

Time-at-Risk: A Metric Family for Legal-Asset Duration

Credit has ratings. Market risk has VaR. Legal assets have had adjectives. This paper proposes the metric family that replaces them: six defined, computable measures of duration risk — each with a formula, a unit, and the decision it feeds — demonstrated with real values computed from the public federal record. Whoever names a risk can price it. Until then, everyone carries it anyway.

11-MINUTE READ • DEFINITIONS + REAL COMPUTED VALUES • PUBLIC FEDERAL RECORD

SERIES

Working Paper WP-001

VERSION

1.0

PUBLISHED

3 August 2026

STATUS

Published

AUTHORS

Criteria Intelligence Research

Corresponding author: Jack Cosentino

01 • FROM ADJECTIVES TO NUMBERS

Unnamed risks get carried. Named risks get priced.

Every institutional asset class that matured did so by converting its defining risk from vocabulary into measurement. Fixed income had “safe” issuers until it had ratings and spreads; portfolios had “acceptable” exposure until VaR gave the board one number with a definition behind it. Legal assets are late to this conversion on their defining risk — time. Duration today is discussed in adjectives: matters run “long,” books feel “slow,” vintages are “delayed.” None of those words can be reserved against, stressed, or put in a covenant.

The companion papers established what time is worth (Paper I) and why duration is a living distribution rather than a number (Paper II). This paper supplies the measurement layer that makes both operational.

02 · THE FAMILY

Six measures. Each with a formula, a unit, and a decision.

TaR90

DAYS · YEARS

$$\text{TaR90} = Q90(T)$$

The 90th percentile of time to resolution for a matter's true cohort. The tail a reserve must actually carry, promoted to a first-class number. Stated at underwriting, monitored thereafter.

FEEDS → RESERVING · PRICING THE TAIL · FACILITY HEADROOM

Excess TaR

DAYS · YEARS

$$\text{xTaR} = Q90(T) - Q50(T)$$

The spread between the tail and the median: the distance between the number the plan is built on and the number one matter in ten delivers. The gap in which financing pressure accumulates.

FEEDS → PORTFOLIO STRESS · REFINANCING RISK · CARRY COST OF THE TAIL

Duration Drift

Δ DISTRIBUTION

$$\text{Drift}(t) = B(t) - B(t_0)$$

The movement of a matter's duration band between the original underwriting estimate and the current evidence-conditioned one — with the original preserved, so the drift is an auditable fact rather than a recollection.

Cash-at-Date

PROBABILITY

$$\text{CaD}(d) = P(T \leq d)$$

The probability that a defined share of a book resolves by a given date, built from matter-level distributions rather than a single assumed average. Realization planning as a distribution.

FEEDS → LIQUIDITY PLANNING · REALIZATION GUIDANCE · DEPLOYMENT PACING

Delay Concentration

SHARE OF EXPOSURE

$$\text{DC} = \max_g \Sigma(T \cdot w \mid g) / \Sigma(T \cdot w)$$

The largest share of a book's time-weighted exposure riding a single common driver — one court's calendar, one doctrine, one procedural regime. Duration risk correlates through shared environments before anything else.

FEEDS → CONCENTRATION LIMITS · CORRELATED-DELAY STRESS · DIVERSIFICATION

Recoverable Time

MONTHS


$$\text{RT} = E[T] - E[T \mid \text{feasible action}]$$

The portion of expected delay that feasible action could remove without damaging position. Defined so it can be measured against outcomes — which is what separates a management claim from a marketing one.

FEEDS → INTERVENTION TRIAGE · THE EMPIRICAL VALUE OF DURATION MANAGEMENT

Excess TaR is a case-type property — and it varies five-fold.

Computed from the public federal record across four national cohorts totaling more than 650,000 resolved matters:




The instruction in the figure: the tail is not a book-level accident, it is a case-type property. Contract matters carry 1.6 years of Excess TaR; antitrust carries 3.8 — a median that looks quick (0.8 years) attached to a tail that runs past 4.5. Two books with identical median durations can carry radically different Time-at-Risk, and only the tail-first metrics can tell them apart.

A book is not “fast” or “slow.” It has a median, a TaR90, and a spread between them — and the spread is where the financing pressure lives.

Realization planning as a distribution.

For the national securities cohort (n = 68,160), the real resolution profile:



by 2 years		69.5%
by 3 years		81.7%
by 4 years		89.8%

Read as a finance instrument rather than a statistic: half of this cohort’s capital timeline resolves inside a year, seven in ten inside two — and one in ten is still unresolved at four. A treasury function planning against the median alone is planning against a number that 50% of matters will miss by construction. Cash-at-Date turns that from a surprise into a schedule.

05 · WHAT MAKES A METRIC MEAN ANYTHING

Definitions are free. Discipline is the product.

A metric family only earns institutional weight under four disciplines. Temporal validation: every underlying estimate tested on periods it never saw, because random splits hide regime change. Censoring treatment: open matters accounted for, never silently dropped, because dropping them biases every duration statistic optimistic. Preserved history: the estimate at each point in time kept on the record, so Duration Drift is a fact and accuracy is auditable. And sufficiency floors: below a stated cohort size the system returns tiers rather than false precision, and says so. These are the standards Criterica’s duration measurement is governed by; the full statement of methods and boundaries is at [/duration](#) and [/methodology](#).

Method notes. Cohorts: resolved matters in the public FJC Integrated Data Base, grouped by NOS-coded case family, national scope; quantiles computed on filing-to-disposition days. Recoverable Time is defined here but not computed — by construction it requires intervention histories and outcomes, which public data does not contain; publishing a number for it without that evidence would be exactly the practice this series argues against. No figure on this page predicts any individual matter.

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 • COMPLETE

I • DURATION AND RETURNS

II • THE SEVEN CLOCKS

DURATION INTELLIGENCE

CITE THIS DOCUMENT

Criteria Intelligence Research (2026). Time-at-Risk: A Metric Family for Legal-Asset Duration. Working Paper WP-001, version 1.0. Criteria. <https://critericaintelligence.com/research/time-at-risk>

```
@techreport{criterica_wp001_2026,  
  author = {{Criteria Intelligence Research} and Cosentino, Jack},  
  title = {Time-at-Risk: A Metric Family for Legal-Asset Duration},  
  institution = {Criteria},  
  type = {Working Paper},  
  number = {WP-001},  
  year = {2026},  
  month = aug,  
  note = {Version 1.0},  
  url = {https://critericaintelligence.com/research/time-at-risk}  
}
```

DOI: pending