

The Seven Clocks: Why Litigation Duration Is Not One Number

“This case should resolve in about 36 months” sounds like a forecast. It is actually seven forecasts wearing one coat — seven processes with different drivers, different endpoints, different data, and different responses to intervention. This paper separates them, then demonstrates on a real matter from the public federal record why a duration estimate must be a living distribution, not a number spoken once at underwriting.

12-MINUTE READ • PUBLIC FEDERAL RECORD • 55,431-MATTER COHORT • RECOMPUTED BEFORE PUBLICATION

FIGURE 1 • SEVEN CLOCKS, SEVEN SPEEDS • THE LAST ONE IS THE ONLY ONE A DECADE-LEVEL RESULT CAN SEE

SERIES

Working Paper WP-002

VERSION

1.0

PUBLISHED

3 August 2026

STATUS

Published

AUTHORS

Criteria Intelligence Research

Corresponding author: Jack Cosentino

01 • SEVEN PROCESSES, ONE COAT

Separate the clocks and each one becomes tractable.

Collapsed into one number, duration can only be endured. Separated, each clock has its own evidence base and its own levers. The first three decide when legal risk resolves; the middle two decide when a right becomes enforceable; the last two decide when anyone is actually paid — and how often a portfolio's capital gets to work per decade.

01 • NEXT PROCEDURAL EVENT

Filing to the next consequential motion, ruling, or hearing. The shortest clock — and the one that re-times all six others when it lands.

02 • SETTLEMENT

Filing to negotiated resolution. Its hazard is not constant; it concentrates where procedural pressure changes both sides' incentives at once.

03 • ADJUDICATION

Filing to a decision on the merits. Runs on the court's calendar and the court's congestion, not on either party's preferences.

04 • APPEAL

Judgment to appellate resolution. A discrete extra distribution that only some matters draw, conditional on how they arrived at judgment.

05 • ENFORCEMENT

Judgment to enforceability. Domestically often short; cross-border and sovereign, frequently the dominant clock of the entire matter.

06 • CASH COLLECTION

Enforceability to money received. Legal resolution and cash resolution are different dates, and only one of them pays anyone.

07 • CAPITAL RECYCLING

Commitment to redeployment. The clock a funded portfolio actually compounds on — and the only one visible in a fund's decade-level results.

02 · COMPETING ENDPOINTS

A matter does not have a duration until an endpoint claims it.


The first five clocks are not sequential stages; they are competitors. Settlement, adjudication, and dismissal race each other, and whichever endpoint arrives first determines both the duration and which of the other clocks ever start. This is why a single expected duration is statistically malformed: it averages across races with different finish lines. The proper object is a set of conditional distributions — if this matter settles, when; if it is adjudicated, when; if judgment comes, what appeal and enforcement add — each updating as events eliminate competitors.

One honesty note that shapes everything downstream: the public federal record codes how cases terminate only coarsely, and it cannot reliably distinguish which negotiated outcomes were settlements or who prevailed. That is a finding, not a footnote — it means public-data duration work must be built on what the record actually supports: resolution timing, procedural posture, and survival. Every figure in this paper is built on exactly that, and nothing else.

03 · A REAL MATTER, RE-ESTIMATED FOUR TIMES

Watch a tail case become a median case, one birthday at a time.

A real contract matter from the New York federal courts, drawn from the public record: filed in July 2019, resolved by judgment 700 days later. Below is the duration estimate its true cohort supports at four moments of its life — each band the interquartile range of the 55,431-matter cohort, conditioned on the matter still being open at that age.



At filing, 700 days sits at the cohort's 80th percentile — a tail outcome, correctly flagged as such. Six months of survival moves it to the 68th. At one year, the once-tail matter is *exactly median*: the act of surviving has re-sorted it among its true peers. At eighteen months the band's lower edge has climbed to 1.92 years — and the matter resolves precisely there, at the 25th percentile of what remained. No model was fitted to this matter. Every band is deterministic cohort arithmetic, reproducible by anyone with the public record.

The estimate was never wrong. It was honest four different times, about four different questions — because “how long will this take?” is a different question at every age of a matter.

This is the mechanical case for continuous re-estimation. An underwriting-day number, however good, answers only the day-zero question, and its information content decays the moment the matter starts surviving. The discipline that makes the living estimate trustworthy is preservation: each prior band stays on the record beside its successor, so accuracy is auditable after the fact.

The underwriting question is never “when exactly.”

It is “what share of outcomes is realistically resolved by each date, and what tail am I carrying past it.” For the same New York federal contract cohort, the real answer:

by 6 months		37.6%
by 12 months		61.0%
by 18 months		73.8%
by 24 months		81.4%
by 36 months		90.2%

Figure 3 · share of the 55,431-matter cohort resolved by age. One matter in ten is still unresolved at three years.

05 · PRACTICE, METHOD, LIMITS

Re-forecast on events, not on calendars.

The operational translation of everything above: duration estimates should update when evidence arrives — a survived motion, a certification decision, a scheduling order — not when a review cycle happens to come around. Each of the seven clocks has its own event triggers, and the racing structure of Section 02 means a single event often re-times several clocks at once. Criterica’s duration measurement is built this way; the approach and its boundaries are described at /duration.

Method notes. Cohort: resolved contract-coded matters in the New York federal district courts, public FJC Integrated Data Base, n = 55,431; conditional bands are interquartile ranges of the cohort restricted to matters surviving past each age. The single worked matter illustrates mechanics; coverage claims come from populations, never from an example. Right-censoring: conditional cohorts are built from resolved matters, so recent filing years thin the long-duration cells; the conditional medians shown are robust to this at the cohort sizes used. Nothing here predicts any individual matter’s outcome or asserts who prevailed in any historical one.

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

I • DURATION AND RETURNS

III • TIME-AT-RISK

DURATION INTELLIGENCE

CITE THIS DOCUMENT

Criteria Intelligence Research (2026). The Seven Clocks: Why Litigation Duration Is Not One Number. Working Paper WP-002, version 1.0. Criteria. <https://critericaintelligence.com/research/seven-clocks>

```
@techreport{criterica_wp002_2026,  
  author = {{Criteria Intelligence Research} and Cosentino, Jack},  
  title = {The Seven Clocks: Why Litigation Duration Is Not One Number},  
  institution = {Criteria},  
  type = {Working Paper},  
  number = {WP-002},  
  year = {2026},  
  month = aug,  
  note = {Version 1.0},  
  url = {https://critericaintelligence.com/research/seven-clocks}  
}
```

DOI: pending