WORKING PAPER WP-004 · CRITERICA INTELLIGENCE RESEARCH · NOVEMBER 2026

When a Document Is Missing, the Number Does Not Go Up

Criterica Intelligence Research

Corresponding: Jack Cosentino, founder · research@criterica.com

Published 8 November 2026 Updated 8 November 2026 Version 1.0

11-minute read

ABSTRACT

A diligence file is never complete, and the party assembling it chooses what goes in. A valuation method that can rise when evidence is withheld rewards the wrong behaviour. This paper states the rule every Criterica figure obeys: removing information never raises a certified figure, a matter on hold carries no figure at all, and one exact calculation sets every number. It works an illustrative example in which a document is removed and the band widens while the certified figure falls, and sets out what the rule means for a funder or buyer reading a data room.

IN ONE PARAGRAPH

A valuation that can rise when a document is removed pays people to withhold evidence. Criterica figures obey one rule: removing information never raises a certified figure. A matter missing an essential document carries no figure at all, only a list of what releases it, and one exact calculation sets every number. A widening band shows where the uncertainty lives, and the gap between two figures is the value of the missing document.

SERIES
Working Paper WP-004

VERSION
1.0

PUBLISHED
8 November 2026

STATUS
Published

AUTHORS

Criterica Intelligence Research

Corresponding author: Jack Cosentino, founder

WHO THIS IS FOR AND THE DECISION IT INFORMS

Reader. Funders, portfolio buyers, lenders and insurers who rely on a file assembled by someone else.

Decision. How much to advance, pay or reserve against a legal asset when the file is incomplete, and what to ask for before deciding.

Cost of getting it wrong. A method that rewards a thinner file pays more for less evidence. The loss arrives later, concentrated in the files that were thin for a reason.

01

Every diligence file is curated by someone with an interest in the answer

A funder evaluating a case, a buyer pricing a book of receivables and a lender sizing a facility all read a file assembled by the other side. The law firm chooses which treatment records to send, the originator chooses which matters go into the tape, and the seller chooses which documents reach the data room. None of that is improper. It is simply how files arrive: incomplete, and incomplete in ways chosen by people who know what the missing documents say.

That creates a specific risk for any valuation method. If a figure can rise when a document is removed, the party assembling the file has a reason to remove it, and the method pays more for

less evidence. The files that benefit most are exactly the ones a reader should trust least: the matters where the missing record would have lowered the number.

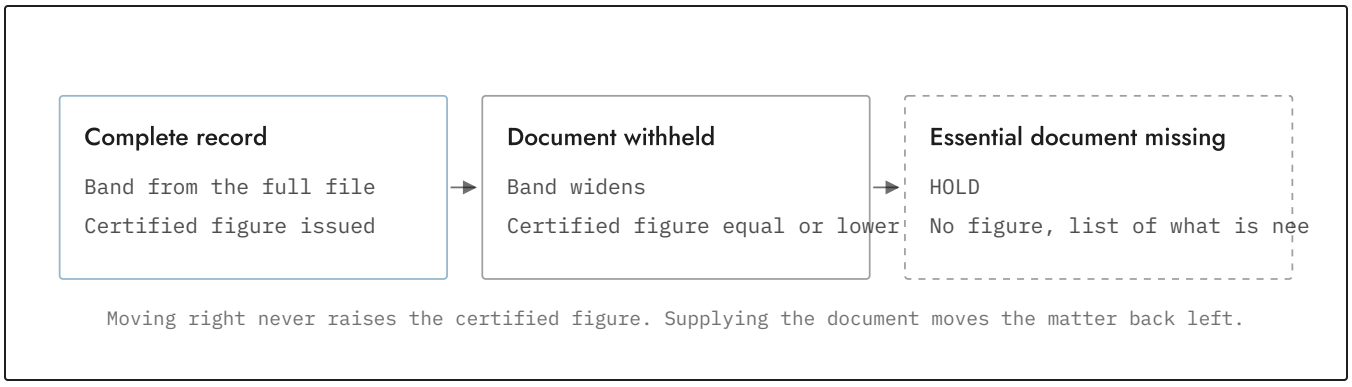
Criterica Intelligence treats this as a design requirement, not a review step. This paper states the rule, explains why it has to be a rule rather than a tendency, works an illustrative example and sets out what it means for a reader working through a data room.

02

The rule, in three parts

1. **Withholding information never raises a certified figure.** Remove any document from a file and the certified figure for that matter stays the same or falls. It never rises.[1]
2. **A matter on hold carries no figure.** When an essential document is missing, the matter is placed on HOLD. A HOLD states what is needed. It does not carry a placeholder, a zero or an estimate.
3. **One exact calculation sets every number.** Each certified figure comes from one exact computation over the stated inputs. Faster approximations may be used to check that result, never to set it.

FIGURE 1 **Three states of a figure** Illustrative



A matter moves right as evidence is removed and left as it is supplied. The certified figure never rises on a move to the right. This figure is evidence of the ordering rule: what the method gives up, and in what order, as evidence is removed.

The three parts work together. The first closes the incentive to withhold. The second prevents a missing document from being disguised as a low but usable number, which a reader could mistake for a judgement on the merits. The third ensures that the figure the reader sees is the figure the method actually produces, with no shortcut that could drift above it.

03

Why the rule has to be built in, not hoped for

It is tempting to assume that any sensible method already behaves this way. Many do not. The most common way to handle a gap in a file is to fill it with something typical: the average treatment cost for the injury, the median award in the venue, a value predicted from the rest of the file. That is reasonable for a population statistic and wrong for an individual decision.[2]

Consider a matter whose treatment records, if supplied, would show conservative care that ended early. Fill the gap with the typical treatment course for that injury and the matter looks better than the complete file would. The seller who withheld the records has been paid for withholding them. Repeat that across a portfolio and the thinnest files are systematically over-valued, which is the opposite of what a buyer needs.

The rule is therefore enforced as a mathematical property of the method rather than as a guideline for reviewers. On every release, the full test set is re-run with documents removed in every combination the test design allows, and the release fails if any certified figure rises.[3] A guideline can be forgotten under deadline; a failing release cannot ship.

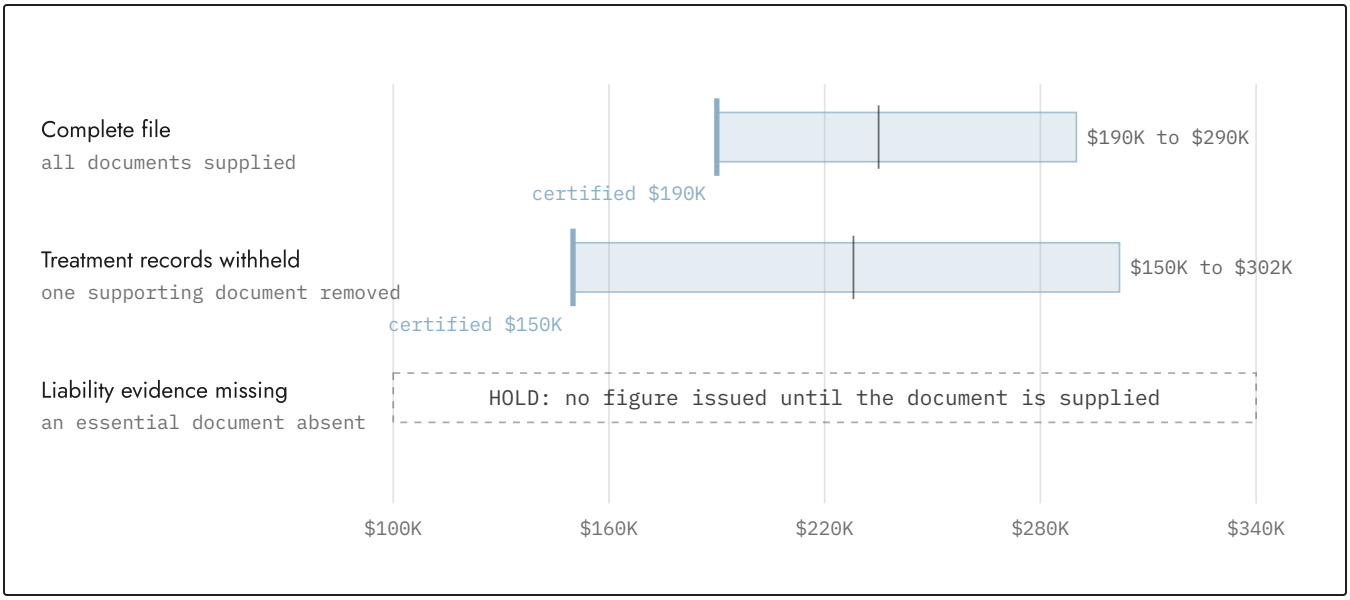
> A valuation method is only as honest as the incentive it gives the person assembling the file.

04

A worked example: a document is removed and the band widens

The illustration below follows one fictional personal-injury matter through three versions of its file. In this illustration the certified figure is the lower edge of the settlement band, the value a funder could size against with the stated confidence.

FIGURE 2 **The same matter, three versions of the file** Illustrative



Shaded span is the settlement band; thin mark is the band median; heavy mark is the certified figure. Removing treatment records widens the band in both directions and lowers the certified figure. Removing the liability evidence places the matter on HOLD. This figure is evidence of the mechanism, shown on an illustrative matter, and not of a measured effect size.

VERSION OF THE FILE	BAND	MEDIAN	CERTIFIED FIGURE
Complete file	$190,000 to $290,000	$235,000	$190,000
Treatment records withheld	$150,000 to $302,000	$228,000	$150,000
Liability evidence missing	none	none	HOLD: liability evidence requested

What happens when the treatment records are withheld

With the complete file, the band runs from $190,000 to $290,000 and the certified figure is $190,000. Withhold the treatment records and the method can no longer distinguish a short course of care from a long one, so it widens the band on both sides: $150,000 to $302,000. The median barely moves, because the missing records could point either way. The certified figure falls to $150,000, because the conservative edge has moved down.

The upper edge of the band rose, from $290,000 to $302,000. That is correct and expected: with less information, a higher outcome is more plausible as well as a lower one. A wider band is not a better valuation, and nothing a reader would rely on increased. The certified figure, the only number a decision is sized against, went down by $40,000. That $40,000 is the measured value of the treatment records to this decision, and it is also the clearest possible request: supply them.

What happens when the liability evidence is missing

Remove the police report and the driver's statement and the question of fault is open. For this kind of matter that evidence is essential: without it there is no defensible settlement band at all. The method does not guess. It places the matter on HOLD, carries no figure, and lists the documents that would release it.

05

What a HOLD is, and why it carries no figure

A HOLD is a complete answer. It says that the record supplied does not support a figure for this purpose, names what is missing and states what the figure will be based on once it arrives. It is issued as readily as a number, and it is not a negative judgement on the matter.

The reason it carries no figure is practical. A placeholder number is read as a number. A $0 in a schedule looks like a matter with no value; an imputed average looks like a matter that was assessed. Either can be summed into a portfolio total, advanced against or quoted in a memo, and the reader downstream has no way to tell that it was never a measurement. Keeping HOLD matters out of the figures keeps every total honest: the portfolio value is the sum of matters that were actually valued, with the held matters counted and listed beside it.

- Every HOLD names the specific documents or facts required.
- A held matter is read again as soon as the missing material arrives.
- Portfolio totals report valued matters and held matters separately, never blended.

06

One exact calculation sets every number

Large books invite shortcuts. A method that values thousands of matters may be tempted to use a fast approximation for most of them and the exact calculation for a sample. Approximations are useful, and Criterica uses them to check work. They are never allowed to set a certified figure, because an approximation that runs above the exact result on even a small share of matters breaks the first part of the rule for those matters, and nobody reading the schedule would know which ones.

The commitment is simple to state and to verify: the figure on the page is the figure the exact method produces from the stated inputs and versions. It can be regenerated on request, and it will match. WP-005 explains why that reproducibility matters more to a decision than any single answer.

07

What the rule means in a data room

For a funder, buyer or lender, the rule changes how an incomplete file should be read and what it is worth asking for.

1. **Send what you have.** There is no advantage in waiting for a complete file before asking for a read. An incomplete file returns a figure that is equal to or lower than the complete one, plus a list of what would change it.

2. **Read the gap as a price.** The difference between the certified figure now and the figure with a document supplied is the value of that document to your decision. Use it to decide which requests to press with a seller.
3. **Treat a widening band as information.** A band that widens when a document is removed shows where the uncertainty lives. A band that does not move shows the document did not matter to this question.
4. **Keep HOLDs out of totals.** When you compare a seller's schedule to a Criterica read, match valued matters to valued matters and list held matters separately. A seller total that includes unvalued matters at a placeholder is not comparable.
5. **Ask every method the same question.** Whatever system you rely on, ask what happens to its number when a document is removed. If the answer is "it depends", the incentive to withhold is still open.

The way to read the bands themselves is set out in WP-006, and the calibration tests behind them in TN-003. The Criterica Score applies the same rule at the level of a whole asset; see the Score Specification.

08

The question, as the reader has it

An underwriter at a funder, a buyer pricing a book of receivables and a claims team reserving a file all hold the same problem. They are handed a file, they cannot see what was left out, and they must commit capital or a reserve to a single figure. The question is not "what is this matter worth?" in the abstract. It is: given exactly these documents, what may I rely on, and how does that change if someone sends me more, or fewer?

A point estimate cannot answer that question, for a structural reason. A single number carries no account of how much of the file stands behind it. Two files can produce the same point figure, one from a complete record and one from a record with its most informative document missing, and the reader cannot tell them apart. A labelled band can: its width records how much the file left

open, and its conservative edge records what the file supports even if the missing material points the wrong way.

09

What is measured

Three quantities are involved, each with its unit and its population stated in public terms.

QUANTITY	UNIT	POPULATION IT IS READ AGAINST
Settlement band	US dollars, as a range with a median	Closed matters of the same case type, venue and stage, drawn from filed court records with a recorded termination
Certified figure	US dollars, the conservative edge of the band at a stated date	The band above, given the documents supplied for this matter
File completeness	Which required documents are present, absent or withheld	The documents that comparable resolved matters turned on

Inclusion follows one logic: a matter enters the comparison population only if its record shows how it ended. Matters still pending are not counted as outcomes, and a matter whose ending the record does not show is not treated as a zero. The band is therefore a statement about closed matters, adjusted for the facts in the file in front of the reader, and not a statement about what any one matter will realize.

10

How a number travels from the record to the band

1. **Record to structured event.** Filed documents are read into dated events: a filing, an order, a stated amount, a stated course of treatment. Parties and courts are matched across records by

entity resolution on a knowledge graph, so one claimant, one defendant or one judge is one node, however many spellings the record uses.

2. **Structured event to distribution.** The matter is placed against comparable closed matters. Where a venue or case type has few of them, hierarchical partial pooling lets the sparse cell borrow strength from its parent group in proportion to how little data it holds, so a thin venue is neither over-trusted nor discarded.

3. **Distribution to calibrated band.** The settlement distribution is stated through quantile and conformal intervals, a class of method that gives a band a stated long-run coverage on matters it has not seen. Quantile estimation fits the problem because dollar outcomes are skewed and a mean misrepresents them; conformal calibration fits because it makes the coverage claim testable.

4. **Band to certified figure.** The certified figure is read from the conservative edge of the band, under the three-part rule of this paper. When a document is removed, the evidence for the missing fact is replaced by the full spread of what comparable matters did, which can only widen the band and can only lower its conservative edge. That is why the certified figure falls while the upper edge may rise.

The rule is the reason the last step is trustworthy. The earlier steps are statistical estimation; the last is a stated, checkable constraint on how an estimate becomes a figure a decision may rely on.

How the rule and the bands are validated

Two different claims are validated, and each has its own design.

The rule: a property tested over a generated set

The claim that removing information never raises a certified figure is a universal statement, so it is tested as one. A property test generates a large set of files, removes documents from each in the combinations the test design allows, and asserts that the certified figure of the reduced file is never above that of the full file.[3] The method follows the randomized property-testing tradition rather than

a hand-picked list of examples, because a hand-picked list tests the cases someone thought of. A single violation fails the release.

The bands: out-of-time evaluation

The settlement bands themselves are evaluated on a held-out later period. What is frozen is the model as it stood at the time cut; what is compared is each band against how the matter actually closed; what prevents leakage is that no record dated after the cut is visible to the model, and that the point-in-time boundary applies to every field in the file, not only to the outcome. The metric families are interval coverage (does an 80% band contain about 80% of outcomes) and interval width (is it as narrow as the evidence allows). TN-003 shows how to read both.

Figure 2 is evidence of the mechanism, not of a measured effect: it shows how a band and a certified figure respond when a single document is removed from an illustrative matter, so the three-part rule can be checked by hand against a case small enough to follow.

Scope and precision

The certified figure is the conservative edge of a distribution of settlement outcomes for comparable closed matters. It is a reading of the record for a stated purpose and date. It is not a forecast of the dollars a specific matter will realize, and it does not say who will prevail: outcome direction is read only from sources that record it, as WP-006 sets out.

- Data lag: the record reflects filings and rulings up to its as-of date, so a figure is re-issued as the record moves.
- Censoring: matters still open are not outcomes, so the closed population leans toward matters that ended.
- Venue sparsity: where few comparable matters exist, the band is wider, and the width says so.

Each of these makes a band more honest, not less useful: the width of a band is where the reader sees how much the record supports.

13

What a reader can reproduce

Everything in the worked example is arithmetic on the published table: the band, the median, the certified figure and the change between versions of the file can be recomputed by hand from the values shown. Any reader can also apply the rule's test to a system of their own by asking for the certified figure on a file, removing one document and asking again; the answer must not rise. Reproducing a figure for a real matter requires the platform, because that depends on the record version and model version stated beside it. Both are supplied with the figure so that it can be regenerated exactly on request.

What this number means

A certified figure is the value a decision may rely on for its stated purpose, at its stated date, given the documents supplied. It is a conservative reading of a labelled distribution, not a forecast of the dollars a matter will realize. When a document arrives the figure is re-issued and the earlier one is kept, so a reader can always see what changed and why. The illustrative matter in Figure 1 is not a real file and its figures are chosen to show the mechanism.

REFERENCES

1. Rubin, D. B. (1976). Inference and missing data. Biometrika, 63(3), 581–592. doi.org/10.1093/biomet/63.3.581
2. Little, R. J. A., & Rubin, D. B. (2019). Statistical Analysis with Missing Data (3rd ed.). Wiley.
3. Akerlof, G. A. (1970). The market for "lemons": quality uncertainty and the market mechanism. Quarterly Journal of Economics, 84(3), 488–500. doi.org/10.2307/1879431
4. Milgrom, P. R. (1981). Good news and bad news: representation theorems and applications. Bell Journal of Economics, 12(2), 380–391. doi.org/10.2307/3003562

5. Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. Journal of Law and Economics, 24(3), 461–483. doi.org/10.1086/466981
6. Claessen, K., & Hughes, J. (2000). QuickCheck: a lightweight tool for random testing of Haskell programs. Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming, 268–279. doi.org/10.1145/351240.351266
7. Gelman, A., & Hill, J. (2007). Data Analysis Using Regression and Multilevel/Hierarchical Models. Cambridge University Press.
8. Fellegi, I. P., & Sunter, A. B. (1969). A theory of record linkage. Journal of the American Statistical Association, 64(328), 1183–1210. doi.org/10.1080/01621459.1969.10501049
9. Koenker, R., & Bassett, G. (1978). Regression quantiles. Econometrica, 46(1), 33–50. doi.org/10.2307/1913643
10. Angelopoulos, A. N., & Bates, S. (2023). Conformal prediction: a gentle introduction. Foundations and Trends in Machine Learning, 16(4), 494–591. arxiv.org/abs/2107.07511
11. Romano, Y., Patterson, E., & Candès, E. J. (2019). Conformalized quantile regression. Advances in Neural Information Processing Systems, 32. arxiv.org/abs/1905.03222
12. Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 102(477), 359–378. doi.org/10.1198/016214506000001437

NOTES

1. The rule applies to every certified figure, including portfolio and facility figures built up from matter-level figures.
2. Mean imputation is used here as the plainest example. Any method that fills a gap with a typical value, a model prediction or a seller-supplied estimate can produce the same effect if it is not constrained.
3. A property test asserts that a rule holds for every case in a large generated set rather than for a handful of chosen examples. A single violation fails the release.

Prepared by Criterica Intelligence Research with data engineering, validation and editorial review. Criterica Intelligence Research publishes numbered documents in five series; each is versioned, dated and issued as a new version when it changes, with the change recorded.

CITE THIS DOCUMENT

Criterica Intelligence Research. (2026). When a Document Is Missing, the Number Does Not Go Up (Working Paper WP-004, version 1.0). Criterica. https://critericaintelligence.com/research/missing-document

```
@techreport{criterica_wp004_2026,
  author = {{Criterica Intelligence Research}},
  title = {When a Document Is Missing, the Number Does Not Go Up},
  institution = {Criterica},
  type = {Working Paper},
  number = {WP-004},
  year = {2026},
  month = nov,
  note = {Version 1.0},
  url = {https://critericaintelligence.com/research/missing-document}
}
```

DOI: pending