

Qualification and Expertise: The Detectable Signals of Professional Competence Distance

Semantic Distance, Realisable Resolution Value, and the Allocation of Professional, Advisory, and Court-Appointed Roles

Qualificação e Especialização: Os Sinais Detetáveis da Distância de Competência Profissional

Distância Semântica, Valor de Resolução Realizável, e a Atribuição de Funções Profissionais, Consultivas e Periciais

Working Paper — SSRN Submission Draft

June 2026

Steven Carroll, Founder, KYC.co / Moral.Money

kyc.co | kyc.co/expertise-check | [SSRN #6996898](https://ssrn.com/abstract/6996898)

Abstract

Institutions that allocate professional and expert roles — courts appointing technical experts, clients instructing lawyers, hospitals assigning clinicians, regulators approving auditors — routinely ask a single question: is this person qualified to do the work? In many institutional settings they ask a second, logically independent question only inconsistently or not at all: how close is this person's demonstrated experience to the specific matter at hand? That second question is not unknown to the law — it is, in substance, the question Federal Rule of Evidence 702 and the Daubert/Kumho Tire line of authority require a court to ask of expert testimony, and economists since Spence and Stiglitz have long treated credentials as an imperfect signal of ability rather than ability itself. What is missing, particularly in civil-law jurisdictions and in the routine allocation of counsel rather than expert witnesses, is a structured, ex ante way of asking it. This paper names the conflation of the two questions the Qualification-Expertise Fallacy, and argues that the gap it produces leaves a structured set of detectable signals — observable in a practitioner's career record, in an organisation's own departmental description of itself, and in the technical content of the matter assigned — capable of being measured and acted upon before a failure occurs rather than only litigated after one has.

The paper develops three contributions, each narrower than a claim to conceptual originality and each, it argues, useful precisely because of that. First, it specifies semantic distance — the separation between a practitioner's demonstrated subject-matter experience and the technical content of a specific task — as an observable signal set, distinct from, and complementary to, the relational distance measured by network-proximity analysis, a separate and independently useful instrument addressed only briefly here. Second, it applies the standard law-and-economics model of the expected value of litigation, in which a disputed matter's value is discounted by the probability of a favourable outcome, to this signal set specifically: Realisable Resolution Value (RRV) frames competence distance as the variable principally responsible for that probability discount, giving it a structured and observable basis rather than leaving it to unstructured litigator judgment. Third, grounding both measures in settled doctrine rather than new theory, it connects the signal set to the existing law of agency, professional indemnity

insurance, and vicarious and negligent-supervision liability, to explain why a competence failure in representation can generate liability running in as many as three directions — to the represented party, to the opposing party, and back to the institution that made the staffing decision — and why responsibility for that failure belongs, in the great majority of cases, to the institution rather than the individual practitioner. The paper illustrates the framework with a case study from an ongoing Portuguese proceeding and the output of a prototype measurement instrument, presented as a worked proof-of-concept rather than as independent corroborating evidence, and closes with a proposal for how courts, regulators, firms, and insurers might act on these signals directly.

Section 1: Introduction

1.1 Two Questions, One Asked

Institutions that allocate professional and expert roles ask, almost without exception, a single question: is this person entitled to do this work. They ask it of court-appointed experts, of instructed counsel, of treating clinicians, of certifying auditors. The question is answered by reference to a credential — a degree, a licence, a bar admission, a board certification — and once answered, the matter of competence is, in routine practice, treated as closed.

A second question is almost never asked, despite being logically independent of the first: how close is this person's demonstrated experience to the specific technical content of the matter in front of them. A person can hold every credential the first question requires and still have a demonstrated career trajectory that has never approached the specific problem now assigned to them. This paper is about that second question — about making it answerable, about identifying the signals that answer it, and about what follows, doctrinally and practically, once it is asked.

1.2 A Sister Instrument, Not a Sequel

This paper has a companion: *Network Proximity and the Appearance of Impartiality*, which examined a different defect in the same general territory — the failure of civil law systems to test court-appointed experts for relational proximity to the parties, using Erdős-equivalent network distance borrowed from the mathematics of academic co-authorship. That paper stands on its own. It does not require this one, and this one does not require it. The two measure genuinely different things: network distance asks who a person is connected to; semantic distance, the subject of this paper, asks how close a person's demonstrated experience sits to a given technical problem. A person can fail one test and pass the other in either direction — relationally clear of every party while working far outside their actual competence, or deeply experienced in the subject matter while sitting one professional connection away from a party with an interest in the outcome. Each tool this author has built — the network-proximity instrument at kyc.co/network-proximity and the expertise-check instrument at kyc.co/expertise-check — answers one of these questions, and each does so independently of the other. Having said

that once, this paper does not return to the comparison. What follows is a standalone account of the semantic question, on its own terms.

1.3 Why the Gap Is Worth Treating as a Set of Signals

The central claim of this paper is not merely that qualification and expertise are different things — that distinction is intuitive enough that it is recognised informally in every profession examined here. The claim is that the gap between them leaves a trace: a set of observable, structured signals in a practitioner's career record, in the public-facing description an organisation gives of its own departments, and in the technical content of the specific matter at hand. Once those signals are named and the relationship between them is specified, the gap stops being a matter of impression and becomes a matter of measurement.

This matters because the consequence of the gap is not symmetric and not academic. Where a matter is mismatched to the person handling it, the dispute's own resolution value — what a correct, competently conducted resolution would actually deliver to the party whose rights are at stake — is placed at risk of erosion through delay, procedural error, or simply being argued badly. Where a matter is well matched, that same resolution value is protected. The competence gap is not a separate quantity to be added to or subtracted from the dispute. It is a multiplier on the probability that the dispute's existing value is realised at all. Sections 4 and 5 develop this point as Realisable Resolution Value, and Section 4 also sets out why, when the mismatched person is professional counsel rather than a private individual, the resulting exposure runs upstream — to the firm that made the staffing decision, and from there to the insurer that carries the firm's liability.

1.4 The Contribution, Stated Narrowly

This paper makes four contributions, each deliberately framed as an application or extension of existing work rather than a freestanding theoretical discovery. First, it specifies semantic distance as a measurable signal set — domain match, technical depth, experiential proximity, and capacity fit — derived from observable career and organisational record rather than from impression or reputation; comparable distance-based approaches to expertise matching exist in computer-science literature on

multi-agent task allocation, and this paper's contribution is applying that style of measurement to the legal and professional-liability setting rather than inventing the underlying technique. Second, it applies the standard expected-value-of-litigation model — value discounted by probability of a favourable outcome, a formula with a long history in law-and-economics scholarship — to competence distance specifically, naming the result Realisable Resolution Value and arguing that the signal set in Section 4.1 gives the probability term a structured basis it does not otherwise have. Third, it connects the signal set to the existing, settled law of agency, professional indemnity insurance, and vicarious and negligent-supervision liability, explaining why a competence failure in representation generates liability in up to three directions, and why that liability is correctly understood as institutional rather than personal to the individual practitioner in most cases where it arises. Fourth, it applies the resulting framework to a case study from an ongoing Portuguese proceeding, and reports the output of a prototype measurement instrument as a worked illustration of what an operational version of these signals could look like.

1.5 Structure

Section 2 sets out the Qualification-Expertise Fallacy as a general proposition and situates it against the existing literature and doctrine that already address parts of it. Section 3 develops the structure of the gap across professions and the institutional asymmetry in how it is scrutinised. Section 4 is the paper's doctrinal core: it specifies the semantic distance signal set, develops Realisable Resolution Value as an application of the expected-value-of-litigation model, connects the signal set to the existing doctrine of agency, insurance, and institutional and vicarious liability, and considers why the client is structurally poorly placed to catch the mismatch unaided. Section 5 presents a case study and a worked illustration from a prototype measurement instrument. Section 6 considers the institutional implications for courts, regulators, firms, and automated systems. Section 7 sets out a reform proposal, including consideration of the principal objections to it. Section 8 concludes.

Section 2: The Qualification–Expertise Fallacy

2.1 Two Different Questions, and What Is Already Known About Them

The argument of this paper rests on a single distinction, stated here in its most general form: *qualification determines whether a person may undertake a task; expertise distance determines how suitable they are for undertaking that particular task*. These are different questions, answerable independently, and an institution that answers only the first has not, by doing so, answered the second.

The distinction itself is not new, and this paper does not claim to be the first to draw it. Labour economics has treated credentials as an imperfect signal of underlying ability, rather than as ability itself, since Spence's foundational work on market signalling and Stiglitz's parallel work on screening, and a substantial sociology-of-credentialism literature running from Collins's *The Credential Society* onward has documented the gap between formal qualification and demonstrated capability across many occupations, not only the professions examined here. Evidence law has, in the common-law tradition, gone further than this paper's framing might initially suggest: under Federal Rule of Evidence 702 as construed in *Daubert v Merrell Dow Pharmaceuticals* and extended in *Kumho Tire Co v Carmichael*, an expert's general qualification in a discipline is, in principle, only the first of several admissibility questions, and courts applying this standard have excluded experienced practitioners whose specific specialty did not match the technical question at issue — a treating physician qualified generally in medicine but not in the specific sub-specialty the standard of care required, for instance. This paper's contribution is therefore not the discovery that qualification and expertise diverge; it is showing that the divergence can be specified as an observable signal set, applied systematically rather than discretionarily, and connected to a structured account of the financial and liability consequences that follow when it is ignored — particularly in civil-law jurisdictions and institutional settings where the Daubert-style second inquiry has no formal equivalent.

2.2 Qualification as a Binary Threshold

Qualification is, in every professional context examined in this paper, a threshold variable. A person either holds the degree or does not. Either is admitted to the bar, licensed by the regulator, board-certified, or chartered, or is not. Once the threshold is crossed, qualification does not increase. A lawyer admitted twenty years ago and a lawyer admitted last month hold, for the purposes of the formal credential, the same qualification: the right to practise law. The credential does not encode duration of practice, subject-matter concentration, case complexity handled, or outcome record. It encodes only that a defined minimum standard, set at the point of entry, was met at some point in the past.

This is precisely the property that makes qualification administratively convenient and substantively limited. It is a single bit of information — qualified or not — applied to a domain in which the actual variable of interest, demonstrated competence in a specific subject matter, is continuous and high-dimensional. It is worth confronting directly a feature of professional conduct rules that appears, on first reading, to cut against this paper's claim: the official commentary to the American Bar Association's Model Rule of Professional Conduct 1.1 states that a newly admitted lawyer can be as competent as a long-experienced practitioner for a matter outside their prior experience, provided adequate preparation, study, or association with a specialist colleague is undertaken. That proposition is not disputed here as a statement of what a careful practitioner *can* achieve. This paper's concern is narrower and, it argues, more consequential: whether an institution allocating the matter — the client, the firm's own staffing process, the court — has any structured way of knowing, before the fact, whether that preparation or association has actually occurred, as distinct from assuming it from the credential alone. The Model Rule commentary describes a standard a careful practitioner should meet; it does not supply, and does not purport to supply, a mechanism by which a third party can verify in advance that the standard has been met in a given case.

2.3 Expertise as a Continuous Distance

Expertise, by contrast, is not a threshold. It is a position in a space of subject-matter experience, and the relevant quantity is not whether a person has expertise in the abstract but how far their demonstrated experience sits from

the specific matter under consideration. This is naturally expressed as a distance measure rather than a binary classification: a distance of zero would represent a person whose accumulated experience maps exactly onto the matter at hand; a small distance represents close adjacency; a large distance represents a different domain entirely, however well qualified the person is in their own area.

The relevant question is not whether a professional has *some* relevant background — almost everyone has some plausible connection to almost any subject if the description is drawn broadly enough — but how many conceptual and practical steps separate their demonstrated experience from the technical content the specific matter actually requires. A general civil litigator and a construction-defects specialist are both, in the broadest sense, litigators; the distance that matters is the one measured at the level of the specific technical content in dispute, not at the level of the broad professional category both happen to share.

2.4 The Fallacy Stated

The Qualification-Expertise Fallacy can now be stated precisely: it is the inference from *qualified* to *expert*, treating a binary threshold variable as though it were informative about a continuous distance variable to which it is, in fact, only loosely and unreliably related. The fallacy is not the claim that qualification is irrelevant — a qualified person is, all else equal, more likely to be competent than an unqualified one — but the stronger and unwarranted claim that qualification status renders an inquiry into expertise distance unnecessary.

The fallacy has a precise structural signature. It appears wherever an institution treats the question “is this person qualified” as dispositive of the separate question “is this person suited to this specific matter,” and it appears with particular consequence wherever the institution in question is the one responsible for assessing or relying upon the professional's work — a court adopting an expert's conclusions, a client trusting counsel's technical judgment, a hospital relying on a clinician's diagnosis outside their demonstrated area, a regulator approving an auditor's sign-off.

2.5 Why the Shortcut Works Most of the Time

The fallacy is not always operative, and it is important to be precise about why. Within a narrow range of task difficulty — the range that the qualification threshold itself was calibrated to handle competently — qualification is a reasonably good proxy for expertise, because the variance in subject-matter difficulty across the category is small relative to the competence floor the qualification guarantees. A newly admitted lawyer handling an uncontested small debt claim, a general practitioner treating a routine infection, a recently chartered engineer inspecting a standard residential foundation: in each case the qualification threshold was set with exactly this class of matter in mind, and the shortcut performs well.

This is the condition under which institutions are entitled to rely on qualification alone, and it explains why the shortcut persists: it is not merely administratively convenient, it is substantively adequate across the majority of matters any qualification category actually encounters.

2.6 Where the Shortcut Fails

The shortcut fails as the technical complexity and specificity of the matter increases, because the distance between the easiest and the hardest matter a qualification category might encounter grows, while the qualification threshold — set, of necessity, at the level the least experienced entrant can meet — remains fixed. A complex structural defect dispute, a contaminated-aquifer regulatory dispute, a multi-party planning and land-registry conflict: each sits at the far end of the difficulty distribution within its profession, precisely where the qualification threshold provides the least information about whether a given practitioner is equipped to handle it.

This produces an asymmetry worth stating explicitly: the matters in which the Qualification-Expertise Fallacy is most consequential are also, systematically, the matters with the highest stakes — the largest financial exposure, the most serious procedural consequences, the greatest risk of irreversible harm. The fallacy is not evenly distributed across risk. It concentrates precisely where its cost is highest.

Section 3: The Structure of the Gap

3.1 Qualified, Not Expert

The first quadrant of the framework is the one institutions are least equipped to detect, because it presents no formal irregularity. The professional is properly admitted, properly licensed, in good standing, and a member of a reputable institution. Nothing in the credential file signals a gap. The gap exists, if it exists, in the distance between what the credential certifies — competence to practise the general discipline — and what the specific matter requires.

A useful operational marker of this quadrant is the practitioner's own demonstrated career trajectory measured against the technical content of the assigned matter, independent of seniority. A practitioner can have eight, fifteen, or twenty-five years of qualified standing and still have a demonstrated trajectory that never approaches the specific technical domain at issue, because seniority within a profession is not the same variable as subject-matter concentration within that profession. Years of practice is a proxy for accumulated general competence; it is not, by itself, a proxy for proximity to any particular specialism the practitioner has not actually practised.

3.2 Expert, Not Qualified

The second quadrant is the mirror image and, for institutional purposes, the more visible one, because the absence of formal qualification is itself a recorded fact that someone can check. A construction foreman who has personally observed soil failure across thirty-five years of foundation work, a long-serving regulatory inspector who has reviewed hundreds of compliance files in a narrow technical area, a self-taught network analyst who has built and tested manipulation-detection methodologies over a sustained period: each may have a closer semantic proximity to a specific technical question than a credentialed professional encountering the same question for the first time, while holding no formal credential that would entitle an institution to rely on their conclusion as a matter of right.

This quadrant illustrates why qualification and expertise cannot be collapsed into a single ordinal scale running from “less qualified” to “more qualified.” A person in this quadrant is not partially qualified. They are unqualified and may simultaneously be the single most semantically proximate person available to address the specific question. Institutions that require the credential as a precondition for being heard at all — rather than as one input among several into a competence assessment — discard exactly the information this quadrant contains.

3.3 Domain Examples

The pattern recurs with sufficient consistency across professions that it supports the general claim made in Section 2 rather than a narrower, domain-specific one.

Medicine. A general practitioner is qualified to practise medicine generally. A cardiac surgeon is qualified in the general discipline and has, in addition, a depth of procedure-specific experience that the qualification alone does not confer and that no patient would expect a general practitioner to substitute for in a cardiac matter. The qualification difference between the two practitioners is formally describable; the experience difference is what actually determines outcomes in a cardiac case.

Engineering. A civil engineer holds a qualification covering structures generally. A geotechnical engineer has the specific accumulated experience that a foundation-failure or soil-stability question requires. Both are qualified engineers. The semantic distance between a civil engineer with no geotechnical case history and a geotechnical specialist, when the question on the table is soil behaviour beneath a structure, is large, notwithstanding that both credentials sit in the same general professional category.

Law. A lawyer is qualified to practise law within the scope of their admission. A specialist property litigator, a construction-disputes practitioner, or regulatory counsel has accumulated case-specific experience in technically dense areas that general civil litigation training does not, by itself, confer. The bar admission is identical. The semantic distance from a complex structural-defect and land-registry dispute is not.

Auditing and regulatory compliance. A chartered accountant or auditor is qualified to certify financial statements generally. A specialist in a particular regulated sector — financial derivatives, insurance reserving, environmental liability accounting — has sector-specific technical depth that the general qualification does not encode. Regulatory failures attributable to an auditor signing off on a matter outside their demonstrated sector experience, while fully qualified in the general sense, recur across financial-crisis and corporate-failure literature as a documented pattern rather than a hypothetical one.

3.4 The Asymmetry of Institutional Scrutiny

Across each of these domains, institutions apply asymmetric scrutiny to the two questions. The qualification question is checked rigorously, often as a formal precondition: a court will not appoint an unlicensed expert; a client cannot instruct someone who is not admitted to practise; a hospital will not roster an uncertified clinician for a procedure requiring certification. The expertise-distance question, by contrast, is frequently not checked at all, or is checked only informally, through a CV review or a firm's own departmental assignment process, with no external verification and no consequence attached to getting the assignment wrong.

This asymmetry is not accidental. The qualification question has a clear institutional owner — the regulator, the bar, the licensing body — whose function is precisely to police that threshold. The expertise-distance question has, in most of the settings examined in this paper, no equivalent institutional owner. No regulator audits whether the specific lawyer assigned to a specific file has demonstrated experience proximate to that file's technical content; no Portuguese court-appointed expert roster requires a nominee to demonstrate subject-specific proximity beyond the general qualification that earned them a place on it. This is not true everywhere: the asymmetry this paper describes is considerably narrower in jurisdictions applying something like the Daubert/Kumho Tire standard to expert evidence, where a court is required, at least in principle, to ask the second question before admitting expert testimony at all, and has on occasion excluded an experienced practitioner whose specific specialty did not match the question asked of them. The institutional implications and reform proposal in Sections 6 and 7 accordingly

treat that line of doctrine not as a rival claim to be distinguished away, but as a working example of what asking the second question in a structured, ex ante way can look like — one this paper argues should be extended, in substance, to the civil-law expert-roster and counsel-allocation contexts its case study is drawn from, where no comparable gatekeeping currently exists. The result, in those contexts, is a structural blind spot: the variable most predictive of outcome quality in technically specialised matters is the variable institutions are least equipped, and least required, to examine.

Section 4: The Signal Set, Realisable Resolution Value, and Institutional Liability

4.1 From Distinction to Signal

Sections 2 and 3 established that qualification and expertise are different things, and that the gap between them concentrates precisely in the technically specialised, high-stakes matters where its cost is highest. That distinction, on its own, is of limited practical use. What makes it actionable is that the gap leaves a trace: a set of observable signals, present in a practitioner's own career record, in an organisation's own description of its departments and specialisms, and in the technical content of the matter itself. None of these signals require speculation about a person's competence in the abstract. Each is drawn from material the practitioner, the firm, or the matter itself already discloses.

Four signals are sufficient to specify the gap with useful precision. *Domain match* asks whether the practitioner's demonstrated field of practice overlaps with the matter's field at all. *Technical depth* asks, conditional on domain overlap existing, how far the practitioner's demonstrated experience extends into the specific technical sub-content the matter requires — not litigation generally, but this kind of litigation; not engineering generally, but this kind of failure. *Experiential proximity* asks whether the practitioner has, on the public record, actually handled a matter of this type before, as distinct from being theoretically qualified to do so. *Capacity and resource fit* asks whether the scale and complexity of the matter is proportionate to what the practitioner and the organisation behind them can be shown to have handled previously. Each signal is independently observable; together they support a structured comparison rather than an impressionistic one.

A measure of this kind — distance between a demonstrated profile and a task's requirements, used to flag a competence gap before relying on the match — is not without precedent as a technique, although its application here is new. Comparable approaches appear in computer-science work on multi-agent task allocation, where a task is assigned to the agent whose demonstrated profile sits closest to the task's requirements across several weighted dimensions, including an explicit semantic-mismatch term, and in expertise-location systems that compute the distance between an individual's

activity profile and an aggregate expert profile on a given topic. This paper does not claim to have invented distance-based expertise measurement; it claims to be applying that general technique, for the first time so far as this paper is aware, to the specific problem of legal and professional-liability staffing, where the consequence of an undetected mismatch is not a suboptimal task routing but a negligence exposure running, as Section 4.5 develops, in several directions at once.

4.2 Classifying the Gap

These four signals, taken together, support a classification of the relationship between a practitioner and a matter, rather than a single undifferentiated score. A practitioner whose demonstrated experience *subsumes* the matter — more senior or more specialised than the task requires — is competent despite any apparent distance on a crude reading of their credentials. A practitioner *within domain* is a direct and appropriate match. A practitioner in an *adjacent* domain is related but distinct, and the matter will require knowledge transfer the practitioner has not yet demonstrated. A practitioner whose depth is *insufficient* is in the right general domain but the wrong depth or seniority for this complexity. A practitioner who is simply *mismatched* is in a different field entirely, with no relevant conceptual framework for the work at hand. This classification does the work that a single qualification bit cannot: it locates a specific practitioner within a graded space, rather than sorting them into “qualified” and stopping there.

The prototype instrument illustrated in Section 5.4 operationalises this classification with a 0–100 distance score banded against the five categories above, reproduced below as the instrument's own published scoring guide rather than as this paper's independent proposal, since a worked reference of this kind is more useful shown than re-described in prose.

SCORE REFERENCE GUIDE			
SDD Score	Risk Indication	Subsumption	Meaning
0–25	Excellent match or overqualified — Low	SUBSUMES	Senior person, simpler task. Competent despite apparent distance.
26–50	Acceptable match with minor gaps — Moderate	WITHIN_DOMAIN	Direct match at appropriate depth. Optimal alignment.
51–70	Significant gaps — Elevated	ADJACENT	Related but distinct domain. Knowledge transfer required.
71–85	High mismatch — High risk of error	INSUFFICIENT	Right domain, wrong depth or seniority for this complexity.
86–100	Critical mismatch — No basis for this work	MISMATCHED	Different field entirely. No conceptual framework for this work.

Figure 1. Score reference guide published by the KYC.co Expertise Check prototype, mapping SDD score bands to risk indication and subsumption classification.

4.3 What the Signals Are Not

Two clarifications matter before the signal set is put to use. First, the signals measure distance from a body of demonstrated experience, not distance from a credential, and not a person's intelligence, character, or general competence. A highly capable practitioner can be genuinely mismatched to a specific technical matter through no fault of their own — the mismatch is a fact about the assignment, not a judgment about the person. Second, and more important, the signal set is built to detect a specific failure mode — a gap between demonstrated experience and task content that creates a risk of unintentional error — and it has no purchase on a different and more troubling failure mode: deliberate misconduct by a highly experienced actor who understands the relevant technical and legal terrain precisely, and exploits that understanding in bad faith. A repeat operator who knows exactly how far a regulatory half-truth can be stretched, what a technical report can omit without being caught, and which relationships make a finding against them unlikely, is not semantically distant from the work. He may be the most experienced person in the room, and the problem is not a competence gap at all but an intentional one. That failure mode is real, consequential, and addressed elsewhere; it is not what this signal set is designed to find, and applying this framework to it would mischaracterise the defect.

4.4 Realisable Resolution Value: Applying the Expected-Value Model to Competence Distance

A disputed matter is not, in the ordinary case, a contest with an undetermined value. Where one party has suffered a breach, a structural harm, or an unjust enrichment at another's expense, the dispute has an underlying resolution value: what making the wronged party whole actually requires. Courts do not usually describe this in the language of upside or profit, and this paper does not either; the value in question is restorative, not speculative. It exists independently of who represents either party, and it does not increase or decrease according to the skill of the practitioners involved. What the skill of the practitioners involved determines is something different: the probability that this existing value is actually realised, rather than eroded through delay, procedural error, a mischaracterised technical position, or an argument competently made for the wrong proposition.

This is not a new relationship to propose. Discounting a claim's value by the probability of a favourable outcome is the standard model of the expected value of litigation in law-and-economics scholarship, developed in the settlement and litigation-selection literature associated with Priest and Klein and applied routinely, if often informally, by litigators pricing settlement positions, who already treat the skill of opposing counsel as one of the qualitative factors feeding into their estimate of that probability. What this paper adds is narrower and more specific: rather than leaving the probability term to unstructured professional intuition, it proposes that the term be partly grounded in the observable signal set developed in Section 4.1, so that the contribution of competence distance to that probability can be inspected and discussed rather than absorbed silently into a litigator's overall feel for the case.

Stated formally: let V be the resolution value of the matter — what a correct and competently conducted resolution would deliver to the party whose rights are at stake. Let p be the probability that this value is realised rather than eroded, where p is, on this paper's proposal, partly a function of the competence distance established by the signal set in Section 4.1: p is higher, all else equal, where the practitioner and the organisation behind them are closely matched to the matter's technical content, and lower as that distance increases. Realisable Resolution Value is then $RRV = V \times p$, the same

expected-value structure used throughout the litigation-economics literature, with competence distance proposed as one structured, observable input into the probability term rather than the only one. This is not a separate quantity layered on top of the dispute, and it is not a prediction of how much a party will win. It is an estimate of how much of the value the dispute already contains is likely to survive the manner in which it is handled, holding the underlying merits constant.

The formula's value, on this narrower framing, is in what it makes visible rather than in any claim to numerical precision or theoretical novelty. It makes clear that competence distance and the underlying merits of a dispute are different variables, multiplying against each other rather than substituting for each other: a strong claim handled by a badly matched practitioner does not become a weak claim, but a meaningful share of its value is placed at risk; a modest claim handled by a closely matched specialist does not become a large one, but what value it has is more likely to be realised in full. It also clarifies why the financial exposure framing developed elsewhere in this paper's companion instrument — the FEE module illustrated in Section 5.4 — and the Realisable Resolution Value framing developed here are two views of the same underlying mechanism rather than two competing claims about gain and loss: FEE estimates the cost that materialises when p is low and the erosion occurs; RRV frames the same relationship prospectively, as the share of V a party can expect to keep depending on who is handling the matter.

4.5 Counsel as Insured Fiduciary Proxy: An Application of Settled Agency and Insurance Doctrine

Nothing in this section is a new legal proposition. The law of agency has long recognised that, for the duration of a matter, counsel stands in the shoes of the party they represent, advancing that party's claims through counsel's own filings, technical characterisations, and procedural choices, and this substitution is precisely why counsel is required to carry professional indemnity insurance: the substitution of counsel's judgment for the party's own direct action is understood, institutionally, to carry a risk that must be backed by a capable indemnitor, and the insurance requirement exists because the risk is real rather than theoretical. What this paper adds is not

the doctrine itself but its explicit connection to the signal set in Section 4.1: the same agency and insurance principles that have always applied to representation give a concrete, observable trigger — a measurable competence mismatch — for exactly the kind of failure the doctrine already anticipates.

Once this substitution is taken seriously, the consequence of a competence failure in representation is not confined to a single relationship. It can run in as many as three directions. First, and most obviously, to the represented party themselves: a client whose interests are damaged by an error, omission, or mischaracterisation that a closely matched practitioner would not have made has a potential claim against their own counsel, independent of the outcome of the underlying matter. Second, to the opposing party: where counsel's error causes unnecessary delay, forces costs that would not otherwise have been incurred, or creates a lost opportunity — a financing window missed, an insurable interest that lapses, a construction season lost — the opposing party has arguably suffered harm caused not by the party they are in dispute with, but by that party's representative acting outside their competence. This is a distinct harm from the underlying dispute itself, and it runs to counsel directly rather than through the represented party. Third, back to the underlying dispute: a competence failure can distort the resolution of the original matter in a way that affects how much of V , as defined in Section 4.4, either party actually recovers, independent of either harm described above.




Figure 2. The three directions of liability exposure that follow from a competence failure in representation, set out in this section.

These three directions of exposure are not a novel proposal. They are the ordinary structure of professional liability law in every jurisdiction examined for this paper, and they explain why the indemnity insurance requirement exists in the first place: a profession in which representation substitutes for direct action, and in which that substitution can cause harm to one's own client, to the opposing party, and to the resolution of the dispute itself, is a profession in which the insurer's exposure is real and must be priced. The signal set in Section 4.1, and the RRV framing in Section 4.4, are offered as a way of pricing that already-recognised exposure before the fact, using observable signals, rather than only reconstructing it after a claim has been filed.

4.6 Institutional and Vicarious Responsibility

A competence failure of the kind this paper describes is, in the great majority of cases, not best understood as a personal failing of the individual practitioner. It is better understood as a staffing decision made by the organisation that assigned them. This is, again, not a new proposition: vicarious liability under the doctrine of respondeat superior, and the closely related tort of negligent hiring and supervision, have long held an employer responsible for harm caused by an employee acting within the scope of their assigned duties, including harm that traces to the employer's own decision about who to assign to a given task. The distinction this paper relies on — between the individual's conduct and the institution's staffing decision — is therefore already embedded in the law. What this paper proposes is a way of identifying, before a claim is filed, when that staffing decision has in fact created the relevant risk, using the signal set in Section 4.1 rather than waiting for a court to reconstruct the mismatch after the fact.

Two analogies outside the legal profession make the underlying doctrine plain, and both are drawn from areas where the negligent-assignment fact pattern is already well established in malpractice and tort literature rather than novel to this paper. A hospital that asks a general practitioner to perform a cardiac surgical procedure has not produced a negligent general practitioner; it has made a negligent staffing decision, and the institution —

not the individual clinician — is the appropriate object of scrutiny and liability. A school that asks a student to drive the school bus has not produced a negligent student driver; it has made a negligent decision at the level of the institution responsible for student safety, and no one would direct the resulting liability at the student. The same structure holds where a firm with an advertised real-estate and urban-planning department assigns a construction-defect and land-registry matter to a general civil litigator with no demonstrated experience in either area. The associate has not behaved negligently in the ordinary sense of the word; she has been placed, by a decision made above her, into a position for which the signal set in Section 4.1 would show her to be poorly matched. The negligence, if it exists, is upstream: it belongs to whoever made the staffing decision, and from there to the firm's own management and, ultimately, to the professional indemnity insurer that underwrites the firm's exposure for exactly this kind of decision. This has a direct and practical implication for where the institutional implications developed in Section 6 should be aimed. A reform proposal, a regulatory intervention, or a settlement strategy addressed only to the individual practitioner mistakes the location of the failure under doctrine that already exists. The correct addressee is the institution that made the assignment, and, where the institution's own exposure is itself backed by insurance, the insurer that will ultimately bear the cost of having underwritten that exposure without ever having tested whether the signal set in Section 4.1 would have flagged the mismatch in advance.

4.7 Why the Client Rarely Asks: Credence Goods, Deference, and Firm Economics

Sections 4.5 and 4.6 establish where responsibility lies once a mismatch of the kind this paper describes has occurred. A separate question, prior to that one, is why the mismatch so often goes unchallenged by the one person with the most at stake in catching it: the client. The answer is not that clients are careless. It is that the structure of the transaction gives them very little to work with, for reasons that are well understood in economics under a different name.

Legal services are, in the economic literature on professional services, a textbook instance of what Darby and Karni termed a *credence good*: a good or

service whose quality the buyer cannot reliably assess even after it has been consumed, because the buyer lacks the expertise to evaluate the procedure, only the outcome, and the outcome is itself confounded by the underlying merits of the matter. A client who instructs a firm for a complex construction and land-registry dispute is, in this sense, in the position of a patient who cannot tell whether a surgical outcome reflects the surgeon's skill or the severity of the underlying condition: legal services, alongside medical and mechanical repair services, are among the examples the credence-goods literature itself uses, precisely because the client commissioning the work has no comparable transaction to benchmark it against and frequently no way of knowing, win or lose, whether the work was done well. A client does not get to see how the matter would have gone with a different practitioner assigned. In many cases the client never learns, even after the matter concludes, whether a different allocation would have changed the outcome — the comparison case simply does not exist to be checked against.

This informational gap interacts with an ordinary social norm to make the mismatch even harder to challenge in practice. A client who has sought out a firm specifically because the matter is serious enough to warrant the firm's reputation has, by that choice, already signalled to themselves that the stakes are high and the technical terrain is unfamiliar. Querying the seniority or specialism of the specific individual assigned, absent some overt signal that something has gone wrong, runs against ordinary professional courtesy and against the client's own incentive to maintain a cooperative working relationship with people on whom their matter now depends. Common decency, in the ordinary case, counsels giving the assigned practitioner the benefit of the doubt rather than opening the relationship with a direct challenge to their seniority — and a client is, in any event, poorly placed to make that challenge stick, since the firm's brand is the very thing the client is relying on to answer a competence question they cannot otherwise answer themselves. The result is that the one moment at which the mismatch could most easily be caught — before any work has been done, when the assignment is first made — is also the moment at which social convention and informational asymmetry most strongly discourage the client from raising it.

A further, more mundane factor compounds this. Large-firm billing economics are built around what the profession itself calls leverage: the

practice of billing client matters at senior, premium rates while having a substantial share of the underlying work performed by associates compensated at a small fraction of what the client is charged for their time. This structure is well documented in legal-profession economics literature as a primary driver of large-firm profitability, independent of any wrongdoing, and it creates a structural incentive, again independent of any individual's intent, to route work toward whichever timekeeper is available rather than whichever timekeeper is the closest match on the signal set in Section 4.1. None of this requires bad faith on anyone's part to operate exactly as described; it requires only the ordinary economics of how large firms staff their matters, combined with a client who has no practical way of seeing the staffing decision from the inside.

It is worth being precise about what this section does, and does not, claim. The pattern described here is structural and, on the evidence available, consistent with ordinary firm economics rather than deliberate exploitation: a mismatch can arise from leverage incentives and informational asymmetry alone, without anyone involved intending the client harm Section 4.4 describes. This is a different and considerably less troubling category than the deliberate exploitation of an information or power asymmetry by an actor who fully understands the gap and profits from maintaining it — a category this paper does not address, but which is taken up directly in this author's companion paper on the manufacturing of bad faith, where the asymmetry is not a by-product of ordinary business economics but the object of the strategy itself. The two categories can resemble each other from the client's chair, since in both the client receives less than the relationship implied; they are nonetheless distinct, and a firm correctly identified under the structural account in this section has a materially different position, and a materially better path to resolution, than one operating under the conduct the companion paper describes.

Section 5: The Case Study

5.1 The Setting

The case study concerns adjoining development plots in a Portuguese jurisdiction. The claimant holds an approved architectural project for residential development on two lots and has been unable to obtain construction insurance without certified structural documentation concerning the adjoining property, whose footprint and licensing position changed under a 2021 addendum to the area's loteamento alvará. The matter, on its technical content, requires command of Portuguese construction and urban-planning law, the regulatory architecture governing illegal or unlicensed structural modification, coordination with forensic engineering evidence, Land Registry (Registo Predial) injunctive procedure, and the specific procedural rules of the Julgado de Paz — a specialised small-claims peace court operating under its own statute, with jurisdictional limits distinct from the ordinary civil courts.

There is no relationship of any kind alleged between the assigned legal team and the claimant. Nothing in this case study turns on a connection, a bias, or an undisclosed interest. It is, on its facts, a clean test of the signal set developed in Section 4: a matter with a defined technical content, a represented party, and a publicly verifiable record of the representing practitioner's actual experience.

5.2 The Defence Advanced and the Allocation Decision

The respondent's legal representation was provided by a mid-to-large Iberian firm with departments covering, among other areas, real estate and urban planning law and public law, urban planning, environment and energy law — departments whose advertised specialism maps directly onto the matter's technical content. The firm's real estate and construction practice is, independently of its own marketing materials, externally recognised through independent third-party legal-directory research as a ranked institutional specialism — confirming, from outside the firm's own description of itself, that the specific specialism this matter required was not merely advertised but externally recognised as a genuine institutional strength. The matter was, however, conducted by an associate whose own background, on the public

record, lies in general civil litigation and family and succession law, with earlier career experience in securities regulation and consumer protection — fields that, whatever their own merits, do not overlap with construction law, urban planning, or land-registry procedure.

The defence filed in the proceedings characterised the adjoining property as fully licensed and compliant. That characterisation was contradicted by four independent official sources, none commissioned by the claimant: municipal planning records, a judicial inspection record, and the documented expiry of the relevant project licence. This paper takes no position on the merits of that dispute, which remain for the proceedings themselves to resolve, and nothing in this section should be read as doing so. What is of interest here is narrower and entirely independent of who is ultimately right: the structural fact pattern by which a position on a specialised technical question came to be advanced by counsel whose demonstrated career experience does not extend into that technical area, within a firm whose own institutional departments, had they been engaged, plainly would.

5.3 The Qualification, Stated Plainly

The associate is fully qualified. She is admitted to practise, in good standing, employed by a reputable firm, and operating squarely within the scope of her professional licence. Nothing about her qualification is irregular, and nothing in this paper suggests otherwise. The question this paper raises is not whether she was entitled to appear — she plainly was — but whether her qualification, considered alone, provided the firm, the client, or the court with any meaningful information about her semantic proximity to the specific technical content of this matter. On the analysis developed in Sections 2 and 3, it did not. Bar admission certifies general fitness to practise law. It does not, and is not designed to, certify subject-specific depth in construction defect litigation, urbanismo regulation, or Julgado de Paz procedure.

5.4 A Worked Illustration — Semantic Delegation Distance

To illustrate, rather than to assert, what a computational approach to measuring the signal set in Section 4.1 might look like, a prototype instrument — a semantic delegation distance (“SDD”) tool, developed as a proof-of-concept companion to this paper's framework — was run against the

public-record description of the matter, the assigned associate's professional background, and the instructing firm's own description of its practice areas. The exercise is reported here as an illustration of the kind of structured output the signal set contemplates, not as independent corroborating evidence of the underlying facts: the tool is a prototype, its output is generated by an automated language-model process rather than human professional assessment, it remains a work in progress, and it has not been independently audited. Nothing in this subsection should be read as a factual finding by this paper, by the instrument, or by any court; the figures below should be read in that light throughout.

The instrument now presents its output in the form of a case caption, naming the assigned practitioner and the score in the position a court filing would give to the parties and the outcome — a framing choice of the tool's own design, reproduced here exactly as generated.




Figure 3. Prototype output header, reproduced verbatim, framed by the instrument itself in the manner of a case caption.

The current version of the instrument also introduces a three-way comparison, scoring not only the assigned practitioner and the instructing firm but a named available alternative within the same firm — used here strictly as a same-firm benchmark illustrating that a closer match existed internally, not as any claim about who made the staffing decision, which this paper does not assert and the instrument does not claim to know.

Instructing Firm	28	Moderate Risk
Assigned Practitioner	58	Elevated Risk
Alternative Specialist	32	Moderate Risk
Assignment Delta (Assigned – Alternative):		+26 pts

Figure 4. Three-way competence comparison — instructing firm, assigned practitioner, and a same-firm alternative, with the resulting assignment delta.




Figure 5. Assignment delta — the SDD-point gap between the assigned practitioner and the same-firm alternative, presented by the instrument as the preventable portion of the competence gap.

The instrument returned an overall semantic delegation distance score of 58 out of 100, adjusted upward from a base score by a case-value multiplier reflecting the matter's financial scale, and classified the relationship between the assigned associate and the matter as *insufficient* — in the terminology introduced in Section 4.2, the right general domain but the wrong depth or seniority for this complexity, one band more severe than a simple domain mismatch. A separate allocation friction score of 68 out of 100 was returned for the gap between the firm's own advertised institutional competence and the associate actually deployed on the file, anticipating the institutional point developed in Section 4.6: that an organisation can possess the relevant departmental expertise in the abstract while the allocation decision that assigns a specific person to a specific file fails to draw on it.

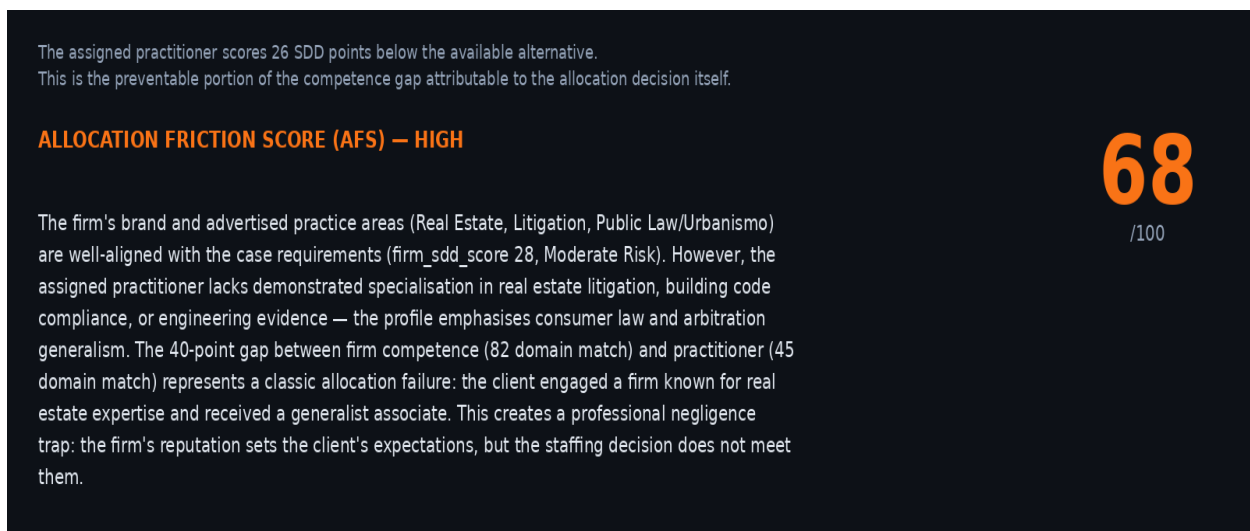


Figure 6. Allocation Friction Score narrative, as generated by the prototype.

The instrument's own summary of its output, reproduced verbatim, is a useful entry point before the remaining figures: “A mid-level generalist litigation associate with no demonstrated real estate or engineering expertise has been assigned to a €480,000 complex structural engineering dispute requiring integrated knowledge of Portuguese building codes, forensic evidence, and real estate law — despite the firm having a senior partner with proven real estate and civil procedure credentials available; this represents a material allocation failure that places €230,000 of claim value at risk.” That sentence, including the case value and at-risk figures it states, is the prototype's own AI-generated output, produced from the inputs described above. It is reproduced here as exactly that — a worked illustration of what the instrument generates, consistent with the qualitative analysis in Section 5.2 — and not an independent factual finding of this paper, of the instrument's developer, or of any court. The €230,000 at-risk figure corresponds to the erosion side of the Realisable Resolution Value relationship set out in Section 4.4: one illustration of what can be lost from V when p is low, not a verified damages calculation, and no reliance is placed on the specific figure itself.

PLAIN ENGLISH VERDICT

A mid-level generalist litigation associate with no demonstrated real estate or engineering expertise has been assigned to a €480,000 complex structural engineering dispute requiring integrated knowledge of Portuguese building codes, forensic evidence, and real estate law — despite the firm having a senior partner with proven real estate and civil procedure credentials available; this represents a material allocation failure that places €230,000 of claim value at risk.

Figure 7. Plain English Verdict, reproduced verbatim from the prototype's output.




Figure 8. SDD metrics dashboard — headline scores, subsumption classification, and negligence risk indicator.




Figure 9. Competence mapping grid — case complexity plotted against practitioner seniority, against the required level for this matter.

The instrument's dimensional breakdown is consistent with the qualitative account given above, and now scores the same-firm alternative alongside the assigned practitioner and the firm. Domain match was scored at 45 percent for the practitioner against 82 percent for both the firm and the alternative; technical depth at 35 percent against 75 and 78 percent; experiential proximity at 40 percent against 80 and 85 percent; capacity alignment at 55 percent against 85 and 88 percent; jurisdictional alignment, where general Portuguese admission is sufficient regardless of specialism, was scored closely across all three at 90, 95, and 95 percent. The pattern across every dimension except the purely jurisdictional one is the same: both the firm's institutional position and the named alternative sit consistently and

substantially closer to the matter's requirements than the assigned practitioner's demonstrated position — the same gap Section 4.6 identifies as belonging, in the first instance, to the institution rather than to the individual.




Figure 10. Competency radar profiles — assigned practitioner (cyan), instructing firm (purple), and same-firm alternative (green), across the dimensions in Section 4.1 plus a jurisdictional-alignment axis specific to this matter.

DIMENSION	ASSIGNED	FIRM	ALTERNATIVE	FRICTION GAP
Domain Match	45%	82%	82%	37%
Jurisdictional Alignment	90%	95%	95%	5%
Technical Depth	35%	75%	78%	40%
Experiential Proximity	40%	80%	85%	40%
Capacity Alignment	55%	85%	88%	30%

Figure 11. Dimensional alignment table — assigned, firm, and alternative scores against the resulting friction gap on each axis.

5.5 Realisable Resolution Value and Institutional Responsibility, Applied

The matter is, in its basic shape, a dispute over harm: an adjoining owner whose structural and licensing position has, on four independent official sources, been other than represented, leaving the claimant unable to obtain construction insurance or proceed with an approved development. On the framework set out in Section 4.4, this matter has a resolution value V — what correctly establishing the adjoining property's actual structural and licensing position would be worth to the claimant in restoring the ability to insure,

finance, and build. That value originates with the underlying harm, and it exists independently of who represents either side.

What the signal set and the worked illustration in Section 5.4 add to that picture is the probability term, p . On the analysis in Sections 4.3 and 4.6, a defence conducted by a practitioner with a measured semantic distance in the adjacent-to-insufficient range, within a firm whose own departmental expertise was not engaged on this file, carries an elevated risk of compounding the underlying harm rather than resolving it cleanly — not because the associate has behaved badly, but because the mismatch identified in Section 5.4 was, on this paper's analysis, a staffing decision made above her. This is a statement of elevated risk, consistent with the RRV framework's treatment of p as a probability rather than a certainty, and not a prediction of how the proceedings will in fact conclude. Each procedural step taken on a mischaracterised technical footing, each filing that does not engage the specific regulatory and forensic content the matter actually requires, is, on that same analysis, a further reduction in p : a further share of V placed at risk of being lost not to the merits of the dispute, but to the manner of its handling.

Section 4.6's analogies apply directly here. The harm in this matter did not originate with the associate, and responsibility for the defence's apparent mismatch to the matter's technical content does not terminate with her either. It runs, as the hospital and school-bus analogies in Section 4.6 illustrate, to whoever within the firm made the decision to assign a general civil litigator to a construction-defect and land-registry matter the firm's own marketing materials say it has specialist departments to handle — and from there, under the doctrine set out in Section 4.5, to the firm's leadership and to the professional indemnity insurer that underwrites the firm's exposure for exactly this category of decision. If the mismatch identified in Section 5.4 in fact contributes to an erosion of V , the resulting exposure is, on the analysis developed in this paper, properly addressed to that institutional and insurance layer, and not solely to the individual practitioner who was assigned the file.

Section 6: Institutional Implications

6.1 Courts

Courts that appoint experts from a professional roster typically verify qualification — registration, professional standing, absence of disqualifying conflicts — without separately verifying semantic proximity to the specific matter assigned. A roster organised by general discipline (engineering, medicine, accountancy) rather than by sub-specialism and demonstrated case history reproduces the qualification-as-proxy structure at the point of appointment itself. A court that adopts an appointed expert's conclusions in the substantial majority of cases is relying on a process that screens for qualification and conflicts but not, as a matter of routine, for semantic fit as defined in Section 4.1. The reform implication is straightforward in principle, though it requires institutional infrastructure that does not currently exist in most civil law jurisdictions: rosters should record sub-specialism and demonstrated subject-matter history with sufficient granularity that an appointing judge can assess semantic proximity, not merely qualification, before an appointment is made rather than after a challenge is filed.

6.2 Regulators and Professional Bodies

Professional bodies are, definitionally, the institutional owners of the qualification threshold, and have correspondingly little institutional incentive to develop or publicise a measure that would reveal the threshold's limits. This is not an accusation of bad faith; it is a structural observation about institutional incentives. A bar association, medical college, or engineering institute exists to certify and defend the qualification it confers; a measure showing that qualification alone is a weak predictor of fitness for specific high-stakes matters does not advance that institutional purpose, even where it would advance the purpose of protecting clients and patients. Regulatory reform in this area is therefore unlikely to originate from the bodies that hold the qualification function, and is more plausibly driven by client-side, insurer-side, or judicial demand for a complementary measure — which is the institutional space this paper's framework, and the SDD instrument illustrated in Section 5.4, is designed to occupy.

6.3 Firms and the Competence-by-Proxy Problem

The case study illustrates a distinct institutional failure mode worth naming separately from the fallacy itself: *competence-by-proxy*, the pattern by which an organisation's aggregate institutional reputation and departmental breadth is implicitly extended, by the client or the court, to whichever individual practitioner within that organisation happens to be assigned a given file. A client instructing a firm with an advertised real-estate and urban-planning department reasonably expects that expertise to be brought to bear on a real-estate and urban-planning matter; if the firm's internal staffing process assigns the file instead to an available generalist, the client receives the firm's reputation without receiving the substance the reputation implies, and frequently has no visibility into the substitution at all. Legal-practice-management research already documents the underlying mechanism, from the client's side, as a feature of how legal services are bought: clients commonly lack the specialist knowledge to judge a lawyer's skill directly, and rely instead on firm reputation, brand, and referral as a substitute. That literature treats the resulting reliance as a business-development and client-relations question; Section 4.7 develops the complementary point that the client's reliance on the firm's brand is not mere inattention but a structurally reasonable response to a transaction in which the client cannot verify quality directly, made harder still by the ordinary social costs of querying a specialist firm's own staffing choices. This paper's contribution is to treat the same reliance as a liability-detection question: the gap between firm-level reputation and individual-level competence is, on the analysis in Section 4.1, an observable signal, and Section 4.6 already supplies the doctrinal answer to where responsibility for it belongs — with the institution that made the staffing decision, not with the individual practitioner who was assigned the file. The prototype instrument's own output on the case file, reported in full in Section 5.4, characterised this same pattern independently, describing it as a case in which “institutional competence does not guarantee individual assignment quality, and brand reputation can mask staffing decisions that violate the principle of proportionate expertise matching — creating a hidden liability vector for professional indemnity insurers and disciplinary bodies” — a description arrived at from the prototype's own analysis of the firm's and practitioner's public materials, not supplied to it as an instruction, and

offered here as a second, independent articulation of the same mechanism this section develops doctrinally, naming the same audience — professional indemnity insurers — that Section 6.4 identifies as the institutional audience for whom this framework is ultimately most consequential.

6.4 Insurers

The professional indemnity insurer occupies a distinctive position in this picture, because it is the party that ultimately prices and bears the consequence of the staffing decisions described in Section 4.6, without, under current practice, any visibility into whether those decisions were sound at the time they were made. An insurer underwriting a firm's professional liability exposure today does so on the basis of the firm's claims history, its general risk profile, and its aggregate scale — not on the basis of whether the firm's internal staffing process reliably matches practitioners to the technical content of the matters it accepts. The signal set in Section 4.1, and the Realisable Resolution Value relationship in Section 4.4, describe exactly the information an insurer would need to price that risk more accurately, and to do so before a claim materialises rather than only when reconstructing one afterward. This is, on the analysis in Section 4.5, the institutional audience for whom this framework is ultimately most consequential: the insurer carries the financial exposure that follows from a mismatch it currently has no structured way to detect in advance.

6.5 AI Systems and the Generalisation of the Fallacy

The fallacy is not confined to human professionals. An AI system deployed for a specialised task — legal research in a narrow regulatory area, medical triage, technical diagnosis — can be described, marketed, or relied upon by reference to a general capability claim (“a large language model trained on legal text,” “a diagnostic model validated on clinical data”) that functions exactly as a qualification does: a threshold claim about general category competence that provides limited information about semantic proximity to a specific, narrow technical task. A general-purpose model with no demonstrated proximity to a specific regulatory regime or a specific class of technical dispute occupies the same *qualified, not expert* quadrant as the newly admitted lawyer in Section 1.1, and the same institutional question

applies: does the deploying organisation ask only whether the system is broadly capable, or does it separately measure how close the system's demonstrated performance sits to the specific task being delegated to it. This is identified here as a distinct line of inquiry — the application of expertise-distance measurement to the evaluation and deployment of AI systems themselves, including questions of training-data provenance and intellectual-property attribution that connect to a separate body of work this author has undertaken on IP economics — to be developed as a dedicated paper rather than absorbed into the present one.

Section 7: Reform Proposal

7.1 A Two-Stage Test

This paper proposes that institutions allocating professional and expert roles in technically specialised, high-complexity, or high-stakes matters adopt a two-stage test in place of the single qualification check currently in routine use. The first stage remains the qualification check: is the person entitled, as a matter of credential and professional standing, to undertake work of this general type. The second stage, currently absent from routine practice, asks the question this paper has developed: what is the semantic distance, measured against the signal set in Section 4.1, between the person's demonstrated subject-matter experience and the specific technical content of this matter.

Neither stage should operate as an automatic disqualifier in isolation. A finding of significant semantic distance is not, by itself, proof of incompetence in the specific matter; it is a finding that the usual basis for assuming competence — accumulated, demonstrated, subject-specific experience — is not present, and that some other basis (supervision by a genuine specialist, engagement of independent technical input, disclosure to the client or court of the gap) should be required before the matter proceeds on the assumption that qualification alone suffices.

7.2 Operationalising the Signal Set

The case study in Section 5.4 illustrates one possible operational approach: a structured comparison of the matter's stated technical requirements against the practitioner's and the organisation's demonstrated background, scored across the dimensions identified in Section 4.1 — domain match, technical depth, experiential proximity, and capacity fit — to produce both a single headline distance figure and a more granular dimensional breakdown capable of identifying *which* aspect of the gap is largest. This paper does not claim that the specific prototype illustrated in Section 5.4 is the only, or even the best, way to operationalise the measure; it claims only that the underlying quantity — semantic distance between task requirement and demonstrated experience — is measurable in principle, that a worked prototype applied to a real fact pattern returns results consistent with independent qualitative

analysis of the same fact pattern, and that the absence of any such measurement in current institutional practice is the gap this paper identifies as the more fundamental problem.

A second operational layer follows from Section 4.4. Where a matter's resolution value V can itself be estimated — from claim documents, contractual terms, or the value of the underlying asset or right in dispute — the same instrument that produces the semantic distance signal can, in principle, combine the two into an estimate of Realisable Resolution Value, giving an insurer, a client, or a court a single figure that expresses how much of a matter's existing value is likely to survive its handling, rather than two unconnected outputs. The prototype's own published output already gestures toward this second layer, directing the reader to a separate, more detailed financial-pathway module for escalation probabilities, settlement timing, and cost-at-risk modelling; that module is not examined in this paper and no reliance is placed on it here, but its existence is noted as a sign that the signal set and the value estimate are, in practice as in principle, intended to be read together rather than separately.

7.3 Objections Considered

Three objections deserve direct consideration. The first is that any computational instrument of the kind illustrated in Section 5.4 is only as good as the inputs available to it — a practitioner's public-facing professional biography may understate or fail to capture genuine relevant experience, and the tool's output should be treated as a structured starting point for human judgment, not a substitute for it. This is accepted without qualification: the proposal in Section 7.1 is a two-stage test conducted by an institution, not an automated determination, and the signal set is an input to that judgment rather than a replacement for it.

The second is that semantic distance, as a measure, does not yet have the kind of established mathematical literature behind it that supports comparable work in network science; it is, at this stage, a conceptual framework with a single worked illustration rather than a peer-validated methodology. This is also accepted, and the paper presents the framework accordingly — as a contribution to be tested and refined, not as a settled instrument. The signal set in Section 4.1 is deliberately built from categories

(domain match, technical depth, experiential proximity, capacity fit) that are independently verifiable from public record, precisely so that the underlying measurement can be checked and challenged rather than taken on faith.

The third objection is more substantive: that a two-stage test risks licensing satellite litigation over every professional appointment, turning every represented party's dissatisfaction with an outcome into a collateral challenge to their own or the opposing side's counsel. The answer is that the proposal is targeted, deliberately, at matters of genuine technical complexity and high stakes — precisely the category identified in Section 2.6 as the category in which the qualification shortcut systematically fails — and is not intended to apply to the routine matters in which, as Section 2.5 explains, the shortcut continues to perform adequately. A test confined to that category, and operated by institutions (insurers assessing risk, courts vetting appointments, firms reviewing their own staffing) rather than by opposing litigants seeking a procedural advantage, does not have the same expansive and destabilising character that a general right of collateral challenge would.

Section 8: Conclusion

Qualification and expertise are not the same measure of competence. That distinction is not this paper's discovery: labour economists have treated credentials as an imperfect signal of ability since Spence and Stiglitz, evidence law has in places already built a second, fit-specific inquiry into the admissibility of expert testimony, and the law of agency, professional indemnity insurance, and vicarious liability has for a long time recognised that a representative's failure can be, in substance, an institutional failure rather than a personal one. What institutions in many of the settings this paper examines still lack — particularly in civil-law jurisdictions, in the routine allocation of counsel rather than expert witnesses, and in the staffing decisions firms make internally and rarely disclose — is a structured, observable, ex ante way of asking the second question at all, rather than leaving it to discretion, marketing materials, or hindsight.

This paper has proposed a signal set capable of doing that: domain match, technical depth, experiential proximity, and capacity fit, each drawn from material a practitioner or organisation already discloses. No new theory was required to connect that signal set to consequence. Realisable Resolution Value applies the standard expected-value-of-litigation model — already decades old in law-and-economics scholarship — to competence distance specifically, giving the probability term in that model an observable basis. The three-directional liability that follows from a competence failure in representation, and the institutional rather than individual location of responsibility for it, both rest on settled doctrines of agency, insurance, and vicarious liability that this paper has connected to the signal set rather than invented. The contribution, stated at this modest and accurate scale, is the connection itself: a way of detecting, before a claim is filed, the fact pattern that agency law, insurance law, and the law of vicarious liability already know how to handle once it has been detected. Section 4.7 adds the reason this detection cannot, in practice, be left to the client alone: legal services are a credence good in the technical economic sense, the client's reliance on firm reputation is a structurally reasonable response to that fact rather than a lapse of vigilance, and the ordinary economics of how large firms staff their matters give the mismatch a route to occur without anyone involved intending it — a structural account this paper is careful to distinguish from

the deliberate exploitation of the same asymmetry, which is a different matter addressed elsewhere.

This paper has a sister instrument, addressed briefly in Section 1.2 and not revisited since: an examination of relational rather than semantic distance, measuring whether a court-appointed expert is connected to a party rather than whether a represented party's counsel is competent to handle the matter assigned to them. The two stand independently, answer different questions, and neither substitutes for the other. The case study, examined in Section 5, demonstrates what the semantic measurement looks like when applied to a single, clean fact pattern: a fully qualified, properly admitted practitioner, with no relational question of any kind, whose own public record places her at a substantial and — on the worked illustration in Section 5.4 — measurable distance from the matter she was assigned. Institutions that ask only whether a person is qualified will find nothing to query in that fact pattern. Institutions willing to ask the second question — a question the law has, in other contexts, already learned to ask — will find a great deal.

Bibliography

Primary Legal Sources

Lei n.º 78/2001, de 13 de Julho — Regime dos Julgados de Paz.

Daubert v Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579 (1993).

Kumho Tire Co. v Carmichael, 526 U.S. 137 (1999).

Federal Rules of Evidence, Rule 702 (Testimony by Expert Witnesses).

American Bar Association (2024) Model Rules of Professional Conduct, Rule 1.1 (Competence) and Comment.

Independent Directory Sources

This Author's Related Work

Carroll, S. (2026) *Network Proximity and the Appearance of Impartiality: Applying Erdős-Equivalent Distance Measures to the Reform of Court-Appointed Expert Standards in Civil Law Jurisdictions*. SSRN Working Paper No. 6905898. Lisbon: Centro de Estudos Judiciários Library (Biblioteca Armando Leandro).

Carroll, S. (2026) *The Manufacturing of Bad Faith*. SSRN Working Paper No. 6972499. (On the deliberate exploitation of an information or power asymmetry by an actor who fully understands the asymmetry and profits from maintaining it — the contrasting category referred to in Section 4.7 of the present paper, which is not addressed here.)

Carroll, S. (2017) 'Interconnected Networks & Bipartite Graphs — Penguin4,' LinkAudit.co Blog. Archived: <https://web.archive.org/web/20210923180914/http://linkaudit.co/blog/interconnected-networks-bipartite-graphs-penguin4/>

KYC.co (2026) *Expertise Check — Semantic Delegation Distance*. Prototype tool. Available at: <https://kyc.co/expertise-check>

KYC.co (2026) *Network Proximity Check*. Prototype tool. Available at: <https://kyc.co/network-proximity>

Academic Sources — Credence Goods, Information Asymmetry, and Firm Economics

Darby, M.R. and Karni, E. (1973) 'Free Competition and the Optimal Amount of Fraud,' *Journal of Law and Economics*, 16(1), pp. 67-88. (Foundational source for the concept of credence goods — goods or services whose quality the buyer cannot reliably assess even after consumption — citing legal, medical, and repair services as paradigm examples.)

Akerlof, G.A. (1970) 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism,' *Quarterly Journal of Economics*, 84(3), pp. 488-500.

Dulleck, U. and Kerschbamer, R. (2006) 'On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods,' *Journal of Economic Literature*, 44(1), pp. 5-42.

Arrow, K.J. (1963) 'Uncertainty and the Welfare Economics of Medical Care,' *American Economic Review*, 53(5), pp. 941-973.

BCG Attorney Search (2025) 'Am Law Economics Primer: How Firms Make Money (Leverage, Rates, Realization).' (On leverage as a documented driver of large-firm profitability, by which client matters are billed at senior rates while a substantial share of the underlying work is performed by associates compensated at a small fraction of the client-facing rate.)

Academic Sources — Signalling, Credentialism, and Expertise Measurement

Spence, M. (1973) 'Job Market Signaling,' *Quarterly Journal of Economics*, 87(3), pp. 355-374.

Spence, M. (1974) *Market Signaling: Informational Transfer in Hiring and Related Screening Processes*. Cambridge, MA: Harvard University Press.

Stiglitz, J.E. (1975) 'The Theory of "Screening," Education, and the Distribution of Income,' *American Economic Review*, 65(3), pp. 283-300.

Collins, R. (1979) *The Credential Society: An Historical Sociology of Education and Stratification*. New York: Academic Press.

Dreyfus, S.E. (2004) 'The Five-Stage Model of Adult Skill Acquisition,' *Bulletin of Science, Technology & Society*, 24(3), pp. 177-181.

Ericsson, K.A., Krampe, R.T. and Tesch-Römer, C. (1993) 'The Role of Deliberate Practice in the Acquisition of Expert Performance,' *Psychological Review*, 100(3), pp. 363-406.

Liu, Y. et al. (2025) 'KABB: Knowledge-Aware Bayesian Bandits for Dynamic Expert Coordination in Multi-Agent Systems,' arXiv preprint. (On distance-based expertise matching between an agent's demonstrated profile and a task's requirements, including an explicit semantic-mismatch dimension, in a computer-science multi-agent setting; cited as a related but independent application of distance-based expertise measurement, not as authority on the legal or liability setting addressed in this paper.)

Academic Sources — Litigation Economics

Priest, G.L. and Klein, B. (1984) 'The Selection of Disputes for Litigation,' *Journal of Legal Studies*, 13(1), pp. 1-55. (Foundational source for the expected-value model of litigation — value discounted by probability of a favourable outcome — which Section 4.4 applies to competence distance specifically.)

Academic Sources — Professional Liability, Agency, and Insurance

Consolo, C. and De Cristofaro, M. (2012) 'The New Italian Rules on Civil Procedure: A Substantial Reform,' *Civil Justice Quarterly*, 31, pp. 348-371.

Jolowicz, J.A. (2000) *On Civil Procedure*. Cambridge: Cambridge University Press.

Taruffo, M. (2012) 'Scientific Evidence in Civil Litigation' in: Edmond, G. (ed.) *Expertise in Regulation and Law*. Aldershot: Ashgate, pp. 34-42.

Von Bar, C. and Clive, E. (eds.) (2009) *Principles, Definitions and Model Rules of European Private Law: DCFR*. Munich: Sellier.

Sealy, L. and Hooley, R. (2020) *Commercial Law: Text, Cases, and Materials*. 6th edn. Oxford: Oxford University Press. (On the law of agency and the legal consequences of one party acting in the place of another.)

Baker, T. (1996) 'On the Genealogy of Moral Hazard,' *Texas Law Review*, 75(2), pp. 237-292. (On the function of liability insurance in pricing and allocating the risk of professional error.)

Cane, P. and Goudkamp, J. (2018) *Atiyah's Accidents, Compensation and the Law*. 9th edn. Cambridge: Cambridge University Press. (On vicarious liability, negligent supervision, and the allocation of responsibility between an individual actor and the institution that directs their work.)