assurance unsupported by the actual transaction.

Figure 3. Review the enterprise boundary before its size calculation



Illustrative evidence dependencies. No actual enterprise or legal size determination is depicted.

8. UK admissibility scope and portfolio decisions

Innovate UK's standard assessment guidance separates the competition manager's eligibility check from the Innovation Lead's scope review. The scope review uses the answer to the scope question. Applications judged within scope proceed to assessment, including proposals the Innovation Lead considers on the edge of the scope requirements. Ineligible or out-of-scope applications receive reasons for exclusion. This published route gives the word borderline a precise constraint: a close scope judgement can coexist with admission to assessment. [1]

Application preparation also has an administrative dimension. The general guidance requires answers to all question criteria and compliance with its application standards. Most competitions close at 11:00 UK time, with departures stated in the brief. The portal closes at the deadline, and applicant-caused technical difficulties do not provide an extension basis under that guidance. The applicable competition brief therefore belongs beside the submitted application in the evidence record. [2]

Entity classification requires similar care. The funding rules distinguish the UK-registered-business category, qualifying registration numbers and a separate evidence route for sole traders. They also require ownership, interests and group associations to be considered when selecting business size. A proposed applicant's legal name, trading structure and group position should consequently be established before applying a category-specific test to its circumstances. The relevant guidance was updated on 5 November 2026. [3]

A useful preparation document is a competition-specific decision record with separate entries for entity category, project scope, submission requirements and the evidence supporting each conclusion. The person preparing it can attach the relevant wording and identify the responsible internal owner. Unresolved items need an explicit next action, such as obtaining a registration extract or asking the programme contact a narrowly framed scope question. A commercial team can then see which work is ready for approval and which commitment depends on further information. This approach also gives legal, technical and finance reviewers a common starting point.

Passing these checks leads to a further selection process. The standard guidance describes scored questions, ranking, a minimum quality threshold and any applicable portfolio approach. Portfolio allocation can produce a successful proposal with a lower score than another application. These provisions give an unsuccessful decision several possible procedural locations. The recorded decision and feedback are needed to identify the relevant location in an actual case. [1]

For management purposes, an eligibility conclusion should have a defined scope and date. A finance director might approve preparation expenditure after receiving satisfactory evidence for the applicant category while retaining a separate approval for the completed proposal. A technical lead might sign off the scope answer with its supporting explanation. The final submission owner can confirm that the required materials are present in the version being submitted. These are proposed internal controls. Their value is a readable account of who considered each issue, the information available at that point and the remaining conditions attached to spending.

9. US entity ownership control and affiliation

The NIH Grants Policy Statement revised in March 2026 describes a more-than-50% direct ownership-and-control route through specified eligible owners, including US citizens or

permanent residents and qualifying business or tribal categories. Its SBIR provisions also describe a majority-ownership alternative involving multiple venture-capital operating companies, hedge funds or private-equity firms, with no single such investor above 50%. The displayed STTR provisions exclude that investor alternative. Programme identity therefore matters when reading an ownership structure. [8]

A corporate-finance review can begin with the legal applicant and trace the people and entities holding its shares and decision rights. The supporting package can include the share register, signed investment documents, constitutional documents and the agreements allocating board or approval rights. Options, convertible instruments and pending transactions deserve their own descriptions, together with a request to establish how the applicable rule treats them. This is an evidence-gathering recommendation. An analyst should describe the instrument actually held and leave a disputed legal classification open for appropriate review. A percentage copied from a presentation needs reconciliation to the underlying documents.

Affiliation creates a further boundary around the information required. NIH's stated employee condition includes affiliates and sets a maximum of 500. Its guidance also describes control arising through ownership, management or contractual relationships. Where eligibility is unclear, NIH states that it will not make an award until SBA provides a determination. These provisions support a separate unresolved-evidence category in an internal review. [8]

The group review should therefore state the perimeter used for its calculations. A preparer can identify every entity considered, explain why an entity was included or left for further consideration, and retain the headcount evidence supplied for each one. Employment figures need a common definition and an identified reference period before they can be combined meaningfully. The applicable counting method should be confirmed from the relevant authoritative requirements. A finance team's existing consolidation schedule may be useful supporting material, although its accounting purpose should be stated clearly when it is used in a grant file.

Corporate transactions make the evidence date commercially important. A proposed investment may change the shareholding shown in a draft schedule, while completion documents establish a later position. Board rights can also change through a separate agreement. Management can ask for a dated description of the current position, the proposed transaction and the conditions still outstanding. The grant preparation owner can then seek advice on the checkpoints relevant to that transaction. A forecast cap table should remain identified as a forecast throughout this process, with a named person responsible for replacing it when reliable completion evidence becomes available.

The practical output is a scoped ownership and affiliation memorandum. It can state which route has been examined, list the supporting documents, show the calculations and isolate questions requiring agency or professional clarification. Approval can be tied to those stated assumptions. This gives an investment committee a usable basis for deciding how much preparation work to authorise and when a further review is needed. No conclusion about a particular company's legal eligibility or award prospects follows from the illustrative review structure developed here.

10. Versions phases and agency-specific limits

The Small Business Innovation and Economic Security Act was approved on 13 April 2026. Its extension provisions move SBIR and STTR authority to 30 September 2031, and NIST's current programme announcement confirms that extension. NIST also states that its website is being updated following reauthorisation. Programme authority and the availability of a particular application opportunity consequently require their own dated evidence. [10, 11]

NIH's notice released on 10 July 2026 rescinds its earlier April application-limit notice. It states an HHS-wide maximum of nine new and resubmission Phase I, Fast-track and Direct-to-Phase-II applications or proposals per small business per fiscal year for the specified submission classes. Its applicability covers due dates on or after 13 April 2026. The publication date, applicability date, agency perimeter and covered submission types all belong in the interpretation of this rule. [12]

Phase-specific conditions add another dimension. The posted SBA directive describes certification at award and other regulatory checkpoints. It states SBIR workshare minima of two-thirds in Phase I and one-half in Phase II, and STTR minima of 40% for the small business and 30% for the research institution. The solicitation is to specify the effort measurement basis. The portal contains legacy sunset wording, so these located passages require version caution when used alongside current law and agency instructions. [9]

A proposed rule register can keep one row for each requirement used in an application decision. Each row should identify the issuing body, programme, phase, document version, publication date, operative date and the precise applicant facts to which the requirement is being applied. A later notice can then be assessed for its effect on that row. The application owner can preserve the earlier entry and explain the adopted interpretation in a dated successor. This provides a clear record of the basis used when management authorised expenditure or approved submission.

For a submission limit, the corresponding internal record can list the organisation making each submission, its submission type, the agency component, the relevant fiscal year and the evidence of submission. The reviewer should establish how the applicable notice treats each item before calculating a total. A draft, a withdrawn application and a submitted resubmission may require distinct descriptions. The source material examined here does not resolve every possible counting question. A doubtful item can be raised in a concise written query with the facts and the proposed interpretation set out together.

For workshare, a project manager can prepare a responsibility schedule showing the proposed work, performing organisation and measurement basis for each activity. Finance can reconcile that schedule to the budget after the relevant definitions have been established. Any requested clarification should identify the exact activity and uncertainty. A final pre-submission review can check that the rule register, responsibility schedule and application refer to the same version of the project. These are recommended preparation steps, with their usefulness assessed through the quality of the resulting evidence and the decisions they support.

11. Separate Gulf programme application from participation

Hub71's current FAQ describes Pre-Seed, Seed and Series A applicants, including international startups and solo founders. It separately identifies factors considered in evaluation. The Access page describes its selection process and application requirements. The FAQ states that

incorporation is unnecessary before applying and describes setup following selection; the Access page says that at least one founder is expected to commit to long-term relocation. [13, 14] The stated expectation, application condition, selection judgement and subsequent participation step should retain their actual roles. The review supplies no signed participant agreement or applicant-specific offer.

A hypothetical founder reviewing that programme should separate existing facts from intended commitments. The file can record the current company's jurisdiction and stage, the founder's actual availability and the conditions described in the public materials. A plan to relocate should be identified as a plan until supported by the relevant action or agreement. The application team should ask what evidence is required at the stage it is entering and which later obligations would arise if selected. This avoids treating the absence of a later-stage document as a present failure without checking when that document is required.

The timing record also needs to preserve documentary precision. At the observation date, the Access page describes a February 2027 cohort deadline without a day. Its review timing is stated as by April, while the FAQ describes April to June. [13, 14] We retain that difference. A team should seek the exact operative cutoff and applicable cohort instructions before committing its submission timetable. The paper supplies no invented day, portal closing time or assumption that the broader period extends the narrower statement. A month-only date remains insufficient for an exact deadline calculation.

We propose a participation checklist with separate sections for application evidence, selection communications and terms accepted after an offer. The checklist should identify the instrument being discussed and the document that establishes any entitlement or obligation. The research does not treat an accelerator's public incentive description as a completed grant agreement or a universally available cash amount. For a financing decision, the company needs the actual offer, governing terms and any conditions affecting its receipt or use. Until then, the funding plan should keep the prospective programme route separate from resources already committed.

This Gulf example also shows why a cross-programme database needs flexible fields. A programme's published company-stage descriptors cannot be assumed to be equivalent to a statutory employee or ownership test elsewhere. The database should preserve the source terms and the analyst's classification separately. It can then support a comparison of evidence required at similar decision stages without pretending that all programmes share one legal definition. The resulting comparison is useful for planning document work, while actual eligibility and participation decisions remain programme-specific.

12. A hypothetical cross-border applicant evidence review

Consider a hypothetical UK business exploring an Innovate UK competition while a proposed US subsidiary considers a separate small-business research opportunity. This constructed example describes a preparation decision only. The group provides a draft organisation chart, a proposed ownership schedule, a preliminary work plan and a consolidated employee figure. Its investment round remains subject to signed documents. The two application files are at an early stage, and the group has received no eligibility decision or award. The immediate management question concerns the additional evidence needed before committing substantial proposal resources.

For the UK file, the preparation owner can first identify the exact applicant category and competition. The current brief can be attached to a decision record, together with the available registration evidence. The technical team can draft a scope answer that connects the proposed activities to the stated objectives. Where the connection remains uncertain, the team can formulate a specific question for the programme contact. The file can record the question, any reply and the scope interpretation retained for review. An internal description of the project as borderline supplies no substitute for that documented sequence.

For the US file, the proposed ownership schedule should be marked provisional. The reviewer can request the documents establishing the eventual ownership and control position, identify the route under consideration and record the relevant programme and phase. A schedule separating current facts from proposed transaction terms would make the remaining uncertainty visible to management. Questions about investor classification or control rights can be assigned to appropriate professional or agency review. Until the necessary documents are available, the internal record can state that an ownership conclusion remains unresolved and identify the evidence that would allow it to be reconsidered.

The consolidated employee figure also needs an explained perimeter. In this example, the finance team has supplied a group total without a list of included entities or a stated counting method. The application owner can request that breakdown and the supporting employment records, then establish the appropriate affiliation and counting rules for the selected route. A project responsibility schedule can separately identify the proposed contributions of the subsidiary, parent and research partners. These documents answer different preparation questions and can be reviewed by people with the relevant financial, legal and technical responsibilities.

The submission history can be assembled once the US opportunity has been identified. The group can list its relevant filings and ask the reviewer to establish which entries fall within the applicable agency rule. The UK application owner can maintain a separate checklist for the required materials and submission deadline. Management can then approve limited work on the open questions, attach a budget to that work and set a further decision point. Such an approval records a preparation commitment. Any later eligibility or funding decision would require the appropriate authority and the actual application evidence.

The document-specific action proposed here is a two-file evidence memorandum with a shared transaction appendix. Each application file would identify its applicant, opportunity, current rule versions, available evidence, unresolved questions, responsible reviewer and next decision date. The shared appendix would describe the proposed investment and distinguish reliable completed events from assumptions still awaiting evidence. This structure gives management an actionable record for allocating preparation effort while preserving the separate conditions governing each application.

13. Represent unresolved evidence explicitly

For the constructed model, we define six mandatory predicates: entity, control, size, activity, cost and submission. They are abstract categories for this exercise. An actual programme could require several separate checks within one category or additional categories altogether. Each predicate is attached to an identified rule, a fact snapshot and its evidence. We assign T when the declared facts support that selected requirement, F when they contradict it and U when the record

leaves the requirement unresolved. The states describe the bounded assessment. They carry no assumed probability and do not measure the applicant's commercial attractiveness.

We define the combined gate as F when any mandatory predicate is F, T when all six are T and U in the remaining cases. This conjunction makes the consequence of a mandatory condition explicit. A supported failing requirement is sufficient for a bounded rule-fail certificate. A rule-pass certificate requires support for every selected requirement. A file containing unresolved requirements and no supported failure remains unresolved. The definition does not assign a funding-body decision to the applicant. Its output concerns only the selected requirements and declared facts at the recorded assessment date.

The unknown state covers several distinct evidential reasons. The analyst may be waiting for an ownership agreement, unable to reconcile two official sources or reviewing a future fact that has not yet occurred. Those reasons should be recorded alongside U. The model uses a common symbol to calculate the conjunction, while the operational record retains the reason and next action. A missing document may lead to a request to the company. A material interpretation conflict may require programme clarification. A future condition needs its relevant checkpoint and responsible owner. The symbol is a compact calculation input; the accompanying record supplies its reason and next action.

Known states should remain traceable when evidence changes. If a newly executed agreement changes a material ownership fact, the team should create a successor assessment and retain the earlier conclusion with its facts. If the team discovers an error in its earlier interpretation, it should record a correction and its scope. The proposed process does not silently replace a prior unknown with a positive answer because the deadline is close. Nor does a commercial sponsor's willingness to proceed alter the evidential state. Management's preparation decision and the selected-rule assessment should be documented in their own fields.

This representation deliberately leaves scope judgements and scored merit in separate records. It does not turn a UKRI on-edge scope decision into a failed mandatory predicate. For a particular programme, a reviewer must first establish which conditions are mandatory, which are assessed and who makes each judgement. Only then can the relevant bounded propositions be encoded. The paper's numerical exercise begins after that interpretive work by assumption. It cannot automate legal interpretation, determine whether an actual document is genuine or certify that the six generic categories exhaust a real call's requirements.

Table 5. Three states for a selected mandatory requirement

State	Defined meaning	Next action in this design
T	Declared facts support the selected requirement	Retain the rule, facts and evidence
F	Declared facts fail the selected requirement	Record the decisive requirement and review the route
U	The selected requirement remains unresolved	Assign the missing evidence or interpretation question
Combined T	All six selected states are T	Issue a bounded rule-pass certificate
Combined F	At least one selected state is F	Issue a bounded rule-fail certificate
Combined U	No F and at least one U	Preserve unresolved status

Stipulated generic model semantics; no funding authority or award decision is encoded.

14. Test the logical consequences of compensating scores

The first exercise enumerates every six-position vector formed from T, F and U. There are three choices at each position, producing 729 distinct vectors. Every vector is counted once. This is a finite logical sensitivity exercise with equal enumeration weight. The vectors are neither a sample of applicants nor an estimate of the joint distribution of programme requirements. Some combinations may have no practical counterpart under a specific programme. The exercise asks how alternative definitions classify the same declared information, using the mandatory conjunction as the specified comparison baseline.

The first alternative maps U to T and then takes an ordinary Boolean conjunction. It provides a transparent illustration of treating unresolved requirements as satisfied. The second alternative assigns one point to T, half a point to U and zero to F, then compares the mean with a threshold. We pre-specify thresholds of one-half, two-thirds, three-quarters, five-sixths and one. The calculations use exact rational arithmetic so that a vector on a threshold is handled consistently. These alternatives are authored test rules. The study makes no claim that a funding agency or commercial product uses either rule.

For each alternative, a passing vector is classified according to the baseline conjunction. A pass with a baseline T is a supported pass. A pass with a baseline U is an unsupported pass because at least one required answer remains unresolved and none is known to fail. A pass with a baseline F is a contradicted pass because a selected mandatory requirement is known to fail. These are defined diagnostic categories. They should not be described as measured legal error rates, false-award rates or estimates of the accuracy of an operational screening system.

Several controls make the enumeration auditable. The model checks the all-positive, all-negative and all-unknown vectors, each otherwise-positive vector with one failure and each otherwise-positive vector with one unknown. It checks that rearranging the six positions leaves the conjunction unchanged. For every vector, it also replaces the unknowns with every possible positive or negative completion and checks the resulting Boolean outcomes. A baseline T must remain positive under completion, a baseline F must remain negative and a baseline U must permit both. Those controls verify the implementation of the declared semantics.

A reader applying the finding should inspect the role of a proposed score before using it. Scores can organise assessed merit or prioritise review work when their meaning is explicit. The defined exercise concerns whether positive answers can compensate for an unresolved or failed mandatory requirement. Its answer follows from the stipulated conjunction and the tested alternatives. An empirical evaluation of a real tool would require that tool's actual rules, labelled cases, the applicable programme versions and a separate validation protocol. The present exercise provides a reproducible diagnostic that can inform that later evaluation.

15. Define the evidence checks and their assumed work

The second exercise uses twelve authored cases. Each case contains six initial states and six declared resolution outcomes in the same category order. A resolution outcome tells the simulator what a covered unknown reveals when a specified check is performed. The analyst's online policy cannot inspect that outcome beforehand. Known initial positive or negative states are required to agree with the declared outcomes. Some outcomes remain U after all available checks. This

permits the exercise to represent a selected document set that leaves a conflict unresolved, without interpreting that conflict as a substantive rule failure.

We define six checks and assign invented evidence work units. An entity-register check costs two units and covers entity. An ownership dossier costs four and covers control and size. A headcount return costs three and covers size. A technical workplan costs three and covers activity and cost. A cost ledger costs two and covers cost. A submission record costs one and covers submission. The coverage and work values are assumptions selected to expose overlap and differing review sequences. They are not staff hours, prices, reported advisory costs or estimates of the burden of any named programme.

When a check is performed, it reveals only the declared outcome for covered predicates that are currently unknown. It preserves known states. An outcome that remains unknown still consumes the check's work. The same check cannot be repeated. Different checks can overlap; for example, both the ownership dossier and the headcount return can supply the declared size outcome. This simplification assumes that the covered documents agree on that outcome. A real contradictory-document case would need a richer evidence model with separate observations and an explicit reconciliation rule. The unresolved cases retain U.

The case set includes a complete initial file, a known entity failure, a known cost failure, individual gaps that resolve positively or negatively, shared document coverage, fully incomplete files and conflicts left unresolved. The cases are deliberately varied stress tests. They are not weighted by prevalence or selected to estimate an average applicant burden. Initial certificates require zero additional work within this model. That zero excludes historical effort already embodied in the initial file. It should be read as work added after the declared starting point, with no implication that real evidence was costless to create.

For each case, the model enumerates all 64 subsets of the six available checks, including the empty subset. It calculates the resulting gate and the sum of the selected checks' costs. Among subsets producing a positive or negative certificate, it records the minimum additional work and every tied minimum witness. The calculation has access to the declared resolution outcomes and is therefore ex post. Its minimum is a lower bound for comparison with a feasible review sequence. It is not a policy that an applicant could implement while knowing outcomes that have not yet been observed.

Table 6. Assumed evidence checks and additional work units

Check	Work units	Covered unknowns
Entity register	2	Entity
Ownership dossier	4	Control and size
Headcount return	3	Size
Technical workplan	3	Activity and cost
Cost ledger	2	Cost
Submission record	1	Submission

All work units and coverage are authored assumptions. They represent no currency, staff hours or observed preparation costs.

Table 7. Declared states in the twelve authored cases

Constructed case	Initial states	Resolution outcomes
Complete bounded file	TTTTTT	TTTTTT
Known entity mismatch	FTTTTT	FTTTTT
Entity evidence gap	UTTTTT	TTTTTT
Control evidence reveals failure	TUTTTT	TFTTTT
Shared group dossier	TUUTTT	TTTTTT
Shared technical workplan	TTTUUT	TTTTTT
Initially incomplete file	UUUUUU	TTTTTT
Incomplete file with a failing control	UUUUUU	TFTTTT
Control conflict remains unresolved	TUTTTT	TUTTTT
Activity evidence remains unresolved	TTTUUT	TTTUTT
Known cost-rule failure	TUTTFU	TTTTFT
Size evidence reveals failure	UTUTTU	TTFTTT

Position order: entity, control, size, activity, cost, submission. T=supported; F=failed; U=unresolved. Outcomes are authored case assumptions, revealed to online policies only by performed checks.

16. Compare feasible review orders with the certificate bound

We specify two feasible online policies before execution. Entity-first follows the entity register, ownership dossier, headcount return, workplan, cost ledger and submission record. Cheap-first sorts checks by their assumed cost, using that same declared index to resolve ties. Its sequence is submission record, entity register, cost ledger, headcount return, workplan and ownership dossier. These fixed orders make the comparison reproducible. They do not represent an estimate of professional best practice, and the model provides no claim that either is globally optimal for unknown real-world cases.

Each policy begins with the declared initial states. It stops immediately if those states already establish a certificate. Otherwise it scans its order once. A check covering no remaining unknown is skipped. A relevant check is performed if its cost fits the remaining work budget, including an exact-budget tie. An unaffordable check is skipped and a later affordable check may still run. The policy stops when it establishes a positive or negative certificate or reaches the end of the scan. This rule exposes the consequences of affordable work that consumes capacity before a decisive later check.

Budgets range from zero through fifteen work units. Fifteen equals the total cost of all six checks, although a policy may stop earlier or skip redundant checks. The model retains the executed checks, skip reasons, expenditure and state transition after every step. A still-unresolved result is reported as unresolved regardless of work spent. For every online result that establishes a certificate, the implementation checks that the expenditure stays within budget and is at least the ex-post minimum. A policy certificate paired with an ex-post unresolved minimum would indicate an inconsistency requiring investigation.

Ex-post feasibility has a simple budget property: once a certifying subset is affordable, it remains affordable at larger budgets. The fixed online policies have a different property because a larger budget can change which earlier checks they buy. An earlier purchase can leave insufficient capacity for a later decisive check. We therefore inspect the full budget sequence and retain any loss of certification as the budget rises. Such a result is a feature of the specified skip-on-budget policy, with no implication that a larger real advisory budget reduces an applicant's legal eligibility.

The implementation is deterministic, with fixed inputs and no random seed, fitted parameter or learned classifier. Design, inputs and code are separately identified, and the executed results are retained for all tables and figures. Later editorial or visual work should reuse those saved arrays. Any correction would require a preserved prior result and a new version with an explained change. This discipline lets a reader distinguish the assumed case design from the consequences calculated under that design and from any later empirical claim requiring different evidence.

17. Read the logical diagnostic within its artificial state space

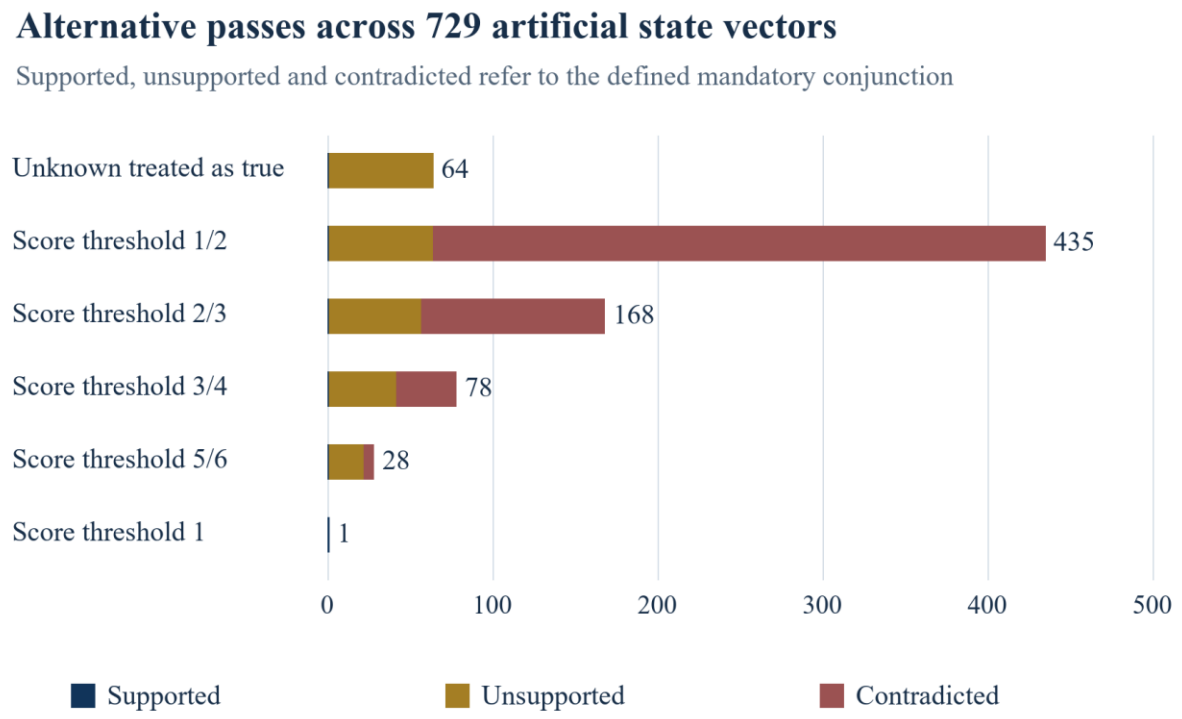
The complete enumeration contains one baseline rule-pass vector, 63 unresolved vectors and 665 rule-fail vectors. The counts follow the declared conjunction: the 64 vectors containing only T and U include one all-T vector, and every remaining vector contains at least one F. Mapping unknowns to true passes all 64 of those T/U vectors. It therefore produces one supported pass and 63 unsupported passes, with no contradicted passes. The result describes what that authored alternative does to unresolved information. It provides no estimate of how often an applicant file contains an unknown requirement.

The half-point score at a one-half threshold passes 435 vectors: one supported, 63 unsupported and 371 contradicted. At two-thirds, the corresponding counts are one, 56 and 111, giving 168 total passes. At three-quarters, there are one supported, 41 unsupported and 36 contradicted passes. At five-sixths, one supported, 21 unsupported and six contradicted vectors pass. A threshold of one passes only the all-T vector. These are saved results from the exact rational comparison, and every pass category is defined relative to the specified mandatory conjunction.

The six contradicted passes at five-sixths have a simple interpretation. Each has one F and five T states. The positive scores compensate for the failing mandatory predicate under the alternative mean-score definition. The conjunction classifies the same vector as rule-fail. This worked boundary makes the choice of rule visible. For a real programme, the reviewer must identify whether a condition is mandatory or assessed and then implement the governing rule accordingly. The artificial comparison cannot establish how an actual funding body handles a particular application or whether a commercial screening system contains such a scoring rule.

The result also explains why a high overall screening score requires a stated interpretation. The score summarises a set of answers. Each mandatory proposition requires its own supporting evidence. We recommend reporting unresolved and failed mandatory conditions individually, with merit assessment in its separate record. Where a company wants a score for prioritising its own review work, it should describe the score's purpose and retain the original states. The present test supplies a checkable example of information lost when a selected mandatory gate is replaced by an explicitly different compensating calculation.

Figure 4. Passes under the declared alternative screening rules



Saved deterministic model results. Counts are logical diagnostics; they are not applicant frequencies or real-tool accuracy rates.

Table 8. Pass classifications across the 729 artificial vectors

Alternative rule or threshold	Supported	Unsupported	Contradicted
Unknown treated as true	1	63	0
1/2	1	63	371
2/3	1	56	111
3/4	1	41	36
5/6	1	21	6
1	1	0	0

Exact saved model output; every count is a vector count relative to the stipulated mandatory conjunction. No population probability or actual-tool accuracy is estimated.

18. Read certificate work by the declared starting file

Three constructed cases already have a certificate at the starting point: the complete bounded file, the known entity mismatch and the known cost-rule failure. Each requires zero additional evidence work in this model. The first is rule-pass; the other two are rule-fail. Their inclusion is important when reading the aggregate counts. A certificate count includes a documented failure as well as a supported pass. It is not a count of applications eligible for funding, and the zero-work results exclude the historical effort represented by the initial evidence.

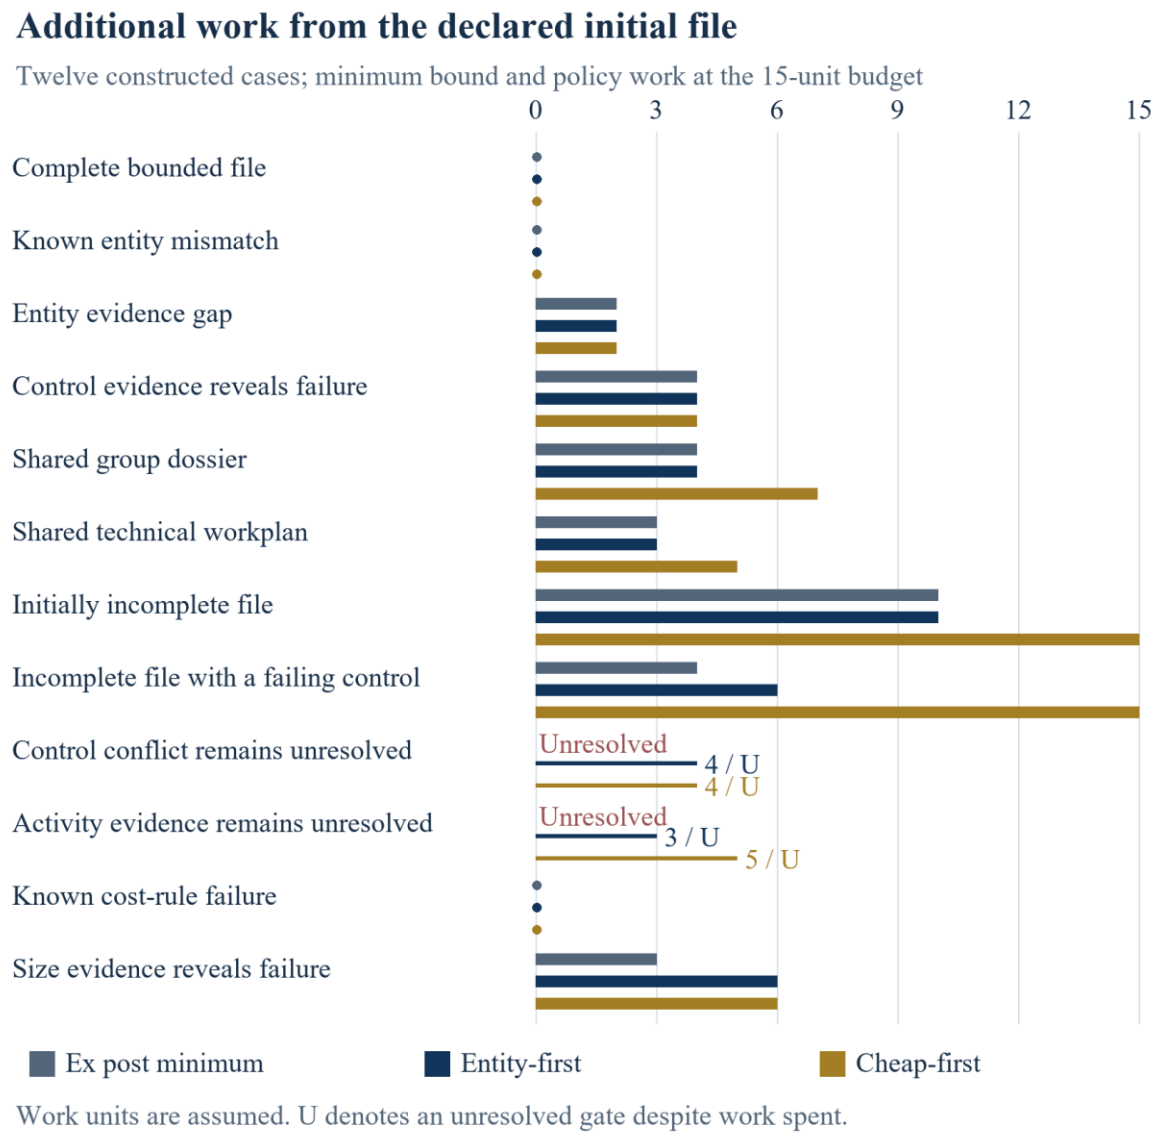
Among the initially unresolved cases, the entity evidence gap has a minimum of two work units, the control failure and shared group dossier each require four, and the shared technical workplan requires three. The initially incomplete file with all positive outcomes requires ten units. The incomplete file with a hidden control failure has an ex-post minimum of four. The size-failure case requires three, using the headcount return. The control-conflict and unresolved-activity cases have null minima because none of the selected check subsets can resolve their gate. Those nulls are retained separately from zero.

At the full fifteen-unit budget, both fixed policies reach certificates in ten of twelve cases, including the three initial certificates. Their work differs. In the shared-dossier case, entity-first uses four units, while cheap-first spends three on headcount before the four-unit ownership dossier, for seven total. In the shared-workplan case, entity-first uses three units; cheap-first first buys the two-unit ledger and then the three-unit workplan, for five. The assumed overlap makes some earlier checks redundant once the later bundle is purchased. This mechanism follows from the declared coverage, with no observed advisory expenditure implied.

The incomplete all-positive file requires ten units under entity-first and fifteen under cheap-first. The hidden-control-failure file requires six and fifteen respectively, against its four-unit ex-post bound. Entity-first reviews the register before discovering the failing control through the ownership dossier. Cheap-first reaches that dossier after the other relevant purchases. The ex-post minimum identifies the decisive document using known case outcomes; the online policies lack that advance knowledge. The comparison therefore records a cost difference under the chosen orders. It establishes no globally optimal online strategy or estimated saving for a real applicant.

For the two unresolved cases, additional work still has an observable cost within the model. The control conflict consumes four units under either policy and remains U. The activity case consumes three units under entity-first or five under cheap-first and also remains U. The output should therefore report both work spent and unresolved status. Management receiving such a record would need a different evidence route, an authoritative clarification or a decision to discontinue the selected review. The model says only that its six checks cannot settle these declared cases; it supplies no conclusion about every possible real-world investigation.

Figure 5. Additional evidence work and the ex post certificate bound



Twelve authored cases and invented work units. Initial certificates require zero additional work. Null minima denote unresolved cases within the selected checks.

Table 9. Additional evidence work at the fifteen-unit policy budget

Constructed case	Ex post minimum	Entity-first work / state	Cheap-first work / state
Complete bounded file	0	0 / T	0 / T
Known entity mismatch	0	0 / F	0 / F
Entity evidence gap	2	2 / T	2 / T
Control evidence reveals failure	4	4 / F	4 / F
Shared group dossier	4	4 / T	7 / T
Shared technical workplan	3	3 / T	5 / T
Initially incomplete file	10	10 / T	15 / T
Incomplete file with a failing control	4	6 / F	15 / F
Control conflict remains unresolved	Unresolved	4 / U	4 / U
Activity evidence remains unresolved	Unresolved	3 / U	5 / U
Known cost-rule failure	0	0 / F	0 / F
Size evidence reveals failure	3	6 / F	6 / F

Invented additional work units. T/F are bounded rule-pass/rule-fail certificates; U remains unresolved. Three initial certificates have zero additional work. Null minima are not zero or a claim of real-world impossibility.

19. Inspect the full budget sequence before choosing an order

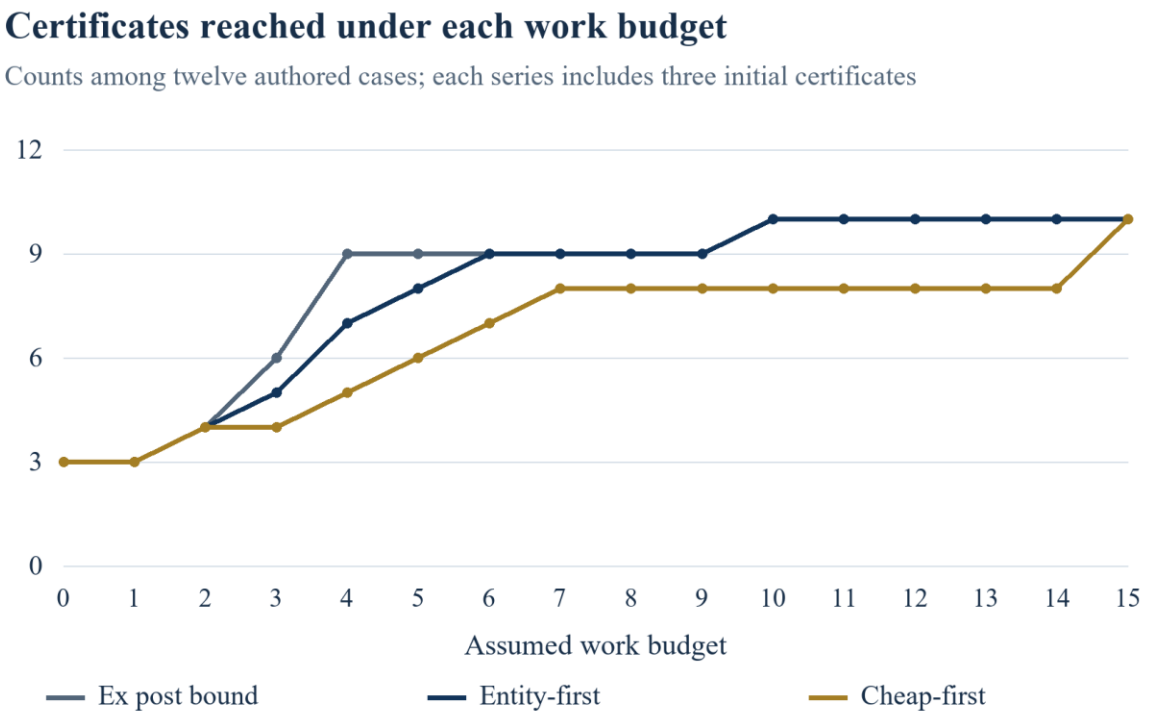
At a zero-unit budget, the ex-post calculation and both policies retain the three initial certificates. At a budget of three, six cases have an affordable ex-post certificate; entity-first reaches five and cheap-first reaches four. At four units, nine cases have a certifying subset, while entity-first reaches seven and cheap-first reaches five. The difference reflects the checks purchased in sequence before a certificate is available. The table reports every budget from zero to fifteen so that readers can examine thresholds and flat regions without relying on a selected comparison.

Entity-first reaches nine certificates at six units and ten at ten units. Cheap-first reaches eight at seven units and stays at eight through fourteen. At fifteen, it reaches ten. The ex-post total reaches ten at ten units. Each total includes three initial certificates; the remaining certificates are acquired by the checks. The two persistent unresolved cases keep the maximum below twelve. These counts are over the fixed authored case set. There is no basis for converting ten of twelve into a predicted rate for applicants, a geographic comparison or an expected portfolio award outcome.

No adjacent-budget loss of certification occurs for either fixed policy in these twelve cases. The specified policy still permits such behaviour in other case designs, because a newly affordable earlier check can change the remaining capacity available for a later one. We retain that methodological distinction and report the observed absence of an irregularity in this declared set. The ex-post feasibility bound has a separate mathematical monotonicity property: an affordable certifying subset stays affordable as budget increases. A policy that chooses checks sequentially must be evaluated using its actual information and purchase rules.

The saved run passed 2,115 deterministic controls covering logical properties, input consistency, trace budgets, lower bounds and the stated boundary checks. These controls test the implementation against the declared model. They do not validate the assumed evidence costs, establish that the six categories cover a real programme or measure legal interpretation accuracy. The practical next step for an organisation is a pilot using its actual document work and approved rules. Any proposed policy comparison should preserve the initial information, the tasks performed and unresolved outcomes so that the resulting claims have a defensible evidential basis.

Figure 6. Certificates achieved within each assumed work budget



Counts over twelve authored cases include three initial certificates. They are not award probabilities. No adjacent-budget loss of certification occurred in these twelve cases.

Table 10. Budget frontier across the twelve authored cases

Budget	Ex post certifiable	Entity-first	Cheap-first
0	3	3	3
1	3	3	3
2	4	4	4
3	6	5	4
4	9	7	5
5	9	8	6
6	9	9	7
7	9	9	8
8	9	9	8
9	9	9	8
10	10	10	8
11	10	10	8
12	10	10	8
13	10	10	8
14	10	10	8
15	10	10	10

All counts include three initial certificates. Each policy result may be rule-pass or rule-fail. New certificates equal the displayed total minus three. No adjacent-budget loss occurred in this case set.

20. Issue a certificate that another reviewer can reconstruct

We propose a decision certificate consisting of the application identity, assessed requirements, fact snapshot, evidence references, conclusion and authorised next action. The conclusion should state its scope in ordinary language. For example, a reviewer could confirm that the selected size calculation is supported under the identified accounts and group perimeter while leaving another requirement unresolved. The certificate should distinguish that limited conclusion from a complete eligibility determination. Its purpose is to give the next reviewer enough information to

reconstruct the reasoning without depending on the original preparer's memory or an informal conversation.

The decisive evidence deserves a specific locator. A source reference can identify a clause, page, table or named section. A company document can identify the relevant schedule and period. The assessment should show how the evidence supports the proposition being tested. A file name alone leaves the reasoning incomplete. The record should also disclose which potentially material evidence was unavailable and whether the preparer relied on a declared assumption. If an assumption affects the conclusion, the certificate should explain what would change if that assumption proved incorrect.

Provenance provides a useful vocabulary for organising this record. The W3C PROV Primer describes relationships among entities, activities and agents involved in producing information. It is a Working Group Note, and its concepts provide a design reference here. [16] We propose representing the source document as an entity, the review as an activity and the preparer or approver as an agent. A derived assessment would retain its relationship to those inputs. That structure records how a conclusion was produced. The substantive interpretation still requires a qualified reviewer and the applicable programme authority.

A successor certificate should identify the event that triggered reassessment. Such an event might be a completed investment, a corrected source, a new programme response or a revised project scope. The reviewer should state which propositions were reconsidered and which supporting facts remained unchanged. Preserving that lineage allows management to distinguish a new fact from a correction of earlier reasoning. It also creates a usable audit trail if a board paper or financing model needs to explain why an application's status changed between reporting periods.

We recommend retaining the official outcome letter separately from the company's certificate. The company's record explains its own review and spending decision. The official correspondence records the programme's decision within its stated scope. Differences between them should be investigated and documented. A real review process should specify access controls and retention appropriate to the commercial and personal information involved. The present proposal does not prescribe a universal legal retention period, and it makes no measured claim that a particular system reduces disputes or changes award outcomes.

Table 11. Fields in the proposed decision certificate

Field	Content to retain	Reason
Application identity	Entity, project, call and component	Locate the assessed case
Rule snapshot	Version, clause and applicability	Locate the governing proposition
Fact snapshot	Period, source and material assumptions	Identify what was assessed
State and reason	T, F or U with explanation	Preserve the conclusion's scope
Decisive evidence	Exact document locator	Support reconstruction
Open issue	Missing fact or unresolved interpretation	Assign the next action
Approval	Reviewer, authority and decision date	Record who authorised the action
Successor link	Changed fact, correction or later decision	Preserve the earlier position

Proposed operating design; substantive review and funding authority remain separate.

21. Constrain AI assistance to reviewable evidence work

An AI-assisted workflow could help propose document classifications, extract candidate dates or locate passages relevant to a predefined requirement. We propose treating each such output as a draft observation requiring verification. The workflow should retain the source passage, its location and the extraction that the reviewer actually checked. A model-generated statement about ownership, legal status or programme eligibility should remain outside an approved certificate until an authorised person has verified the underlying facts and interpretation. This is a proposed control design. No AI system was tested for extraction accuracy or application screening in this study.

A practical implementation can separate retrieval, extraction and judgement. Retrieval records which source was obtained. Extraction identifies the candidate text and fields. Judgement explains how the verified passage applies to the applicant and decision date. Keeping those records separate makes a disagreement local: the team can identify whether the wrong document was retrieved, the passage was misread or the rule was misapplied. It should also retain missing-source and access-failure states. A system that cannot obtain a governing document should record that limitation and direct the question to a human owner.

A pilot should begin with an independently reviewed set of documents and clearly defined tasks. An extraction task might ask whether the source states a day, a month or no closing date. Another could ask which legal entity or funding component a passage names. The reviewer can compare the proposed answer with the actual passage and retain the mismatch. These task-level observations can support a later evaluation. They should not be combined into an invented overall eligibility accuracy figure, particularly where the governing interpretation remains disputed or the applicant's material facts are unavailable.

The system should also preserve the distinction between public programme material and confidential applicant evidence. Before using any external processing route, management needs an approved basis for sending the particular information through that route. The proposed design can operate with a controlled document repository and a review queue whose permissions reflect the sensitivity of each file. This paper supplies no assurance about a provider's security or a company's compliance arrangements. Those questions require their own assessment of the actual service, contractual terms, data and applicable obligations.

Automation should have a defined release condition for each field it proposes to update. A verified date can be accepted with its source and precision. A conflicting rule should open an issue for review. A changed company document should trigger the requirements mapped to it for reassessment. The responsible reviewer should approve the resulting certificate and retain the basis for that approval. These are proposed operating rules that can be evaluated using actual review records. Their commercial usefulness should be measured through demonstrated document quality and decisions supported, with any time or cost effects established separately.

22. Design an applicant study that can explain actual outcomes

A future empirical study needs a defined population of application attempts and a dated record of what happened to each attempt. We propose enrolling consenting applicants before their final submission decision and following their selected calls through the relevant stages. The unit of observation should identify the company, project, programme, round and application version.

Repeated applications by the same company should remain linked. This design would allow researchers to distinguish a revised attempt from a new project and to identify preparation that stopped before an official application existed.

Outcome categories should be specified before collection. They could include discontinued preparation, no submission by the intended cutoff, administrative exclusion, a stated eligibility or scope exclusion, an assessed but unfunded proposal, a selected proposal and a later executed agreement. Actual feedback should support the stated reason. Missing feedback should remain recorded as missing. Multiple reasons may apply to one decision, and the coding protocol should explain how they are recorded. A second coder should review ambiguous cases using the underlying documents and retain any disagreement.

The study should preserve dates and censoring explicitly. A company still waiting for a decision at the end of observation has a pending outcome. A company that withdraws consent creates a different form of missingness. A programme that has not completed assessment cannot be treated as having rejected the applicant. The analysis should publish the relevant denominator for every reported proportion and show which cases are excluded and why. Where applicant consent limits the available evidence, the researcher should describe that limit before interpreting differences between groups.

Howell's published study of R&D grants uses ranked US Department of Energy SBIR applicants in a quasi-experimental evaluation. The publisher's abstract supports its role as an example of a design using applicant information to study financing outcomes. [15] Our documentary review does not reproduce that design or transfer its reported effects to eligibility review. A prospective study of the process proposed here would need its own identification strategy. Firms choosing more extensive preparation could differ in resources, technology or project quality, making a simple comparison of outcomes difficult to interpret causally.

We propose separating descriptive and causal objectives in that future work. The descriptive study can document the prevalence of specific missing records and the stages at which stated exclusions occur, within its enrolled population. A causal evaluation of a review intervention would need a defensible assignment or comparison design, ethical safeguards and adequate data. Commercial milestones such as a later equity round, revenue or project completion should be recorded with their own dates and evidence. Their occurrence after a grant decision alone would leave attribution unresolved. The present paper supplies no empirical estimate for any of those relationships.

23. Connect the evidence position to a financing decision

The finance team needs a funding plan that distinguishes available resources from routes still under consideration. We recommend giving the grant application its own stage record and linking each assumed financing event to supporting evidence. Preparation approval, submission, assessment, agreement and payment should have separate entries. Management can then see which project commitments depend on an external decision and which can proceed using existing resources. The proposed presentation provides transparency about dependencies. It does not establish an accounting recognition policy, lender acceptance or a forecast of what the programme will pay.

For a hypothetical hardware startup, management might be choosing whether to fund an additional test campaign before an application is assessed. The decision memorandum should identify the test's commercial purpose, the evidence it would produce and the source of cash for the expenditure. If the test also supports a customer pilot, that use should be described separately from its potential relevance to a grant. The team can compare the proposed expenditure with its actual liquidity and alternatives. The paper assigns no invented customer value, award probability or return to that illustrative choice.

The preparation budget should have an owner and a stopping rule. We propose approving named evidence tasks with a review point after the decisive information becomes available. A known mandatory mismatch can justify reconsidering the route immediately. An unresolved interpretation can justify a bounded clarification task. A complete selected-rule file can support a decision to proceed to the next stage, subject to the remaining requirements. These actions should be authorised in the company's own governance process. A rule-pass certificate describes the evidence checked and cannot create an external funding commitment.

An investment or acquisition can affect several application files at once. The transaction team should identify any application assumptions that depend on entity, ownership, project location or group size and assign review at the relevant checkpoints. The grant preparation owner can then obtain the completed transaction facts and update the affected assessments. The recommendation is to make the dependency visible during transaction planning. This paper makes no universal claim that a transaction changes eligibility, requires consent or triggers a particular contractual consequence. Each such conclusion needs the actual governing terms and case facts.

The final management paper should state the current decision, its evidential basis and the next commitment requiring approval. A concise unresolved-issues schedule can identify the responsible person, required document or clarification and the relevant decision date. This gives the board an actionable view without compressing every open question into a single probability. Where the evidence does not support a funding assumption, the team should retain that limitation in its plan. Subsequent changes can be reported against the preserved earlier position so that management understands which fact or decision changed.

Table 12. Decision ownership and funding-plan treatment

Responsibility	Proposed approval	Funding-plan record
Finance	Preparation budget and liquidity support	Approved expenditure and actual cash source
Technical owner	Project facts and tested claims	Evidence produced and remaining tests
Legal or qualified adviser	Scoped interpretation review	Opinion scope and unresolved conditions
Submission owner	Exact package and required approvals	Submission receipt or unfinished work
Funding body	Programme decisions within its authority	Actual decision, agreement and payment evidence

Recommended role structure; adapt to the company's actual authority arrangements. No projected award is treated as cash received.

24. Act on the specific question that remains unresolved

The documentary findings and constructed results support a scoped operating recommendation. Begin with the exact applicant and call. Identify which conditions are mandatory and when they apply. Attach the source and relevant facts to each selected proposition. Preserve unresolved

evidence and record any supported failure. Keep scope judgement, merit assessment and the eventual funding decision in separate records. The team can then authorise a specific next action with an owner and a review point. This recommendation concerns the quality of the decision record, with no promised award or measured financial benefit.

The constructed exercise exposes two distinct sources of decision error to investigate in practice. A scoring definition can permit a positive aggregate result despite a failed mandatory condition. A review sequence can spend effort on overlapping evidence before reaching the document decisive for the selected case. Their practical significance depends on actual programme rules, evidence costs and case distributions that this study does not observe. An empirical pilot should measure those quantities directly and retain unsuccessful or unresolved attempts, including work abandoned before submission. Successful-award records alone cannot supply that denominator.

For a company facing a live boundary question, the immediate deliverable should be a short, reviewable issue memorandum. It should state the unresolved clause or fact, the documents already checked, the authority required and the commitment waiting on the answer. A clarified interpretation or newly available fact can then enter a successor assessment. The company retains an intelligible record of its decisions, while the funding body retains its own authority over eligibility, selection, agreement and payment.

Frequently Asked Questions

Does a borderline application automatically fail?

The sources do not support that generalisation. Innovate UK's standard process sends applications considered on the edge of scope into assessment when judged within scope. An actual unsuccessful decision needs its stated reason and stage before it can be attributed to eligibility.

What is the difference between failed and unresolved eligibility evidence?

In the proposed framework, a failed state means declared facts contradict a selected requirement. An unresolved state means the record cannot settle it. Missing ownership documents or conflicting official text can remain unresolved while the company requests evidence or clarification.

Can a high screening score override a mandatory condition?

The defined mandatory conjunction requires every selected condition to be supported. The constructed score alternatives use a different definition and can pass vectors with failed or unresolved conditions. These are logical diagnostics, with no claim about an actual agency or software product.

What do the evidence work units measure?

They are invented units assigned to six checks for the constructed study. They represent no currency, staff hours, actual applicant burden or advisory fee. Zero additional work means the declared initial file already establishes a bounded certificate.

Does a rule-pass certificate establish that funding is available?

It concerns only the selected requirements, declared facts and rule snapshot. Scope, merit, selection, agreement and payment retain separate evidence and authority. The paper supplies no applicant-specific legal determination or award probability.

How should conflicting official sources be handled?

Preserve both versions, identify why the conflict matters to the exact application and request clarification from the responsible programme authority. Record the response with its date and scope. Until the issue is settled, keep the affected requirement unresolved.

References

[1] UK Research and Innovation. What happens after you have submitted your application. Standard Innovate UK assessment guidance; updated 7 May 2026. Accessed 10 November 2026. https://www.ukri.org/councils/innovate-uk/guidance-for-applicants/general-guidance/what-happens-after-you-have-submitted-your-application/

[2] UK Research and Innovation. Before you start an application. Updated 8 May 2026. Accessed 10 November 2026. https://www.ukri.org/councils/innovate-uk/guidance-for-applicants/general-guidance/before-you-start-an-application/

[3] UK Research and Innovation. Funding rules. Updated 5 November 2026. Selected organisation and business-classification sections. Accessed 10 November 2026. https://www.ukri.org/councils/innovate-uk/guidance-for-applicants/general-guidance/funding-rules/

[4] European Commission. European Innovation Council Work Programme 2027. Decision label C(2026)7629; official resource published 5 November 2026. Selected Section V, Table 8 and Annex 2 provisions. Accessed 10 November 2026. https://eic.ec.europa.eu/document/download/ed9ab103-1ffb-4d5f-a6ad-168ac3b36900_en?filename=EIC-2027-Work-Programme.pdf

[5] European Innovation Council. FAQs for EIC Accelerator. Selected full-proposal and Seal of Excellence answers; published inconsistencies retained. Accessed 10 November 2026. https://eic.ec.europa.eu/eic-frequently-asked-questions/faqs-eic-accelerator_en

[6] European Commission. SME definition. Current summary table and group-data explanation. Accessed 10 November 2026. https://single-market-economy.ec.europa.eu/smes/sme-fundamentals/sme-definition_en

[7] European Commission. Commission Recommendation 2003/361/EC of 6 May 2003 concerning the definition of micro, small and medium-sized enterprises. Official Journal L 124, 20 May 2003, pp 36-41. Annex Articles 2-6. Official Publications Office text accessed 10 November 2026. https://publications.europa.eu/resource/cellar/6ca8d655-126b-4a42-ada4-e9058fa45155.0004.01/DOC_1

[8] National Institutes of Health. NIH Grants Policy Statement, section 18.5.2 Eligibility. Revised March 2026; selected agency restatement, predating April reauthorisation. Accessed 10 November 2026. https://www.grants.nih.gov/grants/policy/nihgps/HTML5/section_18/18.5.2_eligibility.htm

[9] U.S. Small Business Administration. Policy Directives. Selected section 6(a) provisions; posted text contains legacy sunset wording, so current law and the applicable solicitation require separate review. Accessed 10 November 2026. https://www.sbir.gov/about/policies

[10] National Institute of Standards and Technology. Small Business Innovation Research Program. Current reauthorisation announcement. Accessed 10 November 2026. https://www.nist.gov/tpo/small-business-innovation-research-program-sbir

[11] United States. Public Law 119-83, Small Business Innovation and Economic Security Act. Approved 13 April 2026; selected sections 4, 9 and 10. Accessed 10 November 2026. https://www.govinfo.gov/content/pkg/PLAW-119publ83/html/PLAW-119publ83.htm

[12] National Institutes of Health, Office of the Director. NOT-OD-26-090, Notice of Information: Implementation of HHS Annual SBIR and STTR Application/Proposal Limit. Released 10 July 2026. Accessed 10 November 2026. https://www.grants.nih.gov/grants/guide/notice-files/NOT-OD-26-090.html

[13] Hub71. Access Programme. Current application requirements and cohort timeline; month-only cutoff and review-timing inconsistency retained. Accessed 10 November 2026. https://www.hub71.com/program/access-programme

[14] Hub71. FAQs. Selected application, selection, timeline and onboarding sections. Accessed 10 November 2026. https://www.hub71.com/faqs

[15] Howell ST. Financing Innovation: Evidence from R&D Grants. American Economic Review. 2017;107(4):1136-1164. doi:10.1257/aer.20150808. Publisher abstract consulted 10 November 2026; no full-paper reanalysis or transfer of estimated effects. https://www.aeaweb.org/articles?id=10.1257/aer.20150808

[16] Gil Y, Miles S, editors. PROV Model Primer. W3C Working Group Note, 30 April 2013. Provenance concepts used as a design reference; the primer is not a funding-programme rule. Accessed 10 November 2026. https://www.w3.org/TR/prov-primer/

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

https://www.linkedin.com/in/ckadya/

https://www.matchpoint-partners.com/team/ck-adya.html

Disclaimer

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.