

Duration and Returns: The Arithmetic of Time in Litigation Finance

A multiple is time-blind. An annualized return is not. This paper works through what that asymmetry does to legal-asset economics — using nothing but arithmetic any reader can verify by hand. No model output appears anywhere in it. The magnitudes are large, the mechanics are simple, and the industry's pricing habits have not caught up to either.

14-MINUTE READ • ARITHMETIC ONLY • INDEPENDENTLY RECOMPUTED BEFORE PUBLICATION

FIGURE 1 • ANNUALIZED RETURN VS. DURATION AT FOUR FIXED MULTIPLES • $IRR = MOIC^{(12/T)} - 1$

SERIES

Working Paper WP-003

VERSION

1.0

PUBLISHED

3 August 2026

STATUS

Published

AUTHORS

Criteria Intelligence Research

Corresponding author: Jack Cosentino

The multiple cannot see time. The return is made of it.

Litigation finance prices its work in multiples: a matter is underwritten to return 1.8x, 2.0x, 2.5x invested capital. The multiple is the natural language of the asset, because it is what the legal

outcome actually determines. But allocators do not compare multiples across asset classes; they compare annualized returns. The bridge between the two is a single identity:

$$\text{IRR} = \text{MOIC}^{12 / T} - 1 \quad (T \text{ in months})$$

Everything in this paper follows from that line. A 1.8x multiple is not a return until it has a denominator, and the denominator is time. The same successful case — same merits, same outcome, same cash — is a 34.2% asset if it resolves in 24 months, a 21.6% asset at 36 months, and a 15.8% asset at 48. Nothing legal changed across those three worlds. Only the clock did.

Duration errors never appear in the multiple. They silently reprice the annualized return — which is the number the capital was raised on.

This is why duration is the least priced variable in the asset class. Underwriting interrogates the merits because the merits decide the multiple, and the multiple is visible. Time is treated as weather: acknowledged, endured, rarely underwritten with the same discipline — even though, within the realistic range of case lengths, the clock moves the annualized return more than most contested merits questions do.

02 · THE SENSITIVITY SURFACE

What the clock does at every point on the book.

Table 1 is the whole mechanism at a glance: annualized return by multiple and duration. Read along any row and watch time consume the return while the legal outcome holds perfectly still.

MOIC \ DURATION	24 MO	30 MO	36 MO	42 MO	48 MO
1.5x	22.5%	17.6%	14.5%	12.3%	10.7%
1.8x	34.2%	26.5%	21.6%	18.3%	15.8%
2.0x	41.4%	32.0%	26.0%	21.9%	18.9%
2.5x	58.1%	44.3%	35.7%	29.9%	25.7%

Table 1 · IRR by multiple and duration, single terminal cash flow. Depth of shading tracks magnitude.

Table 2 restates the surface as a question an investment committee can act on: at a given multiple, what is a defined duration improvement worth? The answer is the spread between two committed positions on Table 1 — earned without changing a single legal fact.

MOIC \ COMPRESSION	3 MONTHS	6 MONTHS	9 MONTHS	12 MONTHS
1.5x	+1.4pp	+3.1pp	+5.3pp	+8.0pp
1.8x	+2.2pp	+4.9pp	+8.2pp	+12.5pp
2.0x	+2.7pp	+6.0pp	+10.1pp	+15.4pp
2.5x	+3.8pp	+8.5pp	+14.5pp	+22.4pp

Table 2 · IRR gained per months of duration compression, from a 36-month base.

THE CONVEXITY

A month is not worth the same everywhere.

The decay curve in Figure 1 is convex, and the convexity carries an underappreciated instruction. At a fixed 1.8x, one month of duration is worth 1.73 points of IRR on a 24-month case — and 0.36 points on a 48-month case. The same calendar month is nearly five times more valuable at the short end of the book.

1.73pp

1.03pp

0.68pp

0.49pp

0.36pp

24-mo case

30-mo case

36-mo case

42-mo case

48-mo case

Duration discipline is most valuable exactly where books believe they are already safe: on the short, “easy” matters, where a single unpriced continuance does the most annualized damage.

The operational consequence: triage duration attention by marginal value, not by absolute length. Long matters attract the scrutiny because their totals look alarming; short matters quietly carry the highest per-month stakes on the book.

03 · FOUR EFFECTS, KEPT SEPARATE

Duration moves the economics through four distinct channels. Summing them is how credibility dies.

Most attempts to value duration improvement collapse everything into one headline number, which double-counts and invites the audience to dismiss the whole exercise. The honest structure is four effects with four different units, presented separately and never added:

I · ANNUALIZED-RETURN EFFECT · IRR POINTS

The identity itself: at fixed MOIC, less T means more IRR. Tables 1 and 2 quantify it completely. This is the only effect that needs no assumptions at all.

II • LIQUIDITY-TIMING EFFECT • PRESENT VALUE

Cash that arrives earlier is worth more at the fund's own discount rate, and it de-risks every commitment standing behind it. Priced by each reader with their own rate; we supply no number.

III • REDEPLOYMENT EFFECT • CAPITAL TURNS

Faster resolution returns capital that can be deployed again. Section 05 gives the turns arithmetic and its honest frictionless caveat.

IV • CREDIBILITY EFFECT • QUALITATIVE

Realization guidance that lands narrows the gap between reported marks and market trust. Real, cumulative, and deliberately left unquantified here.

THESE FOUR NUMBERS MUST NEVER BE SUMMED. They are different units measuring different consequences of the same months. One combined figure would be rhetorically larger and analytically indefensible.

04 • THE PORTFOLIO VIEW

Plans are priced off the median. Financing pressure lives in the tail.

A portfolio's duration is not a number; it is a distribution, and its two ends serve different masters. Realization guidance and deployment pacing are naturally built around the middle of that distribution. But reserves, facility headroom, and refinancing risk are governed by its right tail —

the matters that run long. Across listed funders' own public reporting, portfolio weighted-average lives have been extending, which means the industry's tails are growing at exactly the moment its guidance credibility matters most.

MEDIAN • THE PLAN



This is why duration risk needs tail-first metrics, not averages. The 90th percentile of time to resolution — call it TaR90 — is the tail a reserve must actually carry. The spread between TaR90 and the median, Excess TaR, is the gap in which financing pressure accumulates: every dollar-month inside that span is capital that the plan did not price but the balance sheet must fund. A book can hold its median perfectly and still be squeezed, because the squeeze was never a median phenomenon.

The portfolio arithmetic follows the matter arithmetic: compressing the tail is worth more than compressing the middle, both because the convexity of Section 02 is working against the long matters' totals, and because tail months are funded months, paid for twice — once in forgone IRR and once in carry cost.

05 • VELOCITY

The decade view: the same book, one extra turn.

Over a ten-year window, average duration sets how many times capital can work. At a 36-month average, capital turns 3.33 times a decade; at 30 months, 4 times. Under a deliberately frictionless assumption — instant, full redeployment at an unchanged 1.8x — those paths compound to 7.1x and 10.5x respectively. Reality redeploys with friction and pipelines are not infinite, so treat these as upper bounds on a real effect, not projections. The direction and rough scale survive any honest haircut: six average months is worth roughly an extra turn of the book per decade.

The arithmetic prices the month. The book decides how many months are yours.

Nothing above claims that any given portfolio can shorten its durations. The arithmetic only establishes what a month is worth wherever one can be found — and that the value is largest on the matters that look safest. Whether months are recoverable is an empirical, book-specific question, and an honest one: some delay is structural (court congestion, statutory waits) and belongs in pricing, not in plans to defeat it. Some delay is operational — late recognition of drift, slow settlement posture reviews, enforcement preparation that starts after judgment instead of before — and that portion responds to measurement.

Which is the real argument of this paper: before duration can be managed, it has to be measured like the capital variable it is — as a distribution with a stated tail, re-estimated as evidence arrives, with the original estimate preserved so that accuracy is auditable. An institution that measures duration this way learns, matter by matter, which of its months were structural and which were recoverable. An institution that carries duration as a single assumed number learns only at realization, when the information is worthless.

You do not need to believe any model to accept this arithmetic. The only open question is how many of the months are yours to recover — and that question has an empirical answer.

07 · METHOD NOTES AND LIMITS

Where this simplifies, and why the direction survives.

SINGLE TERMINAL CASH FLOW

The identity assumes one outflow at commitment and one inflow at resolution. Real matters deploy in tranches, which shortens effective duration and raises IRR relative to these figures; a full treatment uses XIRR over the actual schedule. The simplification is conservative for the paper's argument and keeps every figure reproducible by hand.

GROSS, NOT NET

All figures are gross asset-level arithmetic. Fees, expenses, and fund-level structuring sit between these numbers and an investor's net return and differ by vehicle; nothing here estimates them.

IRR'S OWN ASSUMPTIONS

IRR implicitly assumes interim proceeds re-earn the same rate; the turns analysis in Section 05 makes that assumption explicit and labels it a frictionless bound. Where reinvestment is weak, effect III shrinks; effects I and II do not.

NO MODEL CLAIMS

No output of any predictive model appears in this paper. Criterica's duration measurement — cohort bands, event-conditioned re-estimates, preserved forecast history — is described at /duration, and its methods and boundaries at /methodology.

Statistics shown reflect historical or illustrative model outputs derived from real case data. They are not predictions or guarantees of any individual outcome. Litigation results depend on facts, jurisdiction, judge, and counsel, and vary case by case. Model

accuracy is subject to selection effects and changing legal dynamics.



THE DURATION RESEARCH SERIES

Paper I of III.

II • THE SEVEN CLOCKS

III • TIME-AT-RISK

DURATION INTELLIGENCE

ALL RESEARCH

CITE THIS DOCUMENT

Criteria Intelligence Research (2026). Duration and Returns: The Arithmetic of Time in Litigation Finance. Working Paper WP-003, version 1.0. Criteria. <https://critericaintelligence.com/research/duration-and-returns>

```
@techreport{criterica_wp003_2026,  
  author = {{Criteria Intelligence Research} and Cosentino, Jack},  
  title = {Duration and Returns: The Arithmetic of Time in Litigation Finance},  
  institution = {Criteria},  
  type = {Working Paper},  
  number = {WP-003},  
  year = {2026},  
  month = aug,  
  note = {Version 1.0},  
  url = {https://critericaintelligence.com/research/duration-and-returns}  
}
```

DOI: pending