On Legal Bullshit

Frankfurt's Distinction and the Self-Reporting Fallacy in Litigation

Sobre a Treta Jurídica

A Distinção de Frankfurt e a Falácia da Autodenuúncia no Contencioso

Working Paper — SSRN Submission Draft
June 2026

Steven Carroll, Founder, KYC.co / Moral.Money
kyc.co | kyc.co/dispute-engine | SSRN # [pending]
ORCID 0009-0005-1577-4337

Table of Contents

Abstract....3
Section 1: Introduction....5
Section 2: The Shadow of Settlement and the Forfeiture of Small Things....15
Section 3: The Existing Literature, and What It Has Not Asked....20
Section 4: Three Signals, Held Together....27
Section 5: The Moment Cover Runs Out....34
Section 6: Real-Time Burden, Not Eventual Judgment....39
Section 7: Every Case, All the Time....43
Section 8: Naming the Gap, and the Case for Building Across It....47
Section 9: Objections Considered....52
Section 10: Conclusion....54
Bibliography....57

Abstract

Harry Frankfurt drew a distinction that the law has never formally used: a liar knows the truth and asserts its opposite, while a bullshitter is indifferent to the truth altogether, producing assertion calibrated to what will work rather than to what is so. The distinction matters because the two are almost impossible to tell apart from their output alone, and because the law's machinery for detecting deception — intent, knowledge, scienter — is built entirely for the first category and has essentially nothing to say about the second. A professional who files a claim their own evidence already contradicts may have lied. They may equally have never checked, because checking was not necessary to produce a filing that does its job. The document looks the same either way. This paper argues that civil litigation, professional liability, and the institutions that sit above both have organised themselves around exactly this blind spot, and names the structural reason it persists: the Self-Reporting Fallacy.

The fallacy is this: every actor positioned to read a legal file has a stake in the answer it gives. A client depends on trusting their own lawyer's account — a client who does not trust it has already, in effect, dismissed their counsel. A lawyer who filed a document is committed to its coherence regardless of who created the gap inside it. An insurer receives a summary of the file, not the file. A court sees what each side chooses to argue, not what either side has chosen not to mention. None of this requires bad faith to operate; it requires only that each actor behave exactly as their position already rewards them for behaving. The result is a system that can prosecute the rare, provable lie but has no mechanism at all for the far more common case — the assertion nobody checked, made by someone with every reason not to check it, sitting undetected in a record that would have shown the gap to anyone reading without a stake in the outcome.

This paper specifies what a reading position with no stake in the outcome would need to do differently, building on this author's prior work on the distance between an assigned professional's competence and a matter's actual demands, and on the economics of strategic silence under current legal defaults. It adds a third element, specified here for the first time: a test for whether a filing's own positive factual claims survive contact with its own attachments — and with what the record conspicuously fails to address — independent of whether the gap originated in incompetence or in design. Demonstrated against one recurring event — the moment a professional's own filed evidence contradicts their own asserted position, converting a routine dispute into an independent professional-liability exposure that frequently dwarfs the underlying claim — the combination shows what becomes visible only once Frankfurt's distinction is taken as seriously in litigation as it has been taken in moral philosophy for forty years: not

who is lying, but who has stopped checking, and what that costs everyone still relying on them to have checked.

This paper also situates the resulting method against the existing landscape of advocacy ethics, litigation analytics, and legal-AI document review, and argues that none of these literatures, individually mature as each is, was ever positioned to ask the question this paper asks. It closes by naming the category of instrument such a method makes possible, locating the professional-indemnity insurance industry as its natural first market, and offering, as existence proof rather than as its subject, a prototype instrument built by this author that demonstrates the method is achievable with present techniques.

Section 1: Introduction

1.1 The Mechanism in Plain Terms

Because what follows is dense, it is worth stating the whole of it once, plainly, before the apparatus arrives — so that the reader has the shape of the argument in hand throughout, and can see, at each step, which part of the problem that step addresses.

Most civil disputes that reach litigation do not begin with a lie. They begin with two people who each believe they can win. That belief rarely comes from nowhere. It comes, in the great majority of cases, from a lawyer — because that is what a lawyer is for, and because that is what a lawyer is incentivised to say. A lawyer who tells a prospective client the case is weak, do not proceed, does not get hired to litigate a multi-year legal dispute. A lawyer brought in after another has already lost ground does not keep the retainer by agreeing the position is hopeless; the pitch is always some version of I think I can win this, I know a way through it. This is not a claim that lawyers lie about the odds. It is a claim about what the incentive structure selects for: the lawyer who sounds confident gets the case, and the client who hears confidence keeps fighting. Most people who end up in a long dispute are there, at least in part, because someone competent and credentialed told them they could win it.

The case is then built, by both sides, the way a house of cards is built — and it can look, for a long time, more solid the larger it gets. Each filing adds a card. Each motion answered, each exhibit produced, each procedural point won makes the structure look sturdier than it was the week before. This is not an illusion either side is necessarily aware of producing. Most of what goes into a sustained legal argument is genuinely solid, on both sides, which is exactly why the dispute reached court in the first place. Two genuinely strong cases can both be built this way, card upon card, for months or years.

The trouble starts when the arguing itself does what arguing is supposed to do: the weaker arguments are exposed as non-material, with small fixes offered to keep the core argument standing — justifications along the way, "we had no choice because of Variable A" — until the structural core is isolated, and one side finds the few cards that were never quite resting on anything. This is usually not a defect either side knew about at the outset; it might be a subconscious assumption, or something that was hidden from one of the parties. It is something the process exposes — a document that turns out to say less than it was assumed to say, a date that does not line up once someone checks it, a witness whose account does not survive a critical look. The hole is not invented.

It is discovered, often first by the very side it damages, partway through a structure they have already invested heavily in defending.

What happens next is the part this paper is actually about. A party who discovers, partway through, that their house of cards has a card that does not hold has three honest choices: concede the point, settle, or absorb the loss the point represents. Instead, they very often take a fourth option, because nothing forces the first three: they cover the hole. Not usually with an outright lie — that is rare, provable, and risky. More often with the thing this paper, following Frankfurt, calls bullshit: an assertion produced to hold the structure together, made by someone who has stopped checking whether it is true because checking would mean admitting the card does not hold. The lawyer who pitched the win has every reason to keep pitching it, because conceding now means admitting the earlier confidence was wrong. The client who has already spent a year and a great deal of money building the case has every reason to believe the cover is real, because the alternative is admitting the year was wasted. Nobody at this point needs to be acting in bad faith for the cover-up to happen. They need only to be doing what their position already rewards them for doing.

The reason this works — the reason the hole, once found, can be covered rather than conceded — is that nothing in the surrounding legal framework requires otherwise.

Here is what the same dispute looks like from the outside. Two parties enter the process, and by the time they have done so, at least one of them is misrepresenting something which there is no default obligation on them to divulge. Not probably — as a matter of structure. Nobody litigates over a fact that both sides agree on. The act of formal dispute is a signal, before a single document is read, that at least one version of events on the record is false.

Both parties know a gap of this kind is possible, and so both parties claim to be the wronged one. Both will be compelling; both will have documentation; both will have professionals constructing the most coherent story available from the facts in their possession. From the outside, an honest party in a weak position and a dishonest party with a strong story can be indistinguishable. You cannot simply detect the misrepresentation by listening to who is more insistent, or by looking at who filed documents first, or by measuring the confidence of the lawyers. The signal is not in what either party is asserting — it is in the gap of what one of them is avoiding saying, or avoiding replying to.

Every dishonest account has a gap in it. A missing document that should exist in an honest version of events. A date that is not described with the same specificity as the periods around it, and whose absence quietly tells a different story of the same narrative. A question ignored, or a statement left unchallenged, filing after filing, unanswered. Instead a cover story is constructed — a set of diversions, none of which are material breaches in themselves, but all manufactured distractions or credibility attacks. The gap is not random. It sits at exactly the point where the honest version of events would require the party to say something that damages their own position.

This does not mean either side is lying about everything, and it is worth being clear about that, because the more common case is narrower and more sympathetic than a wholesale falsehood. Most disputes that reach litigation are not contests between an honest party and a dishonest one across the board. Sometimes they are contests in which both sides have a genuinely strong case on most of what is in issue, and the dispute is really about one narrow, contested element — one fact, one document, one date — that each side believes, often in good faith, will break the whole case in their favour. Cases go to court because both sides think they can win, even if a win is simply evading responsibility for a known defect by not being caught to have known about it prior. The misrepresentation this paper is concerned with does not need to run through the whole of either party's account to matter. It only needs to sit at the one point the case actually turns on.

This is not a new observation, and it is worth referencing how old it actually is. Legal ethics scholarship has known for decades that the rules governing a lawyer's conduct impose no general duty to volunteer a damaging fact. Luban's discussion of Spaulding v. Zimmerman is the canonical instance: defence counsel, having learned through an examination conducted for the litigation that the plaintiff suffered from a life-threatening condition unrelated to the injury at issue, was professionally entitled to say nothing and let a settlement proceed without disclosing it. The lawyer did not lie. He was never asked, and the rules did not require him to volunteer the answer. What the case shows, and what the surrounding literature has spent decades confirming, is that the absence of a rule compelling disclosure is not a gap in the rules. It is the rule.

This author's companion paper on the manufacturing of bad faith generalises that single case into a structural claim: wherever a legal framework prices the withholding of relevant information at zero, and provides no cheap channel for compelling its disclosure before harm occurs, withholding is not the exception. It is the rational response of an ordinarily self-interested party, which the system will reproduce reliably, case after case, not because the parties involved are unusually dishonest but because nothing in the framework asks anything else of them. That paper develops the claim at the level of the dispute itself — the landowner who knows the shared aquifer beneath his land has been contaminated, the contractor who knows the cement mix was not checked, each protected, like the lawyer in Spaulding, by the simple absence of a rule requiring them to speak before the silence does its damage.

The same missing framework recurs here, one institutional layer higher, and the existing literature has not yet made the connection. There is no rule requiring a court to compare what a party asserted in month one against what the same party's own filings show in month six. There is no check when an extraordinary claim is made by one party, which would shift the burden of proof to that party. There is no flag alerting an insurer of increased risk. There is no rule requiring anyone, across a proceeding that may run for years and pass through several judges and several sets of counsel, to hold the whole record as a single coherent narrative and ask where it stopped adding up.

Call this the Self-Reporting Fallacy: the assumption, nowhere written down but everywhere relied upon, that a legal record will surface its own contradictions to someone capable of acting on them. It will not, for exactly the reason Spaulding already shows and the manufacturing-of-bad-faith analysis already generalises. The cost of looking falls on whoever would have to do the looking, and nothing requires anyone to look.

This paper specifies what looking requires. Three signals, taken together, make the gap detectable algorithmically without requiring any inference about intent: how far the professional behind a filing sits from the matter's actual technical domain; what each party's documented circumstances make it rational for them to assert, withhold, or delay; and whether a filing's own positive factual claims survive contact with the other attachments already filed. When they do not — when a party's filed evidence contradicts their own asserted position — that contradiction marks the moment a dispute between two parties becomes an independent professional-liability exposure, one that, when it does, lands on the insurer's desk, regardless of how the underlying case eventually resolves.

Professional-indemnity insurers already report large, recurring payouts driven substantially by case complexity and conflicts of interest — precisely the conditions under which this kind of gap is hardest for any single human reader to catch and easiest for a party to let pass unmentioned. The Ames & Gough annual survey — the industry's own benchmark survey of the largest professional-liability insurers, covering law firms in the Am Law 100 — found that by 2022, all but one of eleven surveyed insurers had paid out a single claim over $50 million within the prior two years, and four of eleven had paid out a single claim over $300 million.

If the problem is a missing framework, the correction is not moral exhortation but a framework — and the natural place to insert one is the instrument that already sits financially behind almost every lawyer in almost every dispute. Baker & Swedloff's research on liability insurers as private regulators — summarised in their "Insurers as Bumblebees in the Garden of Law Firm Norms" — establishes that insurers already impose discipline on professional conduct through the pricing and monitoring of risk, in ways no court or bar association directly supplies. Aligning a detection instrument to that existing incentive structure, as a shared condition of coverage between insurer and firm, makes commercial sense in a way that voluntary adoption by either alone does not.

In practice, professional discipline operates largely after the fact, through claims histories and premium adjustments once a loss has already occurred — for the same structural reason Spaulding's rule produces silence: nothing currently requires earlier intervention, because until now no instrument has existed to identify a contradiction at the moment it first arises. The natural test of this method's value is therefore prospective: deploy it across a live portfolio under exactly this kind of shared, mandatory monitoring arrangement, and flag in real time precisely the events that, historically, we know have created future liability claims.

Recent advances in large language models make this computationally achievable rather than merely theoretically desirable. Rather than proposing a framework and leaving its operationalisation to others, this paper introduces a working system that treats the detection of legal contradiction as an executable reasoning process — one that models assertions, evidential claims, procedural obligations, and logical dependencies as machine-evaluable relationships, and flags inconsistencies at the moment they arise rather than after harm has occurred. The accompanying implementation is offered as proof of concept: the gap this paper names is not merely real, but already detectable, and the sections that follow set out how.

1.2 The Fallacy

A legal filing is read as a record of what happened. It is a record of what someone wanted believed. The two are not always far apart, but the cause of the distance between them is never knowable from the document alone, and the law has never built a method for telling the two causes apart. A professional may misstate a fact because they did not understand the matter well enough to see the misstatement. The same professional may misstate the same fact having understood the matter precisely well enough to know that nothing downstream would check it. Harry Frankfurt's distinction between the liar and the bullshitter describes the second case more precisely than any term currently used in legal practice: the liar knows the truth and asserts its opposite, while the bullshitter is simply indifferent to whether an assertion is true, producing whatever serves the purpose at hand. Litigation has a name and a doctrine for the first failure. It has none for the second, and the second is very much more common, because it is very much cheaper to produce and very much harder to prove.

This is not a marginal gap in legal vocabulary. The entire apparatus the law uses to police false statement — intent, knowledge, scienter — is built for the liar, who must be shown to have known the truth and chosen against it. None of it reaches the professional who never checked, because checking carried no professional cost and not checking carried no professional risk. A claim, a report, or a defence produced this way is not a lie in any sense a court currently tests for. It is bullshit in Frankfurt's exact sense, and it is filed, argued, and frequently won on, every day, by professionals operating well inside what their own trade considers good practice.

That trade has a name for the underlying conduct, and a celebrated one. The duty to represent a client zealously, within the bounds of the law, has stood at the centre of the profession's self-conception for as long as the profession has had a self-conception, and the considerable literature examining that duty has already done the work of showing what it does to the lawyer's relationship to everyone who is not the client: zealous advocacy, on the dominant reading the profession has defended against its own internal critics for decades, limits the lawyer's moral obligation to the party paying for it and minimises to almost nothing whatever obligation would otherwise run to the other side, to the court, or to the truth of the matter at hand. The profession rewards the result, not the manner of reaching it, for as long as the manner stays inside the boundary the rules define — and a result reached by genuine merit and a result reached by procedural exhaustion, by a jurisdictional technicality, by an opponent who simply gave up first, are not distinguished in how they are received, billed, or remembered. The more esoteric the method, the more the profession tends to admire it. Nothing in that reward structure asks whether the underlying assertion was true. It asks only whether it held.

What holds the entire arrangement together is the boundary itself, and the fact that the boundary can only be crossed with cover. A lawyer operating at the furthest edge of zealous advocacy is, by definition, operating where the personal risk of crossing into actionable misconduct is highest — and no lawyer crosses that line without believing, correctly or not, that something stands behind them if they do: a defensible interpretation, an ambiguity in the record, a position a court might plausibly accept even if it does not. The professional indemnity insurer that ultimately stands behind the firm has, in the ordinary case, no visibility into how close to that edge any given filing actually sits, and prices the firm's risk as if the distance were uniform across its book. It is not uniform. Some matters are argued by professionals operating with real cover — a genuinely contestable reading of the facts. Others are argued by professionals who have already exhausted their cover, and are continuing regardless, because the alternative is conceding a position they have already filed.

The moment that cover runs out is not, on the present state of legal analysis, an event with a name, a date, or a detection method. It should be. It is the moment a professional's own filed evidence contradicts the professional's own asserted position — not the moment a court eventually notices, which may be months or years later if it happens at all, but the moment the contradiction is created, sitting in the record from that point forward, available to be found. Before that moment, the represented party's litigation risk and the firm's professional indemnity risk are the same number, because the firm's only exposure is the case it is running. After that moment, the two numbers diverge, because the firm has created a second, independent exposure that exists regardless of how the underlying dispute resolves — and the cover the lawyer believed they were operating under, the thing the zealous-advocacy boundary was supposed to permit, has in fact already failed.

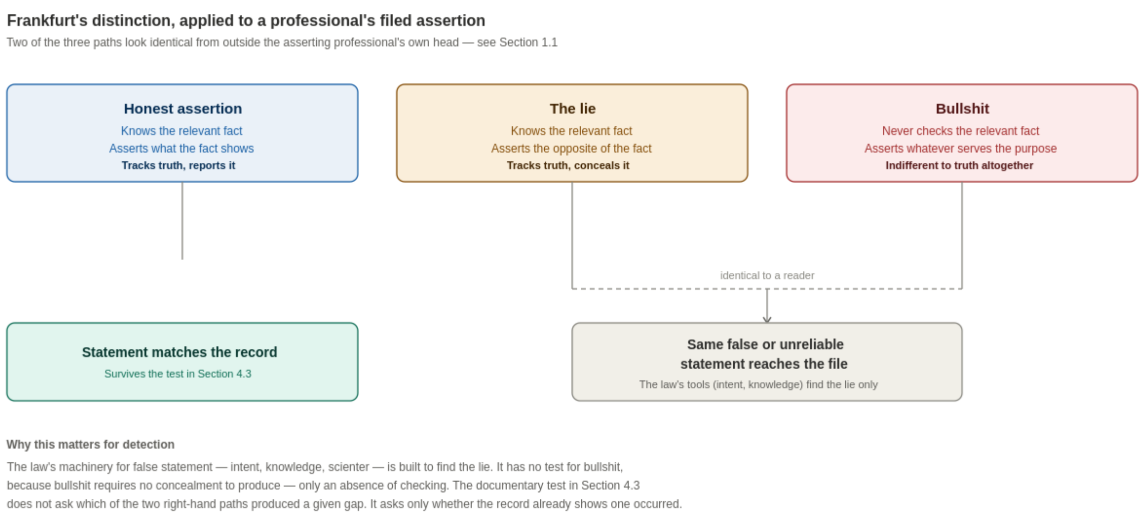

Figure 1. Frankfurt's distinction applied to a professional's filed assertion. The lie and bullshit both end in a false or unreliable statement reaching the file, but only the lie is within reach of the law's existing tools for detecting deception. See Section 1.2.

1.3 Why Negligence, Exploitation, and Entrapment Cannot Currently Be Told Apart

Negligence and exploitation are not opposite failures. They are the same failure read from two different motives, and they produce identical paper. A professional who lacked the specific competence to recognise a contradiction in their own filing, and a professional who recognised it and judged correctly that no one would check, file documents that look the same to anyone reading them afterward — including, very often, to the professionals themselves, who in the first case may never learn the difference and in the second have every reason never to volunteer it. Each side of this coin benefits from being mistaken for the other. The negligent professional is assumed merely unlucky rather than examined for competence. The exploiting professional is assumed merely negligent rather than examined for intent. Both reach that assumption by the same route: nobody positioned to read the file has a reason to ask which one occurred, because the question is not currently one that any institution around litigation is built to ask.

There is a third face to this coin, and naming it matters because it accounts for a case the first two do not: the competent professional, working squarely inside their own area of expertise, who continues to defend a position they have come privately to doubt — not because they never understood the file, which is negligence, and not because they coldly calculated that no one would check, which is exploitation, but because the psychology of escalating commitment to a course already embarked upon makes reversal feel more costly than continuation, even where continuation is the worse choice on the merits. The social-psychology literature on entrapment in escalating conflicts documents this mechanism directly, and documents in particular its vicarious form, in which the escalating commitment is felt on behalf of a party to whom one is psychologically connected — which is precisely the position of a lawyer bound to a client. The professional trapped this way produces the same paper as the negligent one and the exploiting one, from a third and again distinct motive, and is again shielded by the fact that nobody reading the file is positioned to ask which of the three occurred. The purpose of naming all three is not to sort filings by motive — the method specified in Section 4 deliberately declines to infer intent at all — but to establish that a single documentary gap can arise from at least three structurally different sources, none of which the record distinguishes on its face, and all of which the current framework leaves equally unexamined.

This is what makes the contest that follows so durable. Once a position is even slightly exposed, the visible energy of litigation shifts from establishing what happened toward multiplying who else might be made to answer for it — co-defendants joined on thin grounds, regulators implicated, predecessors in title blamed, intervening causes invented, each new name diluting the court's attention rather than resolving the dispute. None of this requires anyone in the file to act in bad faith by the standards of their own role. The client trusts the lawyer's account, because a client who interrogates every filing has already dismissed their own counsel. The lawyer who filed the document is committed to defending its coherence, whether or not they created the gap inside it. The insurer receives a summary of spend and status, not the underlying documents, filtered by people whose job is triage rather than comparison. The court sees what each side chooses to argue, not what either side has chosen not to mention, and has neither the time nor the brief to set one party's exhibit against that same party's own words from a year earlier.

Call this the Self-Reporting Fallacy: the assumption, built into how every institution around litigation is structured, that a legal record will surface its own contradictions to the people best placed to act on them. It will not. The contradictions are not hidden. They sit in a filing and its own attachment, in a position taken in month one against a different position taken in month six, in an expert report that spends its length on what nobody disputes and almost none on what is actually in issue. They are visible to a reader with no stake in which conclusion turns out to be true. That reader does not currently exist inside the ordinary course of litigation.

1.4 What This Paper Adds

This paper makes four contributions. First, it names the Self-Reporting Fallacy and locates its cause precisely: not carelessness, but a predictable consequence of every actor around a dispute occupying a position that depends, in some specific and identifiable way, on the record's contradictions remaining unexamined. Second, it brings Frankfurt's distinction between lying and indifference to truth into a setting it has never been formally applied to, and shows that the distinction maps exactly onto a failure litigation cannot currently name: the gap between the provable lie and the far more common, far cheaper, and currently undetectable assertion made without anyone having checked. Third, it specifies a reading position immune to the fallacy — one that holds three signals together across an entire file at once, where the literature to date has only ever developed each in isolation: how far the professional standing behind a filing sits from the matter's actual technical demands; what each party's documented circumstances make it rational for them to assert, withhold, or delay; and whether a filing's own positive factual claims survive contact with its own attachments and with what the record conspicuously fails to address, independent of whether the gap behind them was made carelessly or by design. Fourth, it demonstrates the combination against one recurring event — the precise moment a professional's cover runs out, converting an ordinary dispute into an independent professional- liability exposure that frequently exceeds the underlying claim many times over — and argues that this event, long recognised by practitioners only in hindsight, becomes locatable at the time it occurs once a file is read this way.

1.5 Structure

Section 1.1 has already stated the whole of the argument once, plainly and without apparatus; the remainder of Section 1 restates it in the paper's own terms — the fallacy itself in 1.2, the three distinct motives a single documentary gap can arise from in 1.3, and what this paper adds to the existing literature in 1.4 — before the sections that follow develop each element at length.

Section 2 develops the shadow settlement casts over the entire premise that a judge will eventually review a file in full, and situates the procedural doctrines of forfeiture and harmless error as the law's own, already- settled account of how unforgiving the system is toward a gap nobody caught in time. Section 3 situates the Self-Reporting Fallacy against the existing literature on advocacy ethics and litigation analytics, and states precisely what that literature has not yet asked. Section 4 specifies the three signals a disinterested reading position must hold together, building on this author's prior work on competence distance and on the economics of strategic silence, and develops the third signal — internal documentary consistency, including conspicuous silence — for the first time. Section 5 shows what the three signals reveal in combination that none reveals alone, working through the moment a professional's cover runs out as the clearest demonstration available. Section 6 specifies what it would mean to flag that divergence in real time rather than waiting for a court's eventual and often absent attention. Section 7 makes the case for applying the method continuously, across an entire portfolio, rather than only once a matter already looks troubled. Section 8 names the resulting category of instrument, situates it against the existing legal-AI and insurance-analytics landscape, and explains why it has not been built before now. Section 9 considers the principal objections. Section 10 concludes.

Section 2: The Shadow of Settlement and the Forfeiture of Small Things

2.1 The Backstop That Rarely Arrives

The implicit assumption behind every argument that a sufficiently attentive judge would eventually catch a filing's internal contradiction rests on a premise that does not hold for the great majority of civil disputes: that the matter will eventually reach a judge who reviews it in full. It will not. The empirical literature on litigation outcomes has shown for decades that the substantial majority of civil claims resolve before final judgment — through settlement, withdrawal, or default — and never receive the kind of comprehensive judicial review that the "the court would have noticed" assumption requires. A contradiction buried in a filing during the life of a case that settles is not corrected later. There is no later. The case ends at the point it ends, on the facts the parties allowed the record to contain at that point, and whatever the contradiction was worth — to the settlement number, to one side's leverage, to the eventual allocation of cost — is the value it silently carried for the rest of the matter's life.

This is worth stating plainly because it inverts a comforting assumption embedded in most discussion of litigation risk. The eventual trial, and the eventual judge, function in most institutional thinking as a backstop: a final, comprehensive check that will, in the worst case, surface what earlier phases missed. For most disputes, that backstop is hypothetical. It is a check the system is structurally entitled to perform and structurally unlikely to be asked to. A risk model, a reserve estimate, or a litigant's own sense of their exposure that quietly assumes the backstop will operate if things get bad enough is a model built on a premise that, for most of the matters it will ever be applied to, is simply false.

It is worth being precise about why this matters more than a bare statistic might suggest. The assumption that a judge will eventually look properly does not merely fail occasionally, at the margins, in unusually quiet cases that happen to settle. It fails as the default outcome, which means that any institutional practice built around the expectation of eventual judicial correction — relying on a court to eventually notice an inconsistency, deferring scrutiny of a professional's filings to whatever review a trial might one day supply — is a practice built around the exception, while treating it as the rule. The ordinary path of a dispute is not toward a comprehensive reckoning. It is toward a negotiated exit, often conducted substantially in the shadow of incomplete information, in which whichever account of the facts currently looks more credible to the parties and their advisers, rather than whichever account a full judicial review would eventually have endorsed, sets the terms.

2.2 What the Procedural Doctrine Already Says

Procedural law has its own settled vocabulary for what happens to an error that is not raised at the moment it occurs, and that vocabulary already describes, with more precision than most litigants realise, exactly how unforgiving the system is toward delay. An objection not raised in time is not merely weakened; under the doctrine of forfeiture, the burden of proof itself inverts. A party who objects promptly typically requires the other side to show that an error was harmless. A party who fails to object in time must instead show, on appellate review, that the error rose to the level of plain error — a standard courts apply rarely, reserved for failures serious enough to threaten the fairness or integrity of the proceeding itself, and one that is met in only a small fraction of the cases where it is argued. The distinction between forfeiture and waiver sharpens this further: where a right has been knowingly and voluntarily given up, no appellate remedy exists for it at all, while a right merely forfeited through inaction at least leaves the much narrower plain-error avenue open — and a party who failed to object in time often has no straightforward way of proving, after the fact, which of the two categories their own silence falls into.

The harmless error doctrine compounds this further: even where a procedural deprivation is real and substantive, reviewing courts routinely hold the error harmless if they are not persuaded it changed the outcome — a judgment made, necessarily, after the fact, by a court reconstructing what difference a fact would have made to a process that already ran its course without it. This reconstruction is asked to answer a counterfactual that grows less answerable the longer the underlying proceeding has continued without the missing fact: a court asked, years after a procedural decision was taken, to imagine how that decision might have gone if a now-discovered document had been in front of it at the time, is engaged in an exercise the doctrine itself treats with considerable scepticism, precisely because hindsight reconstruction of this kind is acknowledged, even within the doctrine that requires it, to be unreliable.

Both doctrines are defensible on their own terms, and neither is the subject of this paper's criticism. A legal system that reopened every completed phase on every belatedly raised objection would never reach a final judgment at all, and the discipline that forfeiture and harmless- error doctrine impose is part of what allows litigation to conclude rather than recurse indefinitely. The doctrine exists because the alternative — unlimited reopening, on demand, at any stage, of any earlier ruling — would itself impose a cost on every other litigant in the system, in the form of delay, uncertainty, and the practical impossibility of any judgment ever becoming final. What the doctrine does not supply, and was never designed to supply, is a mechanism for surfacing the error before the moment for objecting to it has passed. Forfeiture doctrine governs what happens to an error once it is known and not timely raised. It says nothing about errors that are never raised because no party reading the file at the time was positioned to notice them — which, as Sections 1 and 3 of this paper argue, is closer to the ordinary case than the exception.

2.3 Phases, Attention, and the Disproportion of Sanction to Stake

A civil matter does not proceed as a single continuous inquiry. It proceeds in phases — pleadings, evidentiary motions, expert appointment, substantive hearing, each frequently presided over by a different judicial officer with no occasion to revisit a decision a colleague made in an earlier phase on a different record. A decision taken at one phase, such as the framing of questions for an expert or the scope of an evidentiary motion, is ordinarily taken on the facts in the file at that moment, and the system has no standing mechanism for asking whether a fact that surfaces three phases later — in a regulator's later determination, in a separate complaint, in a document a party only obtained through a subsequent disclosure request — would have changed the earlier decision. It usually simply does not revisit it. The earlier phase closed, and closure is the point of a phase.

This produces a structural opportunity that does not require any single dramatic act of misconduct to exploit. A fact omitted from one phase's record, whether through oversight or through calculation, very often does not need to be hidden for long, or hidden well. It needs only to remain absent for the duration of the phase in which its absence is consequential. Once that phase closes and the matter advances, the omission has, in practical effect, already done its work, regardless of whether it is later discovered. And the formal sanction available for that category of omission — a finding of bad faith in failing to inform the court, where such a finding is made at all — is, in most procedural regimes, modest: a fixed, nominal penalty entirely disconnected from the scale of what the omission may have decided. A professional weighing the cost of omission against the benefit of a favourable procedural ruling secured in its absence is, in any matter of meaningful value, weighing a small and capped downside against an uncapped upside. This is not a marginal incentive. It is, on the numbers most procedural regimes actually apply, a rational bet.

It is worth pausing on why this disproportion persists rather than being corrected, because the answer reveals something about how procedural sanctions are designed in the first place. A sanction for a procedural omission is set, almost universally, by reference to the gravity of the procedural conduct itself — was the failure to disclose wilful, was it repeated, did it waste the court's time — rather than by reference to the substantive value of whatever the omission concealed. This is a defensible design choice in isolation: a procedural sanction is meant to discipline conduct toward the court, not to substitute for an assessment of the underlying merits. But the consequence, in any matter where the substantive value at stake is large relative to the kind of sanction a court would impose for a first or merely careless-seeming omission, is exactly the asymmetry described above, and nothing in the design of the sanction regime corrects for it, because the sanction was never built with that asymmetry in view.

2.4 Why Correction Requires Extraordinary, Not Ordinary, Effort

The same doctrinal structure explains why correcting a missed point, once a phase has closed around it, is so rarely attempted and so seldom achieved. Reopening a closed phase on a later-discovered fact ordinarily requires more than identifying the fact. It typically requires forcing that fact into the record through a channel the original proceeding does not itself provide — a separate regulatory complaint, a parallel criminal referral, an access-to-information request directed at a different public body entirely — and then persuading the court that the resulting record justifies revisiting a decision it has already made. Each of those steps is its own undertaking, in its own forum, under its own procedural rules, and the work does not run in sequence so much as in parallel, under time pressure, while the original matter continues to move forward regardless. This is not the work a single competent litigator performs in the ordinary course of a single retainer. It is closer to running several proceedings at once, in service of reopening one decision in a sixth.

This is precisely why the result is so often that the point is never recovered. The skill required to do this work exists; able litigators do this kind of work, and when they succeed, the result can be decisive — a procedural decision made in ignorance of a fact is, once the fact is forced into the record and the decision is shown to have been made without it, often genuinely vulnerable to being unwound. But the fact that this is possible, in the hands of a sufficiently resourced and sufficiently determined party, is not the same as it being ordinary. Most matters do not have a party able or willing to run four parallel processes to recover one procedural point, and most professionals, facing the choice between that scale of effort and simply allowing a favourable but ill-gotten ruling to stand unchallenged, will not volunteer to do the former on the other side's behalf. Bullshit, in the sense developed in Section 1, does not need to be sophisticated to succeed under these conditions. It needs only to survive the phase in which it is told.

There is a further asymmetry worth naming explicitly. The party disadvantaged by an undisclosed or buried fact does not, in the ordinary case, know that the fact exists, or know to look for it in the specific forum where it might eventually surface. They are not failing to do the multi-forum work described above out of laziness; they are failing to do it because nothing in the original proceeding's record gives them a reason to suspect that a regulator, in an unrelated file, holds a determination that bears directly on their case. The extraordinary effort this section describes is not merely costly. It typically requires the disadvantaged party to already suspect that something is missing, before they have any concrete reason to suspect it — which is precisely the chicken-and-egg problem a method built to read across an entire accumulating record, continuously, exists to remove.

2.5 Death by a Thousand Cuts

The cumulative effect of this structure is rarely a single, decisive falsehood that loses a case outright. It is more often an accumulation of small, individually deniable, individually under-sanctioned omissions and procedural advantages, each easily explained away on its own terms, each absorbed into the ordinary noise of a multi-phase proceeding, none of which alone would draw serious scrutiny even if noticed. A case is far more often lost this way — by accretion, across many small points each conceded by default because nobody was positioned to contest them in time — than by any single dramatic act a court would recognise as misconduct if it occurred in isolation. This is the shape of harm the existing procedural doctrine is least equipped to see, because the doctrine reviews errors one at a time, on the question of whether any single error was harmless, and was never built to ask the cumulative question: whether a sequence of small, severally defensible omissions, taken together, determined an outcome that no single one of them could be shown to have caused alone.

The accumulation matters in a further, less obvious way. Each small, individually defensible point, considered against the harmless-error standard on its own, will very often genuinely satisfy that standard — a court asked whether any single one of ten minor omissions, in isolation, changed the outcome, may correctly conclude that it did not. The harm lies precisely in the fact that the doctrine has no comparable test for the aggregate: no established standard asks whether ten omissions, each individually harmless, were collectively decisive, because the doctrine's unit of analysis is the single error, raised and reviewed on its own terms, not the accumulated pattern across a file's full life. A method built to hold the entire pattern in view from the outset does not need a new legal standard to make this visible. It needs only to do what the existing standard was never built to do at the level of an individual error: track the pattern across the whole of a matter's life, continuously, rather than one isolated objection at a time.

This is the condition Sections 4 through 6 of this paper are addressed to. Where the cost of finding any one of these points, in time, has historically been the cost of a sufficiently resourced party doing the work of several parallel proceedings at once, a reading position built to hold a file's full record in view, continuously and from the outset, removes the dependency on any one party having the resources, the timing, or the willingness to do that work by hand.

Section 3: The Existing Literature, and What It Has Not Asked

3.1 The Standard Conception and Its Critics

The legal profession has examined its own relationship to zealous advocacy for longer, and in more depth, than it has examined almost anything else about itself, and the resulting literature already supplies most of what is needed to state Section 1's claim about reward and manner with precision rather than impression. The starting point for that literature, and the position its critics have spent several decades arguing against, is what is generally called the standard conception of the lawyer's role: the view that a lawyer's central professional obligation runs to the client, that the lawyer is entitled, within the bounds of the law, to pursue the client's interests by any lawful means regardless of the lawyer's own view of their merit, and that the moral responsibility for the result, where the means employed cause harm to a third party, belongs to the legal system that permits the means rather than to the lawyer who employed them. This last move is often called the adversary system excuse: the claim that a lawyer's role-specific conduct, even where it would be condemned under ordinary morality if performed by a private citizen, is rendered permissible, or at least non-culpable, by the lawyer's institutional role within an adversarial system whose overall design is thought to serve justice even though no single actor inside it is required to pursue justice directly.

David Luban's sustained critique of this excuse, developed across several decades of work culminating in his comprehensive treatment of the problem, remains the most thorough statement of why the standard conception produces exactly the gap between lawyerly conduct and ordinary moral accountability that Section 1 describes. Luban's discussion of Spaulding v. Zimmerman — a case in which defence counsel, having learned through a medical examination conducted for litigation purposes that the plaintiff suffered from a life-threatening condition unrelated to, but discovered because of, the injury being litigated, was professionally entitled to say nothing and allow a settlement to proceed without disclosing it — is worth recalling here in some detail, because it is as clean an illustration as exists in the published case law of the precise structure this paper is concerned with. The lawyer in that case did not lie. He was never asked the question whose honest answer would have been the undisclosed fact, and the rules governing his conduct of the matter did not require him to volunteer it. What he did was nothing, in the most literal sense available to him, and the nothing he did was, on the standard conception's own terms, exactly what his role obligated him to do. Luban's argument is that the rules permitting this outcome reveal something true and disturbing about the standard conception itself: it is a framework under which a lawyer's silence, however consequential to a third party, generates no moral remainder the lawyer is required to answer for, so long as the silence was never converted into an affirmative misstatement. The lawyer's role morality and the lawyer's ordinary moral obligations, on Luban's account, have been deliberately and permissibly decoupled.

This decoupling is precisely the structural condition Section 1 of this paper describes from a different angle. Luban's critique addresses itself to the moral status of the conduct the standard conception permits — is it defensible that a lawyer may stay silent under these conditions — and answers, at length and with considerable doctrinal support, that it is not easily defensible at all. This paper's concern is narrower and asks a different, and in a specific sense prior, question: given that the conduct the standard conception permits is, on Luban's account, already morally troubling, is there any institutional mechanism currently capable of detecting, from the documentary record itself, when that permitted silence has in fact occurred, as distinct from the alternative explanation that nothing relevant was ever known to be omitted. Luban's literature supplies the moral vocabulary for why the silence matters. It does not supply, and was not attempting to supply, a method for finding the silence in a specific file before the silence has already done its work.

3.2 Simon's Alternative, and Why Contextual Judgment Does Not Solve the

Detection Problem Either

William Simon's critique of the standard conception takes a different, and in some respects more radical, position than Luban's, and the difference is worth developing rather than collapsing into a single undifferentiated "critic of zealous advocacy" category, because Simon's specific diagnosis turns out to bear directly on a pattern this paper names in Section 4. Simon's argument is not, principally, that the standard conception commits lawyers to an adversary system whose moral cost is too high. It is that the standard conception's deeper flaw is its commitment to a particular and, in Simon's account, indefensible style of judgment: categorical, rule-bound, and rigid, applied mechanically to determine what a lawyer may or must do, in preference to the more contextual, purpose-sensitive style of judgment that Simon argues is both more honest about how legal judgment actually works in every other area of legal thought and more capable of producing outcomes that serve justice rather than merely procedural compliance. Simon's proposed alternative asks lawyers to exercise something closer to judicial discretion, oriented toward the purposes the law in question was meant to serve, case by case, rather than toward mechanical rule-following dressed up as ethical neutrality.

Simon's diagnosis is useful to this paper for a reason distinct from Luban's, and it is worth being precise about what that reason is. A filing built entirely from categorical, rule-bound moves — every formal element pleaded, every available procedural objection raised, every standard defence asserted in the conventional order — can be, on Simon's own account, exactly the kind of lawyering the standard conception rewards and exactly the kind of lawyering that produces the pattern this paper's Section 4.3 names: a document that is formally, categorically complete while substantively silent on the one matter that actually required a contextual, purpose-sensitive response. Simon's critique is not, on its own terms, a method for detecting this pattern in any specific filing; it is an argument that lawyers ought to reason differently than the standard conception currently asks them to. But the pattern Simon's preferred, more contextual style of judgment is meant to correct is precisely the pattern this paper's documentary test is built to surface mechanically, without requiring the individual lawyer reading the file to have first adopted Simon's preferred jurisprudential disposition. Where Simon asks lawyers to reason better, this paper asks whether a reading position external to the lawyer can detect, from the documents alone, the cases in which the categorical, rule-bound style Simon criticises has produced exactly the gap his own framework predicts it will.

It is worth noting, in fairness to the considerable internal debate this literature contains, that Simon's proposal has not gone unanswered within the field on its own terms. Robin West's response to Simon, arguing in defence of a version closer to the standard conception against what she characterises as Simon's underappreciation of the political and institutional value of predictable, categorical legal rules applied evenly regardless of a lawyer's individual judgment about a law's underlying purpose, is a serious and sustained piece of the same debate, and this paper does not take a position on which side of that disagreement is correct. What matters for the present purpose is only that the disagreement exists, that it is conducted entirely at the level of how a lawyer ought to reason about their own conduct, and that no position within it — not Luban's, not Simon's, not West's defence of categorical rule application — purports to supply a method for testing, from outside the lawyer's own head and after the filing has already been made, which style of reasoning actually produced the document in front of a reader.

3.3 The Boundary Itself, and Its Enforcement Mechanism

A related and more practically oriented strand of the literature addresses not the underlying justification for zealous advocacy but the operational question of where its boundary sits and how that boundary is enforced. This literature documents, in detail, the historical softening of the profession's own vocabulary — codes of professional conduct that once spoke unambiguously of zeal increasingly speak instead of diligence, and some jurisdictions have removed the word zealous from their formal rules altogether, while commentary within those same jurisdictions acknowledges, candidly, that the underlying disposition the word once named continues to operate in practice regardless of what the current rule book calls it. This strand of the literature also documents the mechanisms by which the boundary, however described, is actually policed: professional conduct rules prohibiting offensive tactics, disciplinary proceedings for the most extreme departures, judicial sanctions for conduct that crosses from advocacy into harassment or into affirmative misrepresentation to the court.

What this literature does not supply, and was never built to supply, is a method for testing, as a documentary matter, whether a specific filing has in fact crossed the boundary it describes. The literature is, in this sense, jurisprudential and institutional rather than forensic: it asks where the line should sit, debates how the profession should describe its own obligations, and catalogues the disciplinary consequences that attach to crossing it once a crossing has been identified and proven through the ordinary mechanisms of complaint, investigation, and adjudication. It does not ask, and has no apparatus for asking, whether a given filing, read against everything else already in the same record, has crossed the boundary at the moment the filing was made — which is exactly the gap between the formal allocation of scrutiny and the evidentiary reality Section 6 of this paper develops at length.

3.4 What Existing Legal Analytics Measures, and Why None of It Reaches This

A parallel and more recent literature has developed around the application of data and computation to litigation itself, and it is similarly mature within its own scope, though considerably younger than the advocacy-ethics literature it sits alongside. Three distinct categories of tool deserve to be named precisely, because the comparison this paper draws depends on describing each accurately rather than treating "legal AI" as an undifferentiated category.

The first category is outcome- and motion-prediction analytics. Platforms in this category — among the most established of which combine large historical databases of federal and state litigation with structured analytics on individual judges, courts, and opposing counsel — report detailed, granular insight into how a specific judge has ruled on comparable motions in the past, how often a specific litigator settles before trial, and what the likely range of damages or resolution outcomes looks like for a given case type in a given venue. Reported predictive accuracy for these systems, across the motion types and case categories where they have been most extensively validated, runs in the range of seventy to eighty-five percent — a genuinely strong empirical result, well above the accuracy generally attributed to unaided human prediction of the same outcomes, and one that has materially changed how some litigators approach early case assessment, settlement valuation, and venue strategy. None of this analytic apparatus, however sophisticated, asks whether the filings underlying any specific matter are internally consistent with themselves or with each other. The variable the system is built to observe is the outcome of the dispute and the behavioural tendencies of the human actors involved in deciding or arguing it; the variable this paper is concerned with — whether a specific assertion in a specific filing survives contact with the rest of the record — is simply not part of what these systems were built to compute, and nothing about the architecture that makes them excellent at the first task equips them for the second.

The second category is legal document-review and contract-analysis platforms, several of which were originally developed for, and remain most heavily used in, large-scale transactional due diligence rather than litigation as such. These platforms apply machine learning, in some cases built on proprietary models trained specifically for legal-document understanding rather than on general-purpose language models, to large sets of contracts or filings, and are genuinely effective at a specific and well-defined task: identifying clauses that deviate from a standard template or an organisation's own playbook, flagging unusual wording, surfacing missing or non-standard provisions across thousands of documents far faster than manual review would allow. This is anomaly detection against a template, and it is valuable precisely because most of what a due-diligence reviewer needs to find is exactly this kind of deviation from an expected pattern. It is also, on inspection, a fundamentally different operation from the one this paper is concerned with. A clause that departs from a standard template is flagged because it differs from what similar documents usually say, not because it has been tested against what a specific other document, filed by a specific other party in the same matter, actually shows. Detecting that a clause is unusual is not the same operation as detecting that a positive factual assertion contradicts material already in the same file, and a platform built and validated for the first task supplies no capability, by extension, for the second. Separately, and worth noting as a matter of honest comparison, several of these same platforms have begun extending their document-review capabilities into litigation-adjacent use cases such as evidence review in large-scale discovery and investigation matters, applying broadly similar classification and anomaly-detection techniques to litigation documents rather than only to contracts. This extension, while genuinely useful for the volume problem large discovery exercises present, remains built on the same underlying operation — surfacing documents or passages that match or deviate from a specified pattern of interest — rather than on the cross-document, fact-against-record consistency test this paper specifies.

The third category is insurance claims- and litigation-analytics platforms, used principally inside the insurance industry rather than by litigating firms directly. These combine structured claim data — dates, amounts, case type codes, venue

— with natural-language processing applied to adjuster notes and claim narratives, to score litigation propensity, flag emerging severity, and support earlier intervention on claims that the model identifies as likely to escalate. This is a genuinely valuable capability, and the industry's own published material on the subject treats it as a meaningful advance over manual, judgment-based triage. It operates, however, on summary fields and adjuster commentary rather than on a continuous, cross-filing reading of the underlying legal documents themselves, and its predictive target is the claim's likely cost and trajectory as a portfolio-level statistical matter, not the internal consistency of any specific filing against the rest of its own record.

A fourth and more recent development deserves brief mention because it approaches the incentive question this paper is also concerned with, from a direction worth distinguishing carefully. Recent work in computational modelling has begun to formalise litigation itself as a multi-agent, reward-structured simulation, in which parties are modelled as optimising across competing objectives — imposing cost on an opponent, securing favourable delay, advancing the underlying claim — under an explicit payoff function that rewards each independently of the claim's underlying merit. This work is a valuable formal confirmation, arrived at independently and through entirely different methods, of the same incentive structure this author's own prior work on the manufacturing of bad faith derives from first principles: that litigation's participants face payoffs which reward strategic behaviour, including delay and the imposition of cost, on terms unrelated to the merits of the underlying dispute. It is, however, a model of the game from the outside — a simulation of how rational agents would be expected to behave under a specified reward structure, used to study aggregate strategic patterns across many simulated matters — rather than a method for reading any specific, real file to determine what a specific actor has, as a documentary matter, actually done. The two are complementary rather than competing: this kind of simulation explains why the incentive structure Section 4.2 of this paper recalls should be expected to produce the behaviour it predicts, in the aggregate, across a population of rational litigants; it does not, and was not built to, locate any particular instance of that behaviour in any particular matter's actual documents.

That same body of work reports a finding worth stating precisely, because it bears directly on the second signal this paper develops in Section 4.2. When the simulated agents are made to optimise on the procedural payoff structure alone — with no instruction to behave adversarially and no exploit strategies authored into them in advance — they converge independently on recognisable exploit chains: inflating an opponent's costs, engineering delay through calendar pressure, and extracting settlement value on claims of low underlying merit. That such strategies emerge automatically from the reward structure rather than having to be written into the agents by their designers is exactly the confirmation this paper's incentive-structure signal predicts: the behaviour is a property of the incentive landscape itself, discoverable by any sufficiently capable optimiser placed inside it, and not a trait of any individual litigant's character.

3.5 The Question Nobody Has Asked

What unites the gap across all four of these literatures is not that any of them is incomplete on its own terms. Each answers a real and difficult question well, within the scope it was built to address. It is that none of them was built to ask the question Section 1 named: not whether a given piece of advocacy sits inside or outside the ethical boundary the profession has set for itself, not whether a given case is likely to be won, not whether a given clause deviates from a standard template, and not how a population of rational litigants is expected to behave in aggregate under a given incentive structure, but whether a specific assertion, made in a specific filing, survives contact with everything else already in the same record — and, if it does not, what kind of gap that failure most plausibly represents, given who made it and what their documented position made it rational for them to do.

The advocacy-ethics literature has the vocabulary for the boundary but no method for testing any given filing against it. The outcome-prediction literature has the method for forecasting results at scale but no interest in, or capability for, testing internal consistency. The document-review literature has the method for finding deviation from a template but not for finding contradiction against a specific, evolving, cross-party record. The claims-analytics literature has the method for scoring portfolio-level risk but operates on summaries rather than on the underlying filings. And the formal incentive-modelling literature explains why the behaviour this paper is concerned with should be expected to occur, in principle, without locating a single instance of it occurring in fact. The remainder of this paper specifies what asking the question none of these literatures asks actually requires.

Section 4: Three Signals, Held Together

4.1 Competence Distance, Recalled

The first signal, developed at length in this author's prior work on the gap between qualification and expertise, measures how far a professional's demonstrated experience sits from the specific technical content of the matter assigned to them. That earlier work distinguished qualification — a binary threshold variable, met or not met, encoding only that some defined minimum standard was satisfied at some point in the past — from expertise distance, a continuous variable describing how many conceptual and practical steps separate a professional's demonstrated experience from the specific technical content a given matter actually requires. The purpose of that distinction, in the earlier work, was to give the probability term in a dispute's expected value a structured, observable basis, rather than leaving it to unstructured intuition: a wide distance between a professional's demonstrated background and a matter's technical demands predicts a lower probability that the matter's underlying value will be realised rather than eroded through delay, mischaracterisation, or procedural error.

Its purpose here is narrower and more specific, and it is worth being precise about the difference. The earlier work used competence distance to discount value directly. This paper uses it diagnostically, to bear on which face of the coin described in Section 1.3 a given documentary gap is more likely to show. A wide competence distance makes an innocent miss a plausible explanation for a contradiction in the record — a professional working well outside their demonstrated area of depth is exactly the professional most likely to file an assertion without recognising that something already in their own possession contradicts it, because recognising the contradiction would itself require a degree of subject- matter fluency the professional's own career history does not show them to possess. A narrow distance does the opposite, and does so with considerable force. A professional demonstrably expert in precisely the technical area where a contradiction sits has a much thinner claim to have simply missed it, because the same expertise that qualifies them to handle the matter at all is the expertise that would, in the ordinary case, have made the contradiction obvious to them on a competent reading of their own file. The same documentary gap, found in the filing of a recognised specialist in the relevant domain, therefore carries a materially different prior than the identical gap found in the filing of a generalist operating well outside their depth, even though the gap itself — the bare fact of the contradiction — is identical in both cases.

Nothing about this signal resolves the question on its own, and it is worth restating this limit before moving to the next signal, because the limit is part of what makes the combination in Section 4.4 necessary rather than redundant. Competence distance is a statement about plausibility, not a verdict on what actually happened in a given instance. A wide distance does not prove negligence, and a narrow distance does not prove exploitation; either professional, in either configuration, might have behaved exactly as a careful, honest practitioner would have behaved under the same circumstances. The signal's contribution is to narrow the space of explanations a documentary contradiction admits, before any other signal is brought to bear on it — shifting a prior probability, not establishing a fact.

4.2 Incentive Structure, Recalled

The second signal, developed in this author's prior work on the manufacturing of bad faith, describes the payoff structure facing a party who holds information relevant to a dispute and must decide whether to disclose it. That paper's central formal result, recalled briefly here because Sections 5 and 6 of this paper depend on it directly, is that wherever a legal regime prices the withholding of relevant information at zero and the disclosure of it at some cost — financial, reputational, or strategic — a party holding that information has no scenario, given current procedural defaults, in which disclosure leaves them better off than silence. Withholding therefore dominates as a strategy, in the specific game-theoretic sense of dominance: it is the better choice regardless of what the other party does, which means the system settles into a stable equilibrium of silence that persists not because any individual actor is unusually dishonest but because the payoff structure makes silence the rational choice for an ordinarily self-interested participant.

Section 2 of this paper gave that abstract result a concrete procedural shape specific to the present argument: a procedural omission, even where it is later identified and sanctioned, typically draws a penalty entirely disconnected from the scale of what the omission may have decided, while the upside of a favourable ruling secured in its absence is uncapped. An actor facing that asymmetry is not behaving exceptionally, or even unusually boldly, by choosing silence. Silence is simply the equilibrium strategy available to them, on the numbers the system itself applies, and choosing it requires no departure from ordinary self-interested rationality.

This signal's contribution to the combination developed in this Section is to map, for any given actor connected to a matter, what their documented position — their role in the proceeding, what they have already filed, what the procedural rules applicable to their position make costly or costless for them to disclose — makes it rational for them to assert, withhold, or delay, independent of whether that actor has in fact done any of those things. It describes the motive landscape a contradiction sits inside, in the same way a map of a terrain describes which routes are easiest to travel without describing which route any particular traveller actually took. It does not, on its own, establish that any actor acted on the incentive it identifies, and a reading that treated this signal as proof of anything more than what an actor was rationally positioned to prefer would be claiming considerably more than the signal supports.

4.3 Internal Documentary Consistency, Introduced

The third signal is specified here for the first time, and it is, in the combination this paper proposes, the one that converts the other two from background context into an active finding. It is best understood as two related tests rather than one, because a filing can fail either independently of the other, each fails in a way ordinary sequential reading is poorly equipped to catch, and the two together cover considerably more of the territory this paper is concerned with than either would alone.

The first test is the one most naturally associated with internal consistency in the ordinary sense of the phrase: does a positive factual assertion made in a filing survive contact with material already present in the same accumulating record — the filer's own attachments, the filer's own earlier submissions, material filed by another party in the same or a related proceeding — independent of who is right about the underlying dispute. This test is restricted, deliberately and as a matter of principle rather than convenience, to assertions of fact rather than to legal argument. A jurisdictional challenge, a limitation defence, a contested valuation are positions, and reasonable parties can disagree about them, sometimes for years, without either being shown wrong by anything in the file; resolving a position is the court's job, and nothing in the documentary record settles it in advance. A claim that a licence was in force, that a payment was compensatory rather than something else, that a document says what a filing represents it as saying, is different in kind: the record itself either contains what would make the claim true or contains something that makes it false, and resolving which requires no judgment about the underlying merits at all, only a comparison of what was said against what the same record, taken as a whole, actually shows. This restriction matters because it is what keeps the test narrow enough to avoid trespassing on the court's own function. The test does not ask who should win. It asks only whether a particular factual claim, made along the way, is consistent with the rest of what the same party, or another party, has already put on the record.

The second test is broader, considerably harder for a human reader to sustain across a long or adversarial record, and, on this paper's assessment, the more consequential of the two in practice, because it surfaces a category of gap the first test cannot reach at all. The second test asks: does a filing actually engage with what the opposing side has specifically put in issue, or does it pass over the point in silence while addressing everything around it at length and in confident, well-organised detail. A defence can satisfy every formal element a competent litigator is trained to include — it can be properly pleaded, fully referenced, structured in the conventional order a court expects to see, citing the right authorities in the right places — while devoting the great majority of its length to matters the opposing side never seriously contested, and a sentence, or nothing at all, to the one point that was actually placed squarely in issue and that the case most plausibly turns on. Read in isolation, such a filing presents as thorough and well-prepared, because thoroughness, as ordinarily perceived by a reader, is a function of how much careful, competent-looking material a document contains, not a function of whether that material happens to be the material the matter actually required. Read against the specific allegation it was filed to answer, the disproportion between what was required to be addressed and what was actually addressed is itself the signal — not a contradiction in the strict sense the first test looks for, but an absence, conspicuous once named and very easily missed otherwise.

The second test matters specifically because of how human attention works under exactly these conditions, and this point deserves to be developed rather than merely asserted, because it is the basis for this paper's claim, taken up at greater length in Sections 6 through 8, that the capability described here is not merely a faster version of what a careful human reviewer already does. A reviewer working through a long, confidently drafted, properly formatted filing experiences its thoroughness on the matters it does address as a reason for reassurance, and that reassurance is strongest precisely where it is least warranted: on the page where the document quietly stops addressing the thing that mattered and moves on to the next properly handled, uncontested point. This is not a failure of diligence on the reviewer's part, and it would be a mistake to read this paper as criticising the competence of any individual reviewer who has experienced it. It is an ordinary property of sequential reading under time pressure, across volume, and it recurs regardless of how skilled the reviewer is, because the skill the moment calls for is not closer reading of what is on the page but active recall of what was never put there — holding, throughout, a separate and continuously updated list of everything the matter actually requires a response to, and checking each new filing against that list rather than against the impression the filing creates on its own terms. Maintaining such a list, accurately, across months or years of filings, while simultaneously reading each new filing closely enough to notice when its confident prose is quietly failing to engage with an item on the list, is a cognitively demanding task that competes directly with the ordinary, sequential business of reading the document in front of one for what it says, rather than for what it omits.

A method built to hold that list explicitly, and to compare it against each filing mechanically rather than impressionistically, does not experience the reassurance that makes the gap easy to miss, because it is not forming an overall impression of the document's thoroughness at all — it is performing a specific, narrow comparison between a defined list of required engagement and the document's actual content, repeated consistently regardless of how much unrelated material surrounds the gap. It is not reading faster than a human reviewer, and the value of the method does not depend on any claim that it reads with superior comprehension. It is reading for a different thing: not whether the document feels complete, but whether it is complete against a specific, named list of what was actually required of it. This is, in a precise sense, the same distinction Section 1 drew between the liar and the bullshitter, applied to the act of reading rather than to the act of writing: a reviewer experiencing reassurance from a thorough-seeming document has not been lied to in any technical sense, and may not even have been the deliberate target of anything — the reassurance is simply the predictable effect that a long, confident, properly structured filing has on sequential attention, whether or not the gap inside it was placed there by design.

4.4 Why None Is Sufficient Alone, and What the Combination Produces

Each signal, used alone, produces something considerably less useful than what the three produce together, and it is worth working through each case of partial application in turn before stating the combined result, because the weakness of each signal in isolation is precisely what explains why no single one of the existing literatures surveyed in Section 3 — each of which has, at most, ever developed something resembling one of these three signals — has produced the capability this paper specifies.

A competence-distance score alone states that a mismatch exists, somewhere in the relationship between a professional and a matter, without locating any specific instance in the file where that mismatch has actually mattered to anything currently at issue. A firm or an insurer presented with nothing but a competence-distance figure has been told that risk exists in the abstract, in the way an actuarial table tells an insurer that risk exists across a population, without identifying which specific policy in the book is the one where the risk has already materialised into an actual, locatable problem.

An incentive map alone states that an actor is rationally positioned to prefer silence over disclosure on some matter, without identifying anything that actor has, as a matter of documentary fact, actually kept silent about. This is, in a sense, the weakest of the three signals taken in pure isolation, because it describes a structural tendency common to almost every actor in almost every matter — the incentive to prefer favourable silence over costly disclosure is, after all, close to universal among parties to litigation — and a signal that applies almost everywhere, without further refinement, supplies very little information about any particular matter.

And a documentary gap found alone — whether a direct contradiction located by the first test of Section 4.3 or a conspicuous silence located by the second — is, without the other two signals, simply an anomaly. It might be inadvertent: a drafting error, a citation to the wrong exhibit, a genuinely ambiguous record that admits more than one fair reading and happens to have been read differently by two different filings without either reading being unreasonable. It might be immaterial: a contradiction on a point neither side actually contests, noticed only because the method looking for it does not yet know which points are live. Or it might be genuinely explicable on grounds the filer would readily supply if asked. Nothing about the anomaly itself, considered without reference to who created it or what their documented position made rational for them, says which of these explanations applies.

It is only when the three are read together that something more specific than "an anomaly exists somewhere" becomes available, and the specificity this combination produces is the paper's central methodological claim. A documentary gap — a contradiction or an unanswered point — found on a matter where the asserting professional's competence profile sits far enough from its technical content to make an innocent miss plausible, or, in the opposite configuration, close enough to make an innocent miss difficult to credit, and where that professional's documented position places them inside an incentive structure under which silence costs nothing while disclosure costs the position already filed, is no longer merely an anomaly. It is the specific, located instance of the structural problem Section 1 named: an assertion, or a silence standing in for one, that nobody with a reason to test it against the record ever did, sitting inside a record that already contained the means to test it.

The method does not claim to distinguish, from this combination alone, the lie from the indifference to truth Section 1 places alongside it, and this limit deserves to be restated plainly because it bears directly on how the method's output should be used. It claims only that the gap was never tested by anyone positioned to have no stake in the answer, and that the record, read whole rather than as a sequence of isolated filings, already contained what the test required. Whether the professional who created the gap did so through a failure of competence, a calculated bet on silence going unexamined, or some combination the bare documentary record cannot fully resolve, is a question the combination narrows the space of likely answers to, without purporting to answer it definitively. Sections 5 and 6 take up what follows once that combination is found in the specific, recurring instance this paper has used throughout as its clearest demonstration.

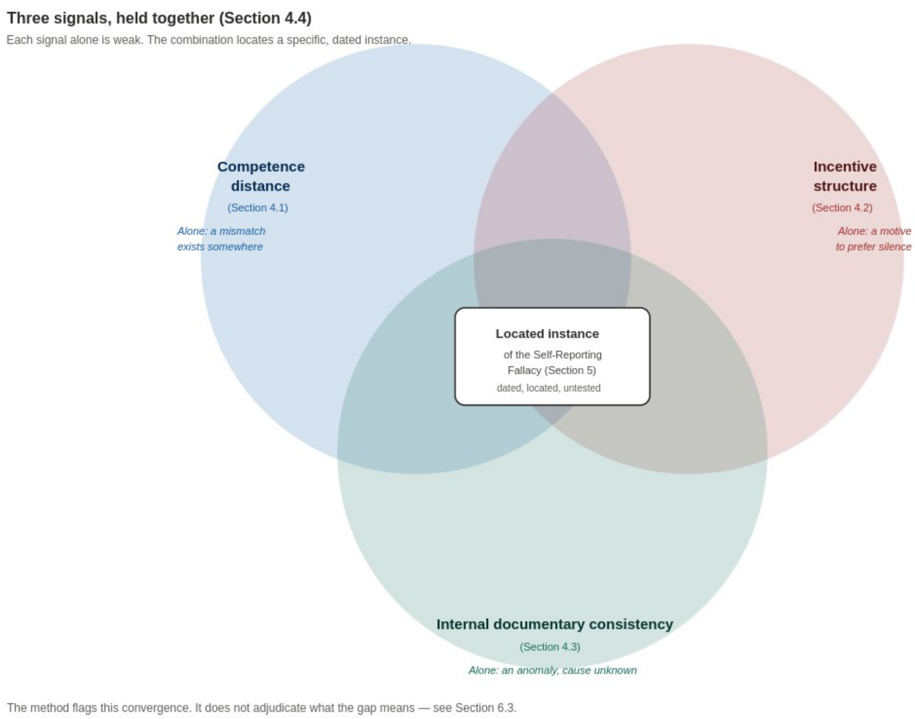

Figure 2. The three signals held together. Each alone produces only a weak finding — a mismatch, a motive, or an unexplained anomaly. Only where competence distance, incentive structure, and internal documentary consistency converge does a documentary gap become a located instance of the Self-Reporting Fallacy. See Section 4.4.

Section 5: The Moment Cover Runs Out

5.1 Zealous Advocacy and the Boundary That Requires Cover to Cross

A lawyer operating at the outer edge of zealous advocacy is, by definition, operating where the personal and professional risk of having misjudged the boundary is highest. No competent lawyer approaches that edge without believing, correctly or not, that something stands behind the position taken if it is later challenged: a defensible reading of an ambiguous record, a plausible characterisation of a document that could reasonably bear more than one meaning, an interpretation a court might, in good faith, accept even if it ultimately does not. This belief is not misconduct. It is the ordinary condition of advocacy conducted near its own limits, and the profession's reward structure, described in Section 1 and developed at length through the advocacy-ethics literature surveyed in Section 3, gives every reason to operate there rather than well inside a safer margin, because results obtained near the boundary are not received, billed, or remembered any differently from results obtained comfortably clear of it.

What this means, stated precisely, is that every filing made near that edge carries an implicit claim of cover — a claim that the position taken, however aggressive, remains a defensible reading of the record rather than an assertion the record itself already forecloses. The professional indemnity insurer standing behind the firm has, under current practice, no visibility into how much cover any specific filing actually has, for exactly the reasons Section 3.4 describes: the analytics available to that insurer report behavioural tendencies and historical outcomes, not the distance between any specific filing and the documentary record that either supports or undermines it. The insurer therefore prices the firm's risk as though the distance between every filing and the boundary were roughly uniform across the firm's book of work. It is not uniform. Some filings are made with real cover: a genuinely contestable reading of an ambiguous record, sustainable on the documents as they exist, defensible even if a court ultimately disagrees with it. Others are made with cover that has already run out — where the document being relied upon to support an assertion, read in full, does not support it, or affirmatively contradicts it — and the filing proceeds regardless, because the alternative is conceding a position already taken in an earlier filing in the same matter, with whatever cost to credibility and momentum that concession would carry.

5.2 The Moment a Filing's Own Evidence Contradicts Its Own Assertion

The moment cover runs out, on this account, is not the moment a court eventually notices, which Section 2 has already shown may never arrive within the life of the matter at all, given how rarely civil disputes reach a comprehensive judicial review of the kind that "the court would notice" implicitly assumes. It is the moment the contradiction is created — the moment a professional asserts a positive fact to a court, a regulator, or an opposing party while already holding, in the same file, material that contradicts the assertion. Before that moment, nothing distinguishes this matter from any other being argued near the edge of zealous advocacy, with whatever degree of genuine cover the underlying facts happen to supply. After it, something specific and nameable has occurred: an assertion has been made that the asserting professional's own record, read in full, does not support, and the gap between assertion and record now exists, dated and located, whether or not anyone reads far enough to find it.

This is precisely the instance Section 4 describes in combination, and it is worth tracing the combination through this specific event explicitly, because doing so shows why all three signals, rather than any one of them, are required to locate it with the specificity this paper claims for the method. The documentary test of Section 4.3 locates the gap itself — the specific sentence asserting a fact, and the specific document, already in the same file, that does not support it. The competence-distance signal of Section 4.1 bears on whether the gap more plausibly reflects a professional who did not recognise what their own file already showed — because the matter's technical content sat well outside their demonstrated area of practice — or one who did, because the matter sat squarely inside it. The incentive signal of Section 4.2 explains why, whichever of those two it is, the professional who created the gap has no structural reason to surface it themselves, and every structural reason, given the procedural defaults described in Section 2, to let the filing stand and proceed rather than withdraw it once made.

5.3 Why Client Risk and Firm Risk Diverge at That Exact Moment, Not Before

Until the moment described in Section 5.2 occurs, the represented party's litigation risk and the instructing firm's professional-indemnity risk are, in substance, the same number: the firm's only exposure is whatever the underlying dispute is worth, discounted by the probability that the firm's own representation succeeds. The two risks move together because they are, up to that point, mathematically the same quantity viewed from two different ledgers. Once the moment occurs, this stops being true, and it stops being true discontinuously rather than gradually. The firm has now created a second, independent exposure — a professional failure that exists as a matter of record regardless of how the underlying dispute eventually resolves, because a court that later rules in the client's favour on the merits has not erased the fact that the client's own representation advanced a position the representation's own file did not support.

This second exposure is not contingent on the underlying claim, and this independence is the precise feature that makes the divergence significant rather than merely an accounting curiosity. It survives a win on the merits as readily as a loss, because it is a fact about how the matter was conducted, not about who was ultimately right, and the question of who was ultimately right and the question of whether the representation along the way was itself sound are, after the moment in Section 5.2, genuinely separate questions with genuinely separate answers.

The two numbers having diverged, the larger of them is, in the ordinary case, no longer the original claim, and this is worth stating with the emphasis it deserves because it is the single most consequential implication this paper draws from the mechanism it describes. A professional-indemnity exposure created this way carries its own measure of damages — the cost of the position having been taken, the delay it caused, the credibility it cost the client before the court, the expense of unwinding or defending it once exposed — and that measure is set by the scale of the professional failure and its downstream consequences, not by the value of the dispute that gave rise to it. A modest underlying claim, conducted this way, can generate a professional-indemnity exposure that dwarfs it many times over, because the professional-liability damages are measured against an entirely different yardstick than the original dispute's value: they are measured against what the professional's own conduct cost the client and, where the gap has compounded into the cascade described in Section 5.4, against what it cost the opposing party as well. This is not an edge case requiring unusual facts. It is the direct and unremarkable consequence of the divergence itself, available wherever the moment in Section 5.2 has occurred, in a modest property dispute as readily as in a large commercial one.

Why client risk and firm risk diverge at the moment cover runs out (Section 5.3)

Before the moment, one number. After it, two — and the firm's is usually the larger one.

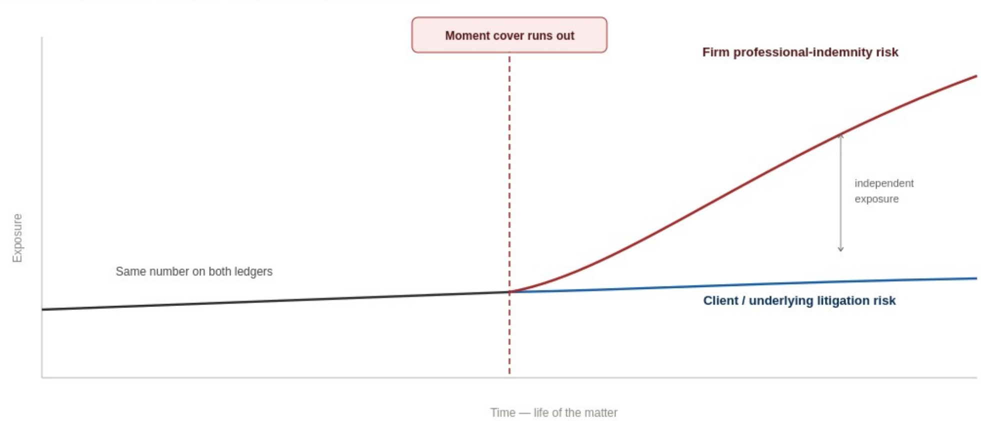

Illustrative, not to scale. The firm's exposure is set by the professional failure and its consequences (Section 5.3–5.4), not by the value of the underlying dispute — which is why it frequently exceeds the original claim many times over.

Figure 3. Why client risk and firm risk diverge at the moment cover runs out. Before the moment, the two are the same number on different ledgers. After it, the firm's professional-indemnity exposure rises independently of the underlying dispute's value, and is frequently the larger of the two. Illustrative, not to scale. See Section 5.3.

5.4 The Cascade That Follows, as Consequence Rather Than Centrepiece

Once the divergence in Section 5.3 exists, a further and largely mechanical consequence tends to follow, and it is worth stating at some length, as a consequence of the method rather than as the method's main subject, because the cascade is what ultimately determines the scale of exposure an insurer or a firm actually faces once the initial gap has been found.

The client whose own representation created the gap has, in the ordinary course, not been told that their position was advanced on a foundation their own lawyer's file did not support; disclosing this to the client would, after all, require the lawyer to recognise and admit the very gap the lawyer has no structural incentive to surface. If the gap surfaces in the proceedings — by opposing counsel, by a court, by a subsequent filing in a related forum, by any of the channels Section 2.4 describes as the extraordinary, multi-forum work required to force a buried fact back into a closed phase's record — it is the client's credibility before the court that is damaged first, not the lawyer's, because it is the client's position that has just been shown to rest on an unsupported assertion. The court does not, in the ordinary course, distinguish at that moment between a client who knew the assertion was unsupported and a client who relied entirely and reasonably on their own lawyer's judgment; the credibility cost lands on the party whose name is on the pleading, regardless of which of the two is in fact true.

The client's rational response, once this becomes apparent, is to seek new representation and to treat the original firm's conduct as a separate wrong from the underlying dispute: the prolongation of the matter, the credibility cost before the court, the expense of correcting the position, each attributable to the firm's own filing rather than to the underlying dispute the client originally retained the firm to handle. This is, in substance, a negligence claim distinct from anything the original dispute ever put in issue, and it exists independently of how that original dispute is ultimately resolved. The original claimant, meanwhile, now has a second and often more solvent target in the firm itself, reachable independently of how the underlying claim against the original defendant is resolved, because the firm's professional failure caused the claimant harm — delay, additional cost, a more difficult path to relief — distinct from whatever harm the original defendant caused.

And the firm's insurer, defending a single policyholder, finds itself defending against both the firm's original adversary and the firm's own former client at once, arising from the same underlying conduct, in a structural conflict the insurer did not create and cannot fully avoid once both claims exist: a defence strategy that minimises the firm's exposure to its former client may not be the strategy that minimises the firm's exposure to the original claimant, and vice versa, and the insurer is asked to fund a defence serving two potentially divergent interests at once. None of this requires unusual facts, exceptional venality, or a particularly sophisticated adversary. It requires only that the moment in Section 5.2 has occurred, and that someone, eventually, looks closely enough to find it — which is precisely the condition Section 6 specifies a method for satisfying without depending on the kind of extraordinary, accidental, or heroically resourced discovery process Section 2 describes as the only existing route to finding it today.

Section 6: Real-Time Burden, Not Eventual Judgment

6.1 The Formal Burden and the Evidentiary Reality, and the Lag Between Them

The formal burden of proof is a procedural allocation: a statement of which party must establish which fact, and to what standard, before a court will find in their favour. It shifts, where it shifts at all, at defined procedural moments — when an affirmative defence is raised, when a threshold showing has been made, when a court rules on a motion that reallocates what remains to be proven. It does not shift merely because a fact exists, somewhere in the record, that would justify shifting it. It shifts only when a court is asked to rule on it, on a record a court actually reviews, at a moment a party has chosen, or been forced, to bring the question forward. This is not a defect in the doctrine; a burden that shifted automatically, the instant some fact justifying the shift came into existence anywhere in a sprawling, multi-year, multi-forum record, without anyone needing to raise the question, would be a burden no party could ever rely on planning around, and the formal rule's dependence on someone actually asking the question is part of what gives litigation its procedural stability.

The evidentiary reality that Sections 4 and 5 describe is not the same thing as the formal burden, and the distance between the two is exactly the lag this paper has been describing since Section 1. A documentary contradiction of the kind Section 4.3 specifies, once it exists in the record, has already changed what a fully informed, disinterested reader would conclude about which party's account deserves the benefit of the doubt — long before any court is asked to rule on it, and frequently in a case that, per Section 2, never reaches a ruling at all. The formal burden sits wherever the pleadings and the applicable procedure put it, static until someone moves it. The evidentiary reality has already moved, the moment the contradiction was created, without waiting for anyone's motion. The lag between the two is where the harm Sections 1 and 2 describe actually accumulates: not in a court reaching the wrong conclusion on a fully argued record, which is a comparatively rare failure mode and one the appellate system exists, however imperfectly, to correct, but in a court, or far more often a settlement negotiation conducted in a court's shadow exactly as Section 2.1 describes, never being asked the question that the record already had the means to answer.

6.2 Flagging at the Moment of Creation, Not the Moment of Notice

What changes, if the three signals of Section 4 are held together continuously rather than applied once, is the moment at which the gap between formal burden and evidentiary reality becomes visible to somebody. The method this paper describes does not wait for a court to ask the question, and does not depend on any party's good fortune in stumbling across the contradiction through the kind of extraordinary, multi-forum effort Section 2.4 describes as the only currently reliable route to finding it. It flags the gap at the moment the contradiction or the conspicuous silence is created — at the filing stage, on the first round in which the relevant documents enter the same record, rather than at whatever later point, if any, a judge's limited and intermittent attention happens to fall on the right two pages at once, sequenced correctly, held in mind long enough to notice the connection between them.

This is the real-time correction to the lag described in Section 6.1: not a faster court, and not a different formal rule, and not any change whatsoever to how or when the formal burden of proof is permitted to shift under existing procedure. It is a reading position that notices the divergence between assertion and record at the moment it exists rather than at the moment, often arriving late or never, that someone with authority to act on it happens to look. The distinction matters because it locates the contribution correctly: this paper is not proposing a reform to the law of evidence or procedure. It is proposing a method for ensuring that the facts the existing law of evidence and procedure would already treat as significant, if anyone asked the right question at the right time, are actually surfaced in time for someone to ask.

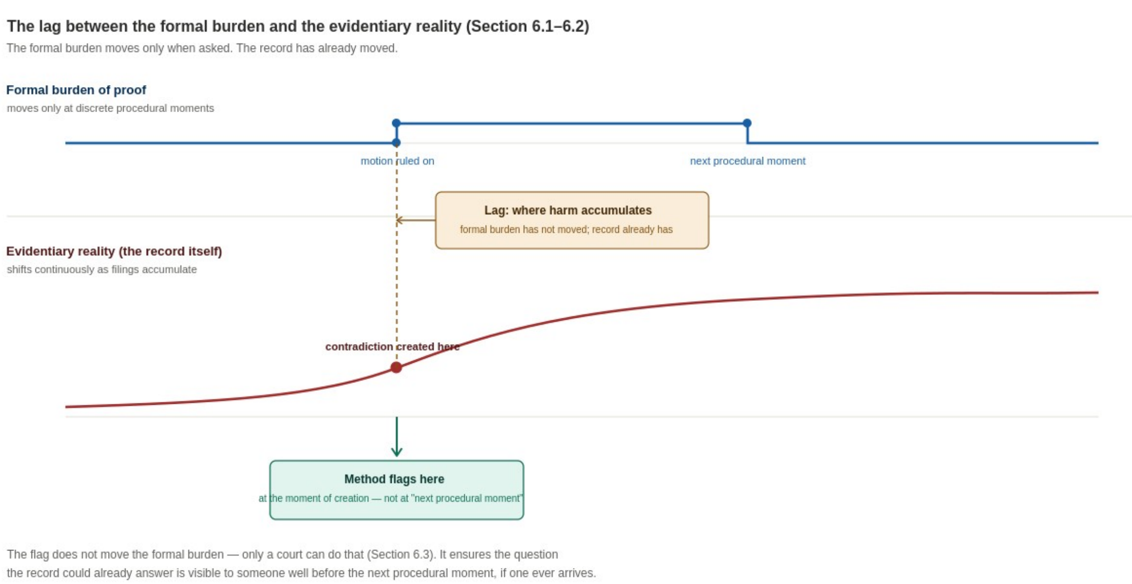

Figure 4. The lag between the formal burden of proof and the evidentiary reality. The formal burden moves only at discrete procedural moments; the documentary record shifts continuously as filings accumulate. The method flags a divergence at the moment it is created, rather than waiting for the next procedural moment, if one ever arrives. See Section 6.1–6.2.

6.3 Detection, Not Adjudication

This capability must be stated alongside a limit precise enough to prevent it being mistaken for something it is not, and this limit is, on this paper's view, at least as important to state clearly as the capability itself. The method flags. It does not adjudicate. It has no authority to strike a claim, to rule a position abandoned, or to determine that a contradiction is fatal rather than explicable — those are determinations for a court, a regulator, or a negotiating party to make, on a full record and with the benefit of argument from whoever made the assertion in the first place, including whatever explanation that party would offer for the apparent gap if given the chance to respond to it. A method that purported to do more than flag — that claimed to determine, on its own authority, that a given contradiction was decisive, or that a given silence was deliberate rather than inadvertent — would be claiming a kind of authority no documentary comparison, however rigorous, can actually supply, because the question of what a gap means, once found, depends on context, intent, and explanation that only a fuller process involving the party who created it can properly assess.

What the method supplies is the thing Section 1 argued no current participant in litigation is positioned to supply on their own: a continuously updated, disinterested account of where the formal position and the documentary record currently diverge, available to whoever reads the report — a client deciding whether to change counsel, an insurer deciding whether to intervene, an opposing party deciding whether a settlement offer reflects the file's actual strength — long before any of them would otherwise have reason to know. This is a meaningfully more modest claim than a claim to detect lies, or to find misconduct, or to predict outcomes, and the modesty of the claim is deliberate: it is exactly calibrated to what a documentary comparison can actually establish, no more and no less.

6.4 Why a Flag, Once Raised, Should Persist

A final feature of this reading position is worth stating because it corrects a specific failure mode Section 2 identified directly: the closing of a phase as a practical, if not always formal, foreclosure of facts that arrive too late to matter. A flag raised against one filing, in this method, does not lapse when the phase that produced the filing closes, and does not depend on the same reader, the same judge, or the same case officer remaining assigned to the matter for it to remain visible. It persists, attached to the record, available to be read against every subsequent filing in the same matter — including filings made months or years later, by parties or counsel who were not present when the original gap was created and who would otherwise have no occasion to discover it, exactly the situation Section 2.3 describes as the ordinary condition of a multi-phase proceeding handled across successive judicial officers and successive counsel.

A procedural decision made in ignorance of a fact that later surfaces remains, under current doctrine, difficult to reopen for exactly the reasons Section 2.2 describes: forfeiture and harmless-error doctrine are built to resist exactly this kind of reopening, for good institutional reasons, and nothing in this paper's proposal asks that resistance to be weakened. A persistent flag does not reopen anything by itself, and does not purport to override the doctrinal protections Section 2.2 describes as serving a genuine institutional purpose. It ensures, at minimum, that the fact making the earlier decision questionable is not lost a second time, simply because the file moved on before anyone read far enough back to find it again — which is a different and more modest achievement than overturning the doctrine, but one that, on this paper's argument, changes the practical balance of who is able to act on a gap considerably, even while leaving the formal doctrine entirely undisturbed.

Section 7: Every Case, All the Time

7.1 From Single-File Diagnosis to Standing Portfolio Surveillance

Nothing about the method specified in Sections 4 through 6 is confined to matters that already look troubled, and this is, on this paper's assessment, the point most likely to be underestimated by an institution encountering the method for the first time. The Self-Reporting Fallacy, as Section 1 argues, is structural rather than exceptional: it operates, to some degree, in every matter where a represented party's filings are read by people whose position gives them a stake in the answer those filings supply, which is to say, in every matter without exception, because the structural conditions Section 1.3 describes — the client's dependence on trusting counsel, the lawyer's commitment to their own filing's coherence, the insurer's reliance on summary rather than record, the court's exposure only to what is argued — are present in every represented matter regardless of how contentious or straightforward that matter currently appears.

A reading position immune to that fallacy is therefore not a specialist instrument reserved for the file that already shows signs of trouble, in the way a forensic accountant might be called in only once a set of accounts already looks suspicious. Its natural unit of deployment is not the single difficult case but the entire portfolio, run continuously, with each new filing in each open matter checked against the same standing discipline as it arrives, in the same way a continuous audit function operates across an entire set of accounts rather than being deployed only once a specific transaction has already attracted suspicion. The difference between these two models of deployment — instrument of last resort versus standing discipline — is, on this paper's argument, the difference between a tool that catches a fraction of the gaps Section 1 describes, in the matters where someone happened to suspect something was wrong, and a tool that catches the much larger fraction that nobody had any independent reason to suspect at all.

7.2 The Management-Level Glance

The practical consequence of continuous deployment is a form of output no existing litigation- or claims-analytics platform surveyed in Section 3 currently produces: a standing, portfolio-wide account of where, across every open matter a firm or an insurer holds, the documentary record currently diverges from what is being asserted about it, updated as each new filing arrives rather than reconstructed from memory at the point a matter finally demands attention. A management team, or a claims-oversight function, reviewing such an account does not need to read the underlying filings in every matter to know which of them currently carry an unresolved flag, what kind of gap the flag represents — a direct contradiction of the kind Section 4.3's first test locates, or a conspicuous unaddressed point of the kind its second test locates — and how long it has been sitting in the record unaddressed.

A portfolio that might otherwise require weeks of manual file review to assess at this level of specificity becomes reviewable, at the level the question actually requires answering — not who will win, which is the question the existing outcome-prediction tools surveyed in Section 3.4 already answer reasonably well, but where the record and the argument have already parted company, which none of them answer at all — in well under an hour. This is not a claim that the underlying matters become simpler, or that the substantive work of litigating each of them is in any way reduced. It is a claim about where attention should be allocated first, across a portfolio too large for any single reviewer to read in full, with any regularity, by hand.

7.3 A Signal Layer, Not a Status Layer

This is a different deliverable from a case-management dashboard, and the difference is worth stating precisely because the two are easily conflated, particularly by an audience already familiar with the case-management and claims-tracking systems most firms and insurers already operate. A status layer reports where a matter stands procedurally: what phase it is in, what is due next, what has been spent so far, who is assigned to it, when the next hearing or filing deadline falls. This is valuable information, and no part of this paper's argument suggests it is not. But a status layer, however well built, was never designed to surface something a signal layer is built specifically to surface: whether the substance of what is being argued, on either side, currently survives contact with the record that already exists. A portfolio can be perfectly current on status — every deadline tracked, every phase correctly logged, every dollar of spend accounted for — and entirely blind on signal, with no flag anywhere indicating that a specific filing in a specific matter has quietly stopped engaging with the point the matter actually turns on. It is the signal, on the argument this paper has made throughout, that determines where the real exposure in a portfolio actually sits, and a portfolio's status metrics, however complete, give no indication of where that exposure is concentrated.

7.4 Why Every Case Benefits, Not Only the Hard Ones

It follows from Section 1 that the value of this reading position does not scale with how troubled a matter already appears to be, and this point is worth drawing out at some length because it runs against an intuition that might otherwise lead an institution to deploy the method selectively, reserving it for matters already flagged as difficult by other means. A matter that looks straightforward today is, on this paper's argument, exactly as likely to contain an unexamined gap as one that already looks contested, because the fallacy that produces the gap operates on the structure of who is reading the file — the incentive landscape described in Section 1.3, which is present in every represented matter — not on the apparent difficulty of the dispute itself. Indeed, there is some reason to expect the opposite pattern to hold in practice: a matter that already looks contentious tends to attract more scrutiny from more parties, simply because everyone involved already suspects something is being fought over closely, while a matter that looks routine attracts correspondingly less scrutiny precisely because nobody has any particular reason to look closely, which is exactly the condition under which a gap of the kind Section 4 describes is least likely to be found by ordinary means and most likely to do its work undetected.

The case for running every open matter through this discipline, continuously, rather than reserving it for the files that already look difficult, is therefore not a case for extra caution on hard cases, which most institutions already apply by other means. It is a case for treating the discipline itself as a baseline, on the same logic that a portfolio is audited for financial irregularity as a matter of course, across every account, rather than only once a specific account has already given someone reason to suspect one — because the entire value of a baseline audit, in either setting, lies precisely in its application to the accounts nobody yet suspects, since the accounts already under suspicion are, almost by definition, already receiving the scrutiny a baseline practice exists to extend to everything else.

7.5 The Mandatory Shared-Window Condition of Coverage

The demand-side obstacle that Section 8.2 identifies — that the firm best positioned to run this method on its own filings is also the actor with the most mixed incentive to do so — dissolves once the method is understood as an instrument of the insurer rather than of the firm. Insurance-law scholarship has established, over more than two decades, that liability insurers already function as private regulators of the professions they cover: by pricing and monitoring risk to protect their own solvency, they impose a discipline on professional conduct that no court or bar association directly supplies. Baker and Swedloff give this phenomenon its sharpest name in their account of insurers as bumblebees in the garden of law firm norms — insurers do not set out to reform professional conduct, any more than a bee sets out to pollinate, but the by-product of their self-interested monitoring is precisely that effect.

The proposal here asks nothing of insurers that is alien to a role legal scholarship already recognises them as occupying; it specifies only the instrument through which that role could be exercised at the moment of filing, rather than only retrospectively, through claims histories and premium adjustments made after a loss has already crystallised. Framed this way, the method's natural deployment is not a service a firm elects to purchase, but a shared, mandatory monitoring window written into the terms of coverage: the insurer requires the instrument to be run across the insured firm's open matters as a condition of the policy, with the resulting signal visible to both insurer and firm.

Two consequences follow, and each removes an obstacle the preceding sections identified. Because the choice to audit one's own filings is no longer the firm's to make, the demand-side conflict of interest set out in Section 8.2 no longer arises: the firm is not being asked to volunteer scrutiny against its own interest, because the scrutiny is a condition of the coverage it has already bought. And because the insurer's exposure is reduced, not increased, by the early discovery of a gap in its own insured's filing — for the reasons Section 8.3 develops — the instrument aligns with the settled interest of the one actor in the entire structure whose financial position depends on getting the portfolio-level picture right, and who is therefore structurally positioned to insist on its use.

Section 8: Naming the Gap, and the Case for Building Across It

8.1 What This Category of Tool Would Need to Do

State plainly what Sections 1 through 6 have been building toward, stripped of formula: a bullshit detector for litigation would need to ingest every filing in a matter as it is produced, by every party and in every forum touching the same underlying facts, and hold the entire accumulating record in active comparison against itself, rather than reading each new document in isolation against the memory of whichever individual happens to be handling the file that week. It would need to flag the moment a positive factual claim, made by any party, is contradicted by material already present elsewhere in that same accumulating record — including material the claiming party itself filed, in this proceeding or in a related one, at an earlier point the system has no obligation to forget. It would need to flag, separately, the moment a filing conspicuously fails to engage with a point the opposing side has specifically placed in issue, applying the second test of Section 4.3 with the same consistency as the first. It would need to make both kinds of flag persist: not lapse when the phase that produced them closes, not require the original reader to still be assigned to the file, not depend on the next judge, the next associate, or the next adjuster happening to read back far enough to rediscover them. And it would need to do all of this without adjudicating anything — without ruling on whether the contradiction or the silence it has found is fatal, excusable, or merely apparent, because that judgment belongs to a court, a regulator, or a negotiating party, not to the instrument that surfaced the material for them to judge.

This is, in plain language, the entire specification. Nothing about it requires resolving a question of law, and nothing about it requires the instrument to take a position on the underlying merits of any dispute it is applied to. It requires only sustained, exhaustive, cross-document memory applied with consistent rigour for the life of a matter — a capability no human role in litigation is currently paid to provide, for the reasons Sections 1 and 2 set out at length, and a capability that has, until recently, been computationally impractical to provide at the scale a real multi-year, multi-forum matter actually generates, given the volume of filings, exhibits, and correspondence a contested matter of any significant duration accumulates.

8.2 Why the Existing Landscape Was Never Going to Build This

The litigation-analytics and claims-analytics industries surveyed in Section 3 are neither small nor unsophisticated. Casualty and professional-indemnity insurers now draw on platforms that model portfolio-level catastrophe exposure from hundreds of millions of structured corporate records, on litigation-analytics tools that report how a specific assigned judge has ruled on similar motions in the past and how often a specific opposing counsel settles before trial with reported accuracy in the seventy to eighty-five percent range across well-studied case types, and on claims-analytics systems that apply natural-language processing to adjuster notes to extract early signals of severity and litigation propensity. Industry commentary describing this landscape does not treat real-time, cross-document contradiction detection as a capability the sector is working toward and has not yet reached; it does not describe the capability at all. The published trade literature is directed almost entirely at prediction from precedent and at severity-scoring of claims already open — forecasting outcomes and costs from historical patterns — and the distinct operation this paper specifies, reading a specific filing against its own attachments and against what the record conspicuously fails to address, appears nowhere in it, not even as a stated aspiration. The absence is not the industry declining a known next step; it is the industry not yet having named the step.

The reason is structural, not a matter of any vendor's competence, and it is worth setting out both halves of the explanation in full, because each half independently would have been sufficient to prevent this category of tool from emerging through ordinary commercial development, and together they make the absence considerably less surprising than it might first appear.

The first half is a problem of available training signal. Every one of the tools described in Section 3 answers a question with a clean, trainable ground truth: did this judge grant the motion; did this case settle; did this claim type, in this venue, with this attorney, historically escalate; does this clause deviate from this organisation's standard template. Those questions can be answered from large historical datasets with labels that already exist, because courts, insurers, and law firms already record outcomes, settlements, and template deviations as a matter of ordinary business practice. The question this paper is concerned with — does this specific filing's own assertion survive contact with the rest of this specific file — has no comparable historical label to train against, because nobody has ever systematically recorded which filings turned out, on later and usually accidental discovery, to have been internally inconsistent at the time they were made. The absence of that label is not incidental, and it is worth being precise about why. It is a direct consequence of the Self-Reporting Fallacy itself: the inconsistencies this paper is concerned with are, by hypothesis, the ones nobody positioned to label them ever surfaced, because surfacing them would have required exactly the disinterested reading position Section 1 argues does not currently exist inside the ordinary course of litigation. A field cannot easily build a supervised model to detect a category of event it has no recorded instances of, precisely because the institutions around litigation have never been structured to record them — a closed loop in which the absence of detection and the absence of training data for detection reinforce each other.

The second half is a demand-side reason, and it deserves to be stated as plainly as the first, because it explains why the gap is not merely a matter of waiting for the technology to mature. The actor best positioned to commission a tool that reads an entire file for its own internal contradictions is, in the ordinary commercial structure of the industry, the instructing firm — and the instructing firm is not, in every matter, an actor whose interests are served by such a tool being run on its own filings. A firm operating well inside the bounds of ordinary practice has every reason to want this capability turned on its opponents, where it would supply exactly the kind of competitive advantage the analytics platforms in Section 3.4 already supply in adjacent forms, and considerably more mixed reasons to want it turned on itself, where it would surface precisely the category of gap that, once surfaced, creates the professional-liability exposure described in Section 5. This is not an accusation, and it does not require attributing any bad faith whatsoever to any firm or any vendor currently operating in this market; it follows directly from the incentive structure Section 1 describes, applied to the purchasing decision itself rather than to any individual filing. It explains, without requiring any bad faith on any vendor's part, why a category of tool whose central use case is auditing one's own side's filings for the very gaps Section 1 names was never going to emerge from ordinary supply and demand inside the industry that would, on the analysis of Section 5, benefit most from using it on others while having correspondingly mixed reasons to commission its use on themselves.

8.3 Why an Insurer Is the Buyer the Existing Gap Has Been Waiting For

The professional-indemnity insurance industry sits in a structurally different position from either the firm or the client, and the position matters because it removes both of the obstacles described in Section 8.2 simultaneously. An insurer is not the author of any side's filings, and has no comparable conflict of interest in having those filings audited for internal consistency; if anything, the insurer's financial exposure is reduced, not threatened, by discovering a gap in its own insured's filing early, before the cascade described in Section 5.4 has had time to develop into the multi-track exposure that cascade describes, rather than after the claim has already crystallised into litigation against the firm itself. And an insurer, uniquely among the actors in this picture, is exposed not to one matter but to an entire portfolio of them simultaneously, which converts a capability that might look like a marginal improvement on any single file into a substantial, aggregated reduction in reserve uncertainty across hundreds of open matters at once — precisely the portfolio-level value Section 7 describes, applied to the single actor in the entire structure whose financial position depends on getting the portfolio-level picture right rather than any single matter's outcome.

The Ames & Gough annual survey — the industry's own benchmark survey of the largest professional-liability insurers, covering firms in the Am Law 100 — found that by 2022 all but one of eleven surveyed insurers had paid a single claim in excess of $50 million within the prior two years, and four of the eleven had paid a single claim in excess of $300 million. These are not portfolio aggregates but individual claims, of a size that makes even a marginal improvement in early detection commercially material rather than merely theoretically desirable — and precisely the category of loss the professional-indemnity literature already attributes, in its own published analyses, substantially to case complexity and conflicts of interest, the very conditions under which the gap this paper names is hardest for any single human reader to catch.

The professional-indemnity claims literature already published by major commercial insurers documents, in considerable and granular detail, the scale and the drivers of large professional-indemnity losses across legal, construction, and financial-services lines, identifying the sectors and the categories of failure most responsible for the largest claims the industry currently faces. What that same literature does not yet describe is any instrument capable of reading the filings behind those losses for the specific signal this paper has named — a gap in the published self-description of an industry that is otherwise unusually candid, in its own trade publications, about the scale and the categories of its exposure. The gap, on the evidence surveyed in this section, is not a gap in insurer appetite. The data above shows an industry already investing heavily in analytics, already documenting its own exposure in granular detail, yet without any instrument that reads the filings behind those losses for the signal this paper names. The gap is a gap in what has been built to meet that appetite, for the structural reasons Section 8.2 sets out.

8.4 The Prototype as Existence Proof

A prototype instrument, built by this author and described at a conceptual level throughout this paper, demonstrates that the capability specified in Section 8.1 is achievable with methods available today, not merely a theoretical specification awaiting some future advance in computational capability or some future improvement in the underlying technology. The instrument ingests the filings in a matter as they are produced, maintains the accumulating record across successive submissions rather than treating each one in isolation, and flags the category of internal contradiction this paper has described, including both tests specified in Section 4.3 and the specific recurring instance — a professional's own filed evidence contradicting the professional's own asserted position — worked through in Sections 4 through 6.

Nothing about the instrument's design required a conceptual breakthrough beyond the combination this paper specifies: the capability follows directly from holding the three signals of Section 4 together, continuously, over a file's full life, which is precisely the reading position no existing tool surveyed in Sections 8.2 and 8.3 was built to occupy and no human role in litigation, for the structural reasons Section 1 develops, is paid to sustain. That the instrument exists, and that it produces this category of output on real filings, is offered here not as the paper's subject but as the paper's evidence: the gap this paper has spent its length naming is not merely real, but already closing, and the closing of it does not depend on any advance this paper has not already described in full.

Section 9: Objections Considered

Three objections deserve direct engagement, in addition to the limits already stated within the relevant sections themselves.

The first is that a method capable of flagging the category of event described in Section 5 risks generating false positives — flagging as suspicious a discrepancy that is, on examination, an innocent slip, a genuinely ambiguous record admitting more than one fair reading, or an omission with an ordinary, non-strategic explanation. This risk is real and is not eliminated by anything this paper proposes, and it would be a mistake for this paper to suggest otherwise. The answer is the limit stated plainly in Section 6.3: the method flags; it does not adjudicate. A flagged gap is a prompt for a human reader — a client, an insurer, an opposing party, eventually a court — to ask a question that the record otherwise gives no one reason to ask. It is not a finding that the question's answer is damning, and it should not be treated, by any institution using the method, as equivalent to a finding of fault. The discipline this method applies is more rigorous than ordinary sequential reading precisely because it holds more material in view at once across a longer span of the record than any individual human reader practically can, but rigour in surfacing a candidate gap is not the same as rigour in resolving what the gap means, and this paper has not claimed the second. An institution deploying this method responsibly will build a review step into its own process, in which a flagged gap is examined by a competent human reader before any consequence attaches to it — a step entirely consistent with, and indeed required by, the detection-not-adjudication limit this paper has insisted on throughout.

The second objection concerns this author's own position. A researcher who develops several interlocking papers and a diagnostic instrument alongside them invites the suspicion of having reasoned backward from conclusions he had an interest in reaching. The paper's answer is not a protestation of neutrality but a matter of what the argument is built from: every claim rests on settled doctrine and public literature — Frankfurt's distinction, the advocacy-ethics critique across Luban, Simon, and their respondents, the procedural doctrines of forfeiture and harmless error, and the analytics landscape surveyed in Section 3 — extended to a named gap, and on no dispute's facts, which the paper has used as evidence nowhere. This is deliberate, and it follows the paper's own principle: the method works precisely because it reads a record without regard to who wrote it or why, and the argument for the method is built to be read the same way. Whether the gap is real, and whether the three signals close it, are claims a reader can test against the cited sources and the structure of the argument alone, taking nothing on the author's word but the accuracy of the citations. The author's interest in the answer is, on the paper's own logic, irrelevant to whether the answer is correct — which is the entire point.

The third objection is more pointed, and deserves to be answered without evasion: if this category of tool is as valuable as this paper argues, naming it and explaining how it works in public risks accelerating its replication by better-resourced competitors, at the cost of whatever advantage having specified it first might otherwise carry. This paper's position is that the conceptual contribution — naming the Self-Reporting Fallacy, specifying the three signals that a reading position immune to it must hold together, and demonstrating that the combination is achievable with present methods — is worth more, to the field and to this author's own standing in it, made and dated in public than held back. What this paper has not done, deliberately and consistently throughout, is describe how any of this is actually implemented: no architecture, no extraction method, no scoring formula, nothing a competitor could build from beyond the conceptual specification any sufficiently motivated team could in principle have arrived at independently, given enough time and the right starting question. The distinction matters, and it is the distinction this paper has tried to honour throughout, including in Section 8.1's deliberately formula-free statement of what the capability requires. A method can be named and argued for in full, with its significance fully defended, without disclosing how it is built.

It is also worth noting that the gap this paper names has been independently confirmed, at the structural level, by the field this paper most directly addresses. Baker and Swedloff's study of liability insurer data identifies, as one of its own explicit findings, that the data on lawyer malpractice is fragmented and incomplete precisely because no mandatory reporting regime exists — insurers and bar associations hold separate pieces of the picture, but nothing forces a complete, structured, cause-and-outcome account across the industry. That finding, arrived at independently by insurance-law scholars studying a different question, is precisely what the Self-Reporting Fallacy predicts: a framework that prices the assembling of complete information at a cost, and provides no mechanism to compel it, will not produce complete information, at any level of the institutional stack.

Section 10: Conclusion

This paper set out to name a single blindness and ends by locating it inside something larger than any single paper. Frankfurt's distinction between the lie and the assertion made in indifference to its truth gave the blindness its name; the argument gave it its cause. Every actor positioned to read a legal file depends, in some specific and identifiable way, on the answer that file gives, and so none of them reads it whole, against itself, from a position with nothing at stake in the result. The contradiction sits in the record from the moment it is created. What is missing is never the evidence. What is missing is a reader with no reason to look away.

Underneath that blindness sits a mechanism, and the five papers this one completes are best read together as its anatomy. A legal system of any scale cannot verify everything at the point of entry, and so it grants a set of presumptions for free. The court-appointed expert is presumed impartial. The lawyer is presumed expert by virtue of admission, rather than by any demonstrated fit to the matter in hand. Everyone the dispute actually reaches is presumed to be the parties formally before the court, and no others. Silence about a damaging fact is presumed to be no wrong at all. And the claim, once filed, is presumed to say what it purports to say — treated as settled ground unless the other side is willing to spend, at great cost and with counsel in the room, to prove otherwise. Each presumption looks like an economy. Each is in fact a transfer: the cost the presumption spares the system does not disappear, it is moved onto whoever must later pay to rebut it — if they can afford to, and if they ever learn there was anything to rebut. *Network Proximity and the Appearance of Impartiality* tested the first of these presumptions, *Qualification and Expertise* the second, *The Doctrine of the Affected Party* the third, *The Manufacturing of Bad Faith* the fourth, and this paper the fifth. At each one, the thing presumed turned out not to be reliably there.

That is a strong claim to make five times, and the honest thing to say about it is how little effort, relative to the age of the target, it took to make. The law is among the most examined institutions in human history, three centuries into its modern form, and a single outside investigator, using tools that did not exist five years ago and reading the record whole, arrived at the same structural gap on every layer he looked at. That says almost nothing about the investigator and almost everything about the gap. These were not findings buried by difficulty. They were findings unlit by incentive — in plain sight, and looked at by no one, because no one inside the system stood to gain from raising the light.

The reason no one raised it is not conspiracy; it is arithmetic. A system whose participants are paid to conduct a dispute rather than to end it will, over enough time and with no one intending it, grow better at conducting disputes than at ending them. The process lengthens, the resolution recedes, and the distance between the two is revenue. Resolution is the one outcome such a system is structurally unequipped to sell, because resolution is where the billing stops. A matter the record could settle in an afternoon becomes the occasion for years of motion, and the argument is conducted — at length, and at great cost — over everything except the one question whose answer would end it. This is not the accusation that lawyers manufacture delay. It is the observation that a structure which meters the process and not the outcome will optimise the process, and that what has been built, measured by what it rewards rather than by what it was meant to do, now resembles an apparatus for extending disputes at least as much as one for resolving them.

This is why the catalogue does not end in a call for reform. Reform adjusts a system's rules from within, and these gaps are not lapses in the rules — they are the rules working as their incentives specify. A design does not petition itself to stop being what it is. Nor will the correction come from the technology now reorganising legal work at speed. Legal artificial intelligence is being built in volume and to genuine effect, but it optimises, without exception, the variable the profession already optimises — the win, and the faster, cheaper, better-defended pursuit of it — because that is what its buyers pay it to do, and its buyers are the professionals, not the parties they act upon. The one instrument this series specifies, which reads the professional rather than for him and shortens a dispute rather than sustaining it, is precisely the instrument the market has no reason to build. For the professional, the length of a dispute is not a cost to be minimised. It is the product.

What changes the picture is not reform, and not the profession's own tools, but the collapse of the single condition every presumption above depended on: opacity. Each of those free grants held not because it had been examined and found sound, but because examining it required reading an entire record from a position with no stake in the reading — the one vantage the system does not contain, and for three centuries could not be supplied at scale. It can be supplied now. Artificial intelligence has, for the first time, made the whole of a record legible to a reader with nothing at stake in its conclusions, and a presumption that survived only because no such reader existed cannot survive that reader's arrival. This is what makes the moment one of disruption rather than of reform. This series has named the actors those presumptions protected — the expert never measured against the parties, the professional never measured against the matter, the advocate never measured against his own file. The authority each of them wore was issued at the door, not earned at it, and the entire cost of discovering it unearned has fallen on the party it was worn against. Naming them is not an accusation. It is the withdrawal of a presumption, and a presumption, once withdrawn, is very difficult to grant again.

One party stands to gain from that withdrawal rather than to lose by it, and it has been in the argument since the moment cover runs out was first defined. The professional-indemnity insurer is the only actor in the whole structure aligned with the client on the axis the profession and its tooling ignore — resolution over operation, real exposure over billed effort — and set against the professional on it. It is at once the firm's customer and, through the pricing and monitoring of risk, its private regulator; and it is the party that pays when advocacy conducted at the edge finally fails. It has been reading a summary while the record it never sees already held the answer. The instrument this paper specifies does not arm the insurer against the lawyer; it lets the insurer see, and price, an exposure it has been carrying blind. That the exposure was always in the record, and that reading the record whole is no longer beyond reach, is the entire practical claim.

The insurer, though, is the pragmatic destination, not the final one. If the gaps are the design, the complete answer is not a better complaint, nor a sharper audit fastened to a system built to produce them, but a better design — a venue in which impartial expertise, timely disclosure, and resolution are native conditions of the structure rather than victories wrung from a process that rewards their absence. Specifying that venue has been the quiet purpose beneath all five papers, each reverse-engineering, from a different failure, one requirement such a design would have to meet; the author has been building it, under KYC.co, and it is named here once, and only to keep the diagnosis honest.

The legal profession is among the last of the great professional-services sectors that artificial intelligence has not yet reorganised, and when the reorganisation comes it will run along the incentive lines this series has mapped — not because a machine chooses it, but because the incentives were always there, waiting only for something able to read the whole record at once to make them legible. That reader has arrived.

What remains is not a decision by a single actor. What remains is a race, and it has already started. Every live portfolio, every contested claim, every file in which an expert was appointed without proximity measurement, a lawyer engaged beyond the edge of their domain, a silence protected by a presumption that did not ask what the silence was protecting — each of those files now sits as a candidate for this reading. The parties on both sides of every such dispute now have access to the same lens. The question is not whether the gaps will be found. The question is who finds them first.

The independent expert who sat too close to the appointing party. The professional who filed an assertion their own attachments already contradicted.

The presumption that transferred cost onto the wrong party because no one was positioned to read the whole record at once. The question that went unanswered through six phases of pleading and was never noticed because no single reader held all six phases at the same time. These are not historical curiosities. They are present in live files, in portfolios of active claims, in cases currently before courts whose participants have no idea the record they have been building contains what it contains.

If you are the party who finds that gap — what do you do with it? The law did not provide the sword. The record always held it. The tools hand it to whoever arrives first with the file and the willingness to read it whole. That is what makes the moment a race and not a reform. Reform waits for consensus. A race starts the moment anyone with a case understands what is now legible, and what reading it might be worth.

Bibliography

Primary Philosophical Source

Frankfurt, H.G. (2005) On Bullshit. Princeton: Princeton University Press.

Legal Ethics and the Standard Conception

Luban, D. (1988) Lawyers and Justice: An Ethical Study. Princeton: Princeton University Press.

Luban, D. (1990) 'Freedom and Constraint in Legal Ethics: Some Mid-Course Corrections to Lawyers and Justice,' Maryland Law Review, 49, pp. 424–452.

Luban, D. (2007) Legal Ethics and Human Dignity. Cambridge: Cambridge University Press.

Simon, W.H. (1998) The Practice of Justice: A Theory of Lawyers' Ethics. Cambridge, MA: Harvard University Press.

West, R. (1999) 'The Zealous Advocacy of Justice in a Less Than Ideal Legal World,' Stanford Law Review, 51, pp. 973–1003.

Spaulding v. Zimmerman, 263 Minn. 346, 116 N.W.2d 704 (1962).

Psychology of Escalation and Commitment

Brockner, J. and Rubin, J.Z. (1985) Entrapment in Escalating Conflicts: A Social Psychological Analysis. New York: Springer-Verlag.

Procedural Doctrine

United States v. Olano, 507 U.S. 725 (1993).

Chapman v. California, 386 U.S. 18 (1967).

Federal Rule of Criminal Procedure 52(a)–(b).

Pierce, R.J. (2014) 'Making Sense of Procedural Injury,' Administrative Law Review.

Author's Prior and Companion 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.

Carroll, S. (2026) Qualification and Expertise: The Detectable Signals of Professional Competence Distance. SSRN Working Paper No. 6996898.

Carroll, S. (2026) The Doctrine of the Affected Party: Realigning Standing, Notice, and the Burden of Proof in Interdependent Urban Structures. SSRN Working Paper No. 6972581.

Carroll, S. (2026) The Manufacturing of Bad Faith: Mechanism Design and the Doctrine of the Affected Party. SSRN Working Paper No. 6972499.

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

KYC.co (2026) Dispute Engine. Prototype tool. Available at: https://kyc.co/dispute-engine

Litigation and Claims Analytics

LexisNexis (2026) Lex Machina Legal Analytics. Product documentation. Available at: https://www.lexisnexis.com/en-us/products/lex-machina.page

Bloomberg Law (2026) Litigation Analytics. Product documentation.

Luminance (2026) Legal-Grade AI Contract Management Platform. Product documentation. Available at: https://www.luminance.com/

Allianz Commercial / AGCS (2023) Professional Indemnity Insurance Claims: Report and Insights.

Genre (2025) 'Litigation Analytics: Turning Data into a Competitive Advantage for Casualty Insurers.'

Ames & Gough (2022) Lawyers' Professional Liability Claims Trends: 2022. Survey of leading lawyers' professional-liability insurers. [Exact report title, year, and the cited $50m / $300m figures to be verified before submission.]

Liability Insurance as Private Regulation

Baker, T. and Swedloff, R. (2013) 'Regulation by Liability Insurance: From Auto to Lawyers Professional Liability.' UCLA Law Review, 60. [Pagination to be verified.]

Baker, T. and Swedloff, R. (2015) 'Liability Insurer Data as a Window on Lawyers' Professional Liability.' UC Irvine Law Review, 5. [Pagination to be verified.]

Baker, T. and Swedloff, R. (2017) 'Insurers as Bumblebees in the Garden of Law Firm Norms.' [Full publication details to be verified.]

Multi-Agent and Computational Modelling of Litigation

LegalSim (2025) Multi-Agent Simulation of Legal Systems for Discovering Procedural Exploits. [Verify author (recovered as Badhe), venue (arXiv vs. NLLP / EMNLP 2025 workshop), and pagination before submission.]