Quantome

Comprehensive quantitative solutions for superior outcomes

WHITE PAPER

Fair Value Engine for SBA 7(a) Loan Pools

A Monte Carlo Simulation Framework for Pricing,
Risk Measurement, and Option-Adjusted Spread Analysis

Q-SBA Analytics Platform
February 2026

Executive Summary

Government-guaranteed lending programs - particularly the SBA 7(a) loan program - have long provided a vital source of capital for small businesses across the United States. When these loans are pooled and securitized, investors gain exposure to a diversified portfolio of small business credit backed by a partial government guarantee. However, determining the fair value of these pools remains a complex undertaking: the loans are variable-rate, subject to prepayment and default risk, and their cash flows depend on the future path of interest rates.

The Fair Value Engine, part of the Q-SBA Analytics Platform developed by Quantome, addresses this challenge with a rigorous, transparent, and fully integrated valuation framework. It combines stochastic interest rate modeling, loan-level cash flow projection, and Monte Carlo simulation to produce defensible fair value estimates, option-adjusted spread (OAS) analysis, and a comprehensive suite of risk measures - all within a single, auditable workflow.

This white paper introduces the framework at a conceptual level, explaining why it exists, how it works, and what it delivers - without requiring deep technical expertise from the reader.

The Valuation Challenge

SBA 7(a) loan pools present a unique valuation problem that sits at the intersection of credit analysis, interest rate modeling, and structured finance. Unlike a simple government bond with fixed cash flows, an SBA pool generates payments that are inherently uncertain. Borrowers may prepay their loans early when interest rates fall, or default when economic conditions deteriorate. The pool's interest-only (IO) strip - a key source of return for many investors - is particularly sensitive to these behavioral dynamics.

Why Traditional Approaches Fall Short

Simple discounted cash flow models that assume a single prepayment speed and a fixed discount rate fail to capture the optionality embedded in these pools. In reality, prepayment and default behavior shift with the interest rate environment: when rates decline, refinancing accelerates; when the yield curve flattens, spread compression changes the economics of early payoff. A credible valuation must account for the full range of possible future interest rate scenarios and the corresponding behavioral responses.

Furthermore, accounting standards such as ASC 820 (Fair Value Measurement) require that fair values reflect market-consistent assumptions and account for the uncertainty inherent in future cash flows. This calls for a simulation-based approach that can quantify not just a point estimate of value, but also the distribution of outcomes and the precision of the estimate.

Framework Overview

The Fair Value Engine is built on a four-stage analytical pipeline. Each stage feeds into the next, creating a coherent end-to-end workflow from raw loan data to a fully documented fair value result.

Stage	What It Does	Key Output
1. Interest Rate Simulation	Generates hundreds of plausible future interest rate scenarios using a calibrated term structure model	Yield curve paths across all key maturities
2. Credit & Prepayment Modeling	Projects loan-by-loan prepayment, default, and recovery behavior under each rate scenario	Monthly credit vectors (CPR, CDR, recovery)
3. Cash Flow Projection	Rolls forward the loan pool's amortization schedule under each scenario, computing principal, interest, and IO strip cash flows	Path-level cash flow waterfalls
4. Valuation & Analytics	Discounts cash flows, computes fair value, solves for OAS, and produces risk and sensitivity measures	Fair value, OAS, duration, WAL, distributions

Stage 1: Interest Rate Simulation

The foundation of any Monte Carlo valuation is a credible model of how interest rates might evolve over the life of the asset. The framework leverages the Dynamic Nelson-Siegel (DNS) model, a widely

respected term structure model used by central banks, asset managers, and academic researchers worldwide.

The Dynamic Nelson-Siegel Model

The Nelson-Siegel model describes the shape of the yield curve at any point in time using three intuitive factors: a long-term level (where rates tend to settle over time), a slope (the difference between short- and long-term rates), and a curvature (the "hump" or "dip" at intermediate maturities). Together, these three factors can reproduce virtually any observed yield curve shape - normal, inverted, or humped.

In its dynamic extension, these factors are not held constant; instead, they evolve over time as mean-reverting stochastic processes. This means the model captures the empirical tendency of interest rates to drift back toward long-run averages while still exhibiting meaningful volatility and co-movement. The model is calibrated directly from U.S. Treasury yield curve data retrieved from the Federal Reserve, ensuring that the simulated scenarios are grounded in actual market conditions.

From a single calibration, the framework generates hundreds of correlated yield curve paths spanning the full maturity spectrum - from short-term rates like Prime and SOFR to longer-duration Treasury benchmarks at 2, 3, 5, 7, 10, 20 and 30 years. These paths serve as the foundation for every subsequent calculation.

Stage 2: Credit and Prepayment Modeling

With a rich set of interest rate scenarios in hand, the next step is to determine how borrowers within the pool will behave under each scenario. The framework supports multiple credit modeling approaches, each suited to a different level of analytical sophistication.

Available Credit Models

Model	Description	Best For
Baseline	Uses constant, user-specified CPR and CDR assumptions across all scenarios	Quick pricing, deterministic analysis
Rate-Sensitive	Adjusts prepayment and default speeds in response to rate level changes along each simulated path	Capturing rate-driven behavioral shifts
Q-SBA Model	A proprietary econometric model calibrated to historical SBA 7(a) loan performance data, incorporating yield curve spread dynamics	Production-grade fair value, regulatory reporting

The Q-SBA Model is the most advanced option and is the recommended choice for fair value determinations. It was developed by analyzing years of historical SBA 7(a) pool performance and identifying the key economic drivers of prepayment and default behavior - including the 10-year/2-year Treasury spread, which has proven to be a strong predictor of borrower refinancing activity. Because the credit model is linked to the same yield curve paths used for discounting, the framework captures the critical interdependence between rates and borrower behavior.

Regardless of which model is selected, the output is a set of monthly credit vectors - prepayment rates (SMM), default rates (MDR), and recovery rates - for every loan in the pool, under every simulated scenario. A user-adjustable PPC (Pricing Prepayment Curve) multiplier allows analysts to stress or moderate the model's base assumptions.

Stage 3: Cash Flow Projection

Armed with rate scenarios and credit vectors, the framework projects the full cash flow waterfall for the pool on a loan-by-loan, month-by-month, path-by-path basis. This is the most computationally intensive stage, and the engine is designed for parallel processing to handle large pools efficiently.

Cash Flow Components

For each period under each scenario, the engine computes the following cash flow components: remaining balance, scheduled principal amortization, voluntary prepayments, involuntary defaults, credit losses (net of recoveries applied after a user-specified lag), gross interest, servicing fees (SBA fee, FTA fee, and lender fee), and the COOF - the IO strip cash flow.

For variable-rate loans - which comprise the majority of SBA 7(a) pools - the coupon rate resets each period based on the simulated reference rate (typically Prime). This means the pool's interest cash flows are inherently path-dependent, varying across scenarios in a way that a fixed-rate model simply cannot capture.

Settlement and Accrued Interest

The framework also handles the practical mechanics of trade settlement. When a pool trades between payment dates, the buyer owes the seller accrued interest and principal from the last payment date through the settlement date. The engine computes these settlement adjustments at the loan level and applies them to the first period's cash flow, ensuring that the resulting "clean" fair value is settlement-date accurate.

Stage 4: Valuation and Risk Analytics

The final stage brings everything together. The framework discounts each path's projected cash flows back to the valuation date and computes a comprehensive set of valuation and risk metrics.

Two Pricing Modes

The framework supports two complementary pricing modes, each suited to a different analytical question:

- **Price from Discount Margin (DM):** Given a spread over the reference rate index, the engine computes the implied fair value and dollar price of the pool. This is the typical approach when an analyst has a target yield and wants to determine what the pool is worth.
- **Solve OAS from Market Price:** Given an observed market price (e.g., from a BWIC or dealer quote), the engine reverse-engineers the option-adjusted spread that equates the

Monte Carlo average NPV to the market price. The OAS solver uses a root-finding algorithm to pinpoint this spread with high precision.

In both modes, the IO strip is valued separately using a Treasury-plus-spread methodology, with the benchmark tenor and spread configurable by the analyst.

Valuation Outputs

For each valuation run, the framework produces the following:

Metric	Description
Fair Value (Clean)	Settlement-adjusted NPV of pool cash flows, averaged across all Monte Carlo paths
Price (%)	Fair value expressed as a percentage of outstanding pool balance
IO Fair Value	Separate NPV of the interest-only strip, priced over a Treasury benchmark
Dirty Price	Clean price plus accrued interest and principal through the settlement date
Implied OAS (bps)	The option-adjusted spread extracted from a given market price
Weighted Average Life	Expected time to receipt of principal, with mean and standard deviation across paths
Modified Duration	Interest rate sensitivity of the pool's value to parallel rate shifts
Effective Convexity	Second-order rate sensitivity, capturing the curvature of the price/yield relationship
Key Rate Durations	Sensitivity to rate changes at specific tenors (2Y, 5Y, 10Y, 30Y) when using DNS simulation
Convergence Metric	A 95% confidence interval on the Monte Carlo estimate, expressed as a percentage of fair value

Critically, the framework does not just produce a single number. For every valuation, it retains the full path-level detail—the NPV, price, WAL, and duration for each individual Monte Carlo scenario. This enables rich distributional analysis, including percentile breakdowns (5th, 25th, 75th, 95th), standard deviations, and visual inspection of the NPV distribution.

Sensitivity and Stress Testing

A fair value without context is of limited use. The framework includes a built-in sensitivity analysis suite that enables analysts to understand how the valuation responds to changes in key assumptions.

Spread Sensitivity

The spread sensitivity grid re-prices the pool across a range of discount margin or OAS shifts (e.g., ±25, ±50, ±100, ±150 basis points), showing how the fair value and dollar price change. This is essential for understanding the pool's spread duration—the sensitivity of value to changes in the required risk premium.

Rate Shift Analysis

The parallel rate shift analysis bumps all discount rates by a specified number of basis points and re-computes the NPV, providing a direct measure of interest rate sensitivity. This is the basis for the modified duration and convexity calculations reported in the valuation summary.

Credit Sensitivity

The PPC multiple adjustment allows analysts to stress the credit model's prepayment and default assumptions by a scalar factor. Running the valuation at, say, 75% and 125% of the base credit curve provides a clear view of how sensitive the pool's value is to faster or slower prepayment behavior - a critical consideration for IO strip investors, whose cash flows are eroded by faster prepayments.

Transparency and Audit Trail

Regulatory and accounting requirements demand that fair value determinations be transparent, reproducible, and well-documented. The Fair Value Engine is designed with this imperative in mind.

Every valuation run produces a complete audit trail that records the valuation date, settlement date, rate model used, loan rate index, pricing mode, discount margin or market price, credit model selection, PPC multiple, base CPR and CDR, recovery assumptions, IO benchmark parameters, the number of Monte Carlo paths, and the convergence metric. This information can be exported alongside the path-level results in an Excel workbook, providing a self-contained package suitable for auditor review.

The convergence metric deserves special attention. Because Monte Carlo simulation produces a statistical estimate rather than an exact analytical solution, it is important to quantify the precision of that estimate. The framework reports a 95% confidence interval on the mean price, expressed as a percentage. By increasing the number of simulated paths, analysts can tighten this confidence interval to any desired level of precision.

Integration Within the Q-SBA Analytics Platform

The fair value framework does not exist in isolation. It is deeply integrated with the broader Q-SBA Suite ecosystem, which includes tools for collateral analysis, delinquency tracking, prepayment benchmarking, and BWIC (Bids Wanted in Competition) pricing.

The Q-SBA Suite base package provides the data layer - fetching live loan-level data from a PostgreSQL database, downloading FTA payoff data from SBA sources, and retrieving current interest rate curves from the Federal Reserve. The Option-Adjusted Spread component provides the analytical engine for cash flow projection, present value calculation, and OAS solving. The Dynamic Nelson-Siegel component provides the term structure model. Together, these components form a cohesive, modular platform that can be deployed locally, on a server, or in a containerized cloud environment.

Component	Role
Q-SBA Foundation	Data acquisition, collateral analysis, rate curves, dashboard infrastructure
OAS	Cash flow projection, present value calculation, OAS solving, sensitivity analysis
DNS	Dynamic Nelson-Siegel yield curve model, calibration, stochastic simulation
Fair Value Engine	Fair value dashboard integrating all components into an interactive workflow

Who Benefits from This Framework

- **Portfolio managers and traders** gain a rigorous, defensible pricing tool that accounts for the full optionality embedded in SBA pools - eliminating the guesswork that accompanies simpler models.
- **Risk managers** benefit from the sensitivity analysis and distributional output, which provide clear measures of interest rate risk, spread risk, and prepayment risk under stressed scenarios.
- **Fund accountants and auditors** receive a transparent, fully documented fair value with a complete audit trail - satisfying ASC 820 requirements and facilitating external review.
- **Institutional investors and allocators** can use the OAS analysis to compare SBA pool opportunities on a risk-adjusted basis against other fixed income alternatives.

Conclusion

Quantome's Fair Value Engine represents a step change in how SBA 7(a) loan pools can be analyzed and valued. By combining a calibrated term structure model, econometric credit modeling, loan-level cash flow projection, and Monte Carlo simulation within a single integrated workflow, it delivers the rigor and transparency that institutional investors, regulators, and auditors demand - while remaining accessible through an interactive dashboard interface.

Whether you are pricing a new acquisition, marking an existing portfolio to market, computing hedge ratios, or preparing for an audit, the framework provides a single source of truth: a fair value grounded in market-consistent assumptions, supported by distributional analytics, and documented with a full audit trail.

About Quantome

Quantome develops comprehensive quantitative solutions -from powerful analytics and innovative data products to automated workflows and actionable market insights- to achieve superior outcomes in specialty credit and structured products markets.

Our Focus and Expertise

Data & Analytics: Comprehensive loan-level data and performance analytics service for lenders, investors, custodians and solution vendors to access standardized historical performance metrics, loan characteristics, and industry benchmarks across the specialty finance space.

Asset Valuation: Independent, data-driven valuation and pricing analytics for complex credit assets across niche lending markets including private credit, equipment leasing, and alternative consumer loans, uniquely combining proprietary models, comprehensive market comparables, and real-time performance data.

Market Intelligence: Competitive insights and trend forecasting across specialty credit and structured products markets through a unified platform aggregating proprietary performance metrics and macroeconomic indicators to provide actionable intelligence.