Doctrinal Distribution: How the Derivatives Safe Harbor Engineers the Value a Failing Firm Must Survive

J. Schidzig · Legal Quant

The U.S. Bankruptcy Code treats derivative counterparties unlike other holders of debt instruments or securities in a failing firm. Through a set of provisions known as the safe harbors, swaps and repurchase agreements escape the automatic stay, the avoidance powers, and the bar on bankruptcy-triggered termination. This paper makes two connected claims. The first is doctrinal: read across the rulings that have tested it since Lehman, the safe harbor moves value in one direction so consistently that the arrangement has the shape of an engineered outcome, recoverable from its distributional result. The second is methodological: the question the safe-harbor debate cannot answer from inside the law — what simultaneous close-out does to a coupled system — is a question about dynamics, and a stochastic engine that can switch the safe harbor off, then on, recovers its contribution to what a failing firm is worth as the difference between the two runs, not as an assertion. The legal analysis says where the value goes; the methods say what that movement does to the distribution a bankruptcy court will read.

I. Derivatives, in brief

A derivative is a bilateral contract whose value derives from an underlying: a commodity (livestock, crops, metals) or a financial reference (rates, indices, equities). It fixes the terms of a future transaction without settling it now, and in doing so transfers a defined risk from one party to another. That transfer, not delivery, is the contract's purpose: relocating exposure.

Two distinctions organize everything that follows. The first is symmetric versus asymmetric risk. Forwards, futures and swaps fix reciprocal claims at inception, so both parties carry offsetting obligations and the position can be neutralized by locking it: these are the linear, or quasi-linear, instruments. Options reverse this. The holder pays a premium for a right without a matching obligation, so rights and duties are non-linear and contingent, and the position requires dynamic hedging because its sensitivity shifts with the underlying and evolves over the position's remaining life. The distinction is not taxonomy for its own sake: it determines whether an exposure can be neutralized once, through a static lock, or must be managed continuously, and it is the same distinction the valuation engine later turns on.

The second distinction is venue. Exchange-traded derivatives are standardized and cleared: a central counterparty interposes itself, absorbs bilateral default risk, and prices are transparent. Over-the-counter derivatives are bilateral and bespoke, terms freely set, with no clearing house between the parties. The consequence is structural. In an OTC contract each party's position depends on the other's creditworthiness, so counterparty risk is not removed but held directly. That direct dependence is the channel through which one firm's distress reaches another, and it is the feature both the safe-harbor analysis and the coupling term act on.

Credit derivatives make the dependence explicit. A credit derivative pays out when a credit event, bankruptcy, failure to pay, or repudiation, affects a reference entity. The credit default swap, the dominant form, functions as insurance against a borrower's default, priced on the seller's assessment of default probability and loss given default. The credit event is the hinge that joins the two halves of this paper: a market instrument's trigger and a bankruptcy filing at once, the point where price dynamics and legal process fall on the same date.

II. Risk topology, read as structure

The risks a derivative carries are conventionally drawn from the creditor's view as a single classification: on one side the quantitative-financial risks — creditworthiness, market, and solvency; on the other the descriptive risks — systemic and legal. Drawn as a taxonomy it is a list of exposures to be aware of.

It is more useful to read the taxonomy as a claim about structure: these categories are not independent.

Creditworthiness, market volatility, and funding liquidity do not move on separate clocks. When conditions deteriorate, a counterparty's rating, the volatility of its book, and its access to short-term funding worsen together, because all three respond to the same underlying shift in market state. A single latent stress regime sits beneath the financial branch, and when it switches it raises default intensity, volatility, and liquidity cost across exposed entities at once. The financial branch is diffusive and estimable: it is the domain of the price and asset processes the engine simulates. The descriptive branch is not. Legal risk, whether documentation is enforceable, which law governs, whether a close-out clause will hold, does not diffuse and does not track a volatility regime. It resolves discontinuously, at a single moment: a filing, a ruling, a date. That difference in kind, continuous versus discontinuous, is what marks the methodological boundary between the two branches. The stochastic engine works the first branch; the legal analysis works the second, and the two meet only where a legal event enters the value path as a discontinuity.

III. The safe harbor and its distributional logic

The Bankruptcy Code's central instrument is the automatic stay. On filing it freezes collection, halts enforcement, and voids contractual clauses that would terminate an agreement on the fact of bankruptcy alone. The stay buys the estate time to reorganize and enforces a single principle, that creditors are paid by rule and in order rather than by speed.[1] Derivative counterparties stand outside that principle. The safe harbors exempt swaps, repurchase agreements, and related financial contracts from the stay, from the avoidance powers that let a trustee unwind pre-filing transfers, and from the prohibition on bankruptcy-triggered termination.[2] A swap counterparty may therefore do at once what every other creditor is forbidden to do: terminate, net its positions, and seize or liquidate collateral, the moment the debtor files.

The consequence is distributional before it is anything else. What the safe-harbored counterparty takes out of the estate is taken before the ordinary claimants are reached. The value is not destroyed; it is relocated, out of the common pool and to the holder of the protected contract. As a mechanism: the safe harbor converts a derivative exposure from a claim against the estate into a right exercised against the estate, and the gap between those two is borne by everyone whose claim remains inside.

Lehman supplied the worked example at scale. At filing it held more than 900,000 derivative contracts, a notional position estimated near US$35 trillion.[3] Most were closed out within days under the safe harbors, each counterparty netting and seizing collateral while the estate's remaining creditors stayed inside the stay. The asymmetry was not a contingency of the case. It was the safe harbor operating as written.

The decade after Lehman tested the harbor's reach, and the result is not the simple contraction it is sometimes taken to be. The clearest expansion came on the close-out right itself. In the synthetic CDO structures Lehman had written, "flip clauses" reversed the payment priority between Lehman's swap entity and the noteholders the instant Lehman filed, subordinating the very entity owed money on the terminated swap. The judge who first presided over Lehman's bankruptcy treated these as unenforceable ipso facto clauses, void because triggered by the filing itself.[4] The Second Circuit reversed course in 2020, holding the priority provisions enforceable under the section 560 swap safe harbor even where they would otherwise be void ipso facto clauses, on the ground that the harbor protects a swap participant's contractual close-out rights and is not bounded by the ordinary anti-ipso-facto rule.[5] The counterparty's bargained priority against the estate now survives the filing that triggers it.

Against that expansion sits an apparent contraction, and the relation between the two carries the argument that follows. The Supreme Court held in 2018 that the section 546(e) avoidance safe harbor protects only the transfer a trustee actually seeks to avoid, not its component steps, so that a financial institution serving as a bare conduit, with no beneficial interest, no longer pulls a transaction into the harbor.[6] Read in isolation, the holding looked like the safe harbor in retreat. But the retreat was sealed at its edges almost as soon as it opened. The same circuit whose pro-trustee ruling the Supreme Court had affirmed went on, in 2024, to read the surviving machinery broadly: "securities contract" reaches privately held stock and ordinary acquisition financing, not only public securities touching the national clearance system, and section 546(e) preempts the state-law fraudulent-transfer claims a trustee might otherwise use to recover the same value under another name.[7] The conduit exception stayed narrow; everything around it stayed wide.

Read for where the value goes rather than for doctrinal label, the rulings run one way. When the question is whether a derivative counterparty keeps value against the estate, it keeps it. When the question is whether the harbor reaches a transaction, it reaches. When the question is whether a trustee can route around the harbor to return value to the estate, the route is closed. The one point at which the harbor yields, the pure conduit with no economic stake, costs the protected class nothing, because a party with no stake was never the intended beneficiary. The contraction that looked like a reversal left the distributional arrangement untouched.

That arrangement has the shape of an engineered outcome. The priority the safe harbor confers was not derived from the first principles of bankruptcy; on the available record it was placed there, urged into the Code by the Federal Reserve, the Treasury, and the derivatives industry across successive enactments.[8] The distributional result, counterparties made whole and the estate absorbing the residual, is what the design produced: read from the outcome, the structure that generates it follows. To call that structure engineered is a claim about the statute's design and the legislative record behind it, not a claim about what any court intended. The courts in all three lines decided questions of statutory wording and believed their results were compelled by the text. Whatever engineering exists sits in the legislative drafting, not in the judging: the courts applied the statute as written; the design, if there is one, was built in before any of them read it.

Whether the safe harbor should stand is unsettled, and the terms of the dispute set up what follows. The case for the safe harbor was always systemic: exempting financial contracts from the stay was said to keep one firm's failure from cascading through its counterparties. The case against it, developed since the crisis, is that the exemption produces the cascade it claims to prevent, by letting counterparties run on a distressed firm at once and by stripping their incentive to monitor that firm in advance, since a protected counterparty comes out ahead in the resulting bankruptcy regardless.[9] Proposals to subject derivatives to a brief stay, long enough to permit an orderly resolution, have been drafted repeatedly without becoming law, and expert review has reached no agreement on whether narrowing the harbor would raise or lower systemic risk.[10] The resolution architecture built alongside the harbor stands structurally intact; the 2025 reconciliation act that revised other parts of the Dodd-Frank framework reached only its consumer-protection margin and left the resolution provisions and the safe harbors unchanged.[11] The bankruptcy-versus-resolution question the harbor poses has not closed. It recurred in 2023, when a systemically connected bank entered FDIC receivership while its holding company turned to Chapter 11.[12]

At the center of the systemic-risk dispute over the harbor sits an empirical question the statutes and the case law cannot answer from inside themselves. Does simultaneous close-out under the safe harbor dampen contagion or amplify it? Doctrine can describe the right to close out and trace where the value goes. It cannot say what the aggregate of those rights does to a coupled system under stress, because that is a question about dynamics, not entitlements. That is the question the rest of this paper takes up.

IV. The stochastic engine

The legal analysis ends at a question it cannot answer: what does simultaneous close-out do to a coupled system under stress. Answering it requires a model of how a firm's value moves, how it can fail discontinuously, and how its distress travels to the firms it is tied to. What follows completes the apparatus, adding only what the legal mechanism requires.

The single-firm baseline

A firm's value, left to ordinary market motion, drifts and diffuses. The standard representation is geometric Brownian motion: value grows at a drift rate and fluctuates with a volatility, both estimated from data, the path continuous.[13] On this baseline rests the entire classical apparatus of capital-structure stress-testing, and for the idiosyncratic, day-to-day evolution of a solvent firm it is adequate. It is also, by construction, unable to represent the thing the legal analysis identified. A continuous path cannot jump. A model whose every trajectory is unbroken cannot show a firm losing a discrete block of value in an afternoon, which is exactly what a safe-harbored close-out does to the estate it acts against.

Adding the discontinuity

The minimal correction is to let the path break. A jump-diffusion augments the continuous process with a second component: at random times, governed by an arrival intensity, the value takes a discrete step, drawn from a distribution skewed toward losses.[14] Between jumps the firm diffuses as before; at a jump it moves discontinuously. The diffusion carries ordinary fluctuation, the jump carries the discrete loss event.

The jump is the footprint of a discontinuous loss in the firm's value path — the close-out, the collateral seizure, the cross-default crystallizing at once. It is not the counterparty's decision to *run*. That decision is strategic, made by a party reading the filing and the behavior of other counterparties, and it lives in the legal analysis, not in the model. The model represents what the *run* does to the path, not why a counterparty chooses to *run*; it does not predict close-out, only the consequence of one when the legal mechanism produces it.

The jump also retires a tension the baseline carried. A process that can jump is a process with a fat tail and a discontinuous path, which is the regime the fat-tail literature describes.[15] Within this environment, the framing and the engine are aligned with each other rather than against.

The same system, more securities

A single firm with a jump can fail discontinuously on its own, but the legal question is about coincidence: close-out at one firm occurring when its counterparties are themselves under stress, recoveries impaired across the network at the same moment. The hidden-Markov layer already supplies the structure for it. A latent regime — calm or stressed, unobserved and inferred from price behavior — defines a set of conditions through its initial and transition probabilities, and securities are the emissions that meet those conditions.[16] The original version *ran* a single emission. Reading several emissions off the same regime is the coupling: one latent stress state, one set of conditions, several securities meeting them.

When the regime switches into the stressed state, volatility and jump intensity rise together across every security sharing it.

Isolating the harbor's contribution

A thick tail, on its own, does not say what put the mass there. The jump parameters are estimated from realized data, and realized data already contains whatever the safe harbor did: the close-outs that happened, the collateral that moved. An engine that reproduces that tail simulates the world with the harbor in it, and cannot by itself separate the harbor's contribution from any other source of discrete loss. To make the doctrinal claim, that the harbor relocates value into the downside, the engine has to recover that contribution rather than assert it alongside a tail it merely displays.

The jump intensity therefore splits into a baseline channel and a harbor channel, with a lever $\theta \in [0, 1]$ scaling the second:

$$\lambda = \lambda_{\mathrm{base}} + \theta \lambda_{\mathrm{H}}(s)$$

Here λ_{base} carries ordinary discrete losses — the operational and idiosyncratic discontinuities present in any firm's path regardless of bankruptcy law — while λ_{H} carries the close-out channel: the losses that fire on a filing and seize value ahead of the estate. The harbor channel does not switch on at a threshold; it activates as stress accumulates, so its intensity is a logistic function of the stress level s the shared regime carries:

$$\lambda_{\mathrm{H}}(s) = \lambda_{\mathrm{H}}^{\max} \cdot \frac{1}{1 + e^{-k(s - s_0)}}$$

with activation midpoint s_0, steepness k, and ceiling $\lambda_{\mathrm{H}}^{\max}$. The sigmoid is what matches the observed behavior of close-out losses: they do not arrive at a constant rate but cluster: rare in calm periods, accelerating through a deteriorating window, saturating in crisis. Because s is the shared regime's stress, every coupled security's harbor channel climbs together — the close-out channel is the one that fires jointly, which is what a run is. The harbor lever and the coupling are then the same mechanism seen twice: the harbor is precisely the channel that makes discontinuities arrive at once. Jumps in the harbor channel are also more severe than baseline ones, $\ln \eta_{\mathrm{H}} \sim N(\mu_{J,\mathrm{H}}, \sigma^2_{J,\mathrm{H}})$ with $\mu_{J,\mathrm{H}} < \mu_J < 0$.

The lever makes the counterfactual explicit. At $\theta = 0$ the close-out channel is absent — the firm as it would move without the safe harbor; at $\theta = 1$ the channel is at full license, the law as it stands; intermediate values are a brief-stay reform that caps how much close-out the regime ultimately permits. Running the engine at both ends with everything else held identical, the harbor's contribution to the downside is the difference in below-floor mass:

$$\Delta_{\mathrm{H}} = P_{\theta=1}(Y_T < F) - P_{\theta=0}(Y_T < F)$$

where F is the liquidation floor. This is the engine's statement of the doctrinal claim: not that a thick tail exists, but that this much of it is the harbor, recovered by differencing rather than asserted in prose.

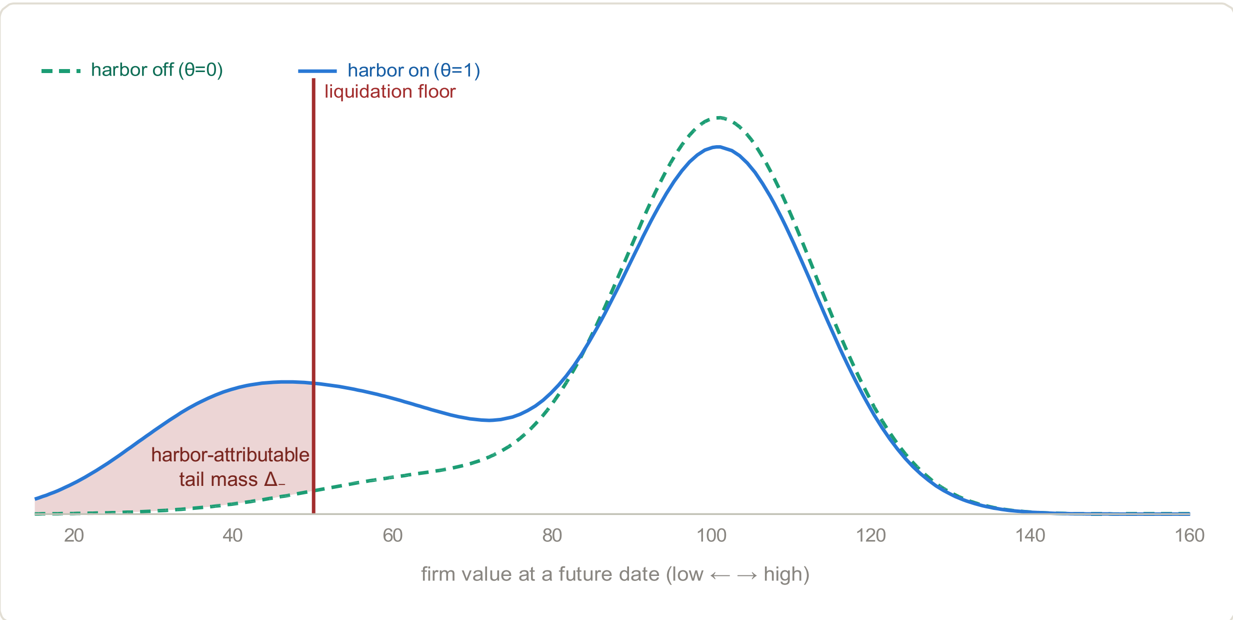

Figure 1. The harbor's contribution, recovered by differencing (illustrative). The dashed curve is the firm with the close-out channel switched off ($\theta = 0$); the solid curve is the same firm with it on ($\theta = 1$), everything else held identical. The shaded band below the liquidation floor is the difference between them, $\Delta_{\mathrm{H}} = P_{\theta=1}(Y_T < F) - P_{\theta=0}(Y_T < F)$— the value the safe harbor relocates out of the estate, read as an area rather than asserted in the text. A continuous lever between the two settings traces the same area as a function of how much close-out a brief-stay reform would permit.

That a nonzero difference

$$\Delta_{\mathrm{H}} = P_{\theta=1}(Y_T < F) - P_{\theta=0}(Y_T < F)$$

follows from toggling θ shows only that the tail moves with the lever; it does not yet show that the close-out channel specifically is what moved it, rather than the act of perturbing a jump intensity at all. A placebo lever, the term borrowed from experimental controls rather than medicine, θ_{P} answers this. It scales a channel matched to the harbor channel in intensity and severity but driven by idiosyncratic stress rather than the shared regime — a placebo that jumps as often and as hard, but not together. By construction the two leave a single firm's below-floor mass almost unchanged: matched in rate and severity, the coupled and the uncoupled channel move one firm's tail by nearly the same amount. The harbor's signature is not in any one firm's tail but across firms — in the probability that a large fraction of coupled counterparties breach the floor at once. There the shared-regime harbor far outruns the uncoupled placebo, and that excess is attributable to the harbor's distinctive feature, simultaneity, and not to discrete loss in general. The placebo isolates the common-mode contribution, which is the mechanism the doctrinal half identifies: the harbor's harm is not that close-outs happen but that they happen together.

One identification point stands plainly. The split into baseline and harbor channels is a structural decomposition, not a counterfactual read from labeled data: realized losses do not come tagged as close-outs. What licenses the split is that close-out losses have a signature — they cluster on filings, co-occur across counterparties, and run severe — and the sigmoid is estimable from the shape of that clustering, the rate at which losses accelerate as a stress proxy rises, rather than from an unobservable tag. The placebo keeps the decomposition honest by confirming the recovered difference

$$\Delta_{\mathrm{H}} = P_{\theta=1}(Y_T < F) - P_{\theta=0}(Y_T < F)$$

tracks the simultaneity that defines the harbor channel. The engine recovers the harbor's contribution under a structural decomposition validated by the placebo, not from a counterfactual observed in the world.

What the engine produces, and for whom

The engine takes a firm and its counterparties, estimates the diffusion, jump, and regime parameters from data, and simulates forward by Monte Carlo, generating a distribution of firm value at a horizon under dynamics that can jump and that deteriorate jointly when the shared regime turns.[17] The output is a firm-value distribution with a realistic downside: it carries the chance that the firm's own distress coincides with a stressed regime in which its counterparties are impaired and recoveries thin.

For the restructuring use this paper has in view, that distribution is the operative object. A debtor valuing its estate must show the court a credible firm value, and the contested cases turn on the downside. But the same distribution reads differently at different stages of the case, and the difference shows the method carries legal weight rather than legal preference.

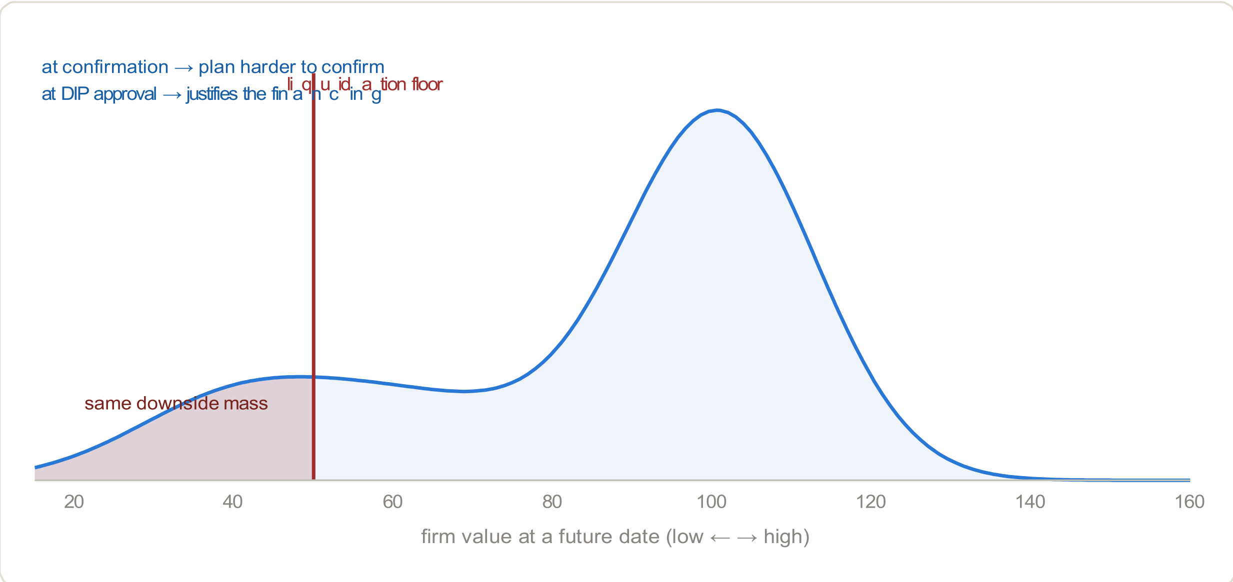

Figure 2. One stress-tested value distribution, read at two stages of the same Chapter 11 case (illustrative). The shaded region is the probability that the firm's value falls below the liquidation floor. At plan confirmation, that downside mass is a liability: it is the range in which the reorganization plan fails feasibility and best interests, so the debtor must show the mass is small. At DIP financing approval under section 364, the same mass is the justification: it establishes that the firm faces a value-destroying downside the financing is meant to bridge, and it sizes the adequate protection a priming lender requires. The engine produces a single object; the procedural question being asked decides whether its downside counts against the debtor or for it.

The two readings are not the debtor talking out of both sides. They are one distribution answering two questions. Confirmation under section 1129 asks whether the firm clears the floor often enough to keep its promises — whether the reorganization plan is feasible and leaves dissenting creditors at least as well off as liquidation would.[18] DIP financing approval under section 364 asks whether the firm faces a downside severe enough to justify the loan and whether a priming lender can be adequately protected against it. A distribution with a thick but bounded left tail can answer both questions at once, because they concern different regions of the same curve: the central mass shows the firm survives in the ordinary case, and the tail shows the downside is real enough to justify protecting a priming lender against it. The engine's value is that it produces a single object honest enough to be read at both stages without the debtor having to tell two stories. The methods reach legal consequence not by predicting a cascade but through the valuation on which a plan, and the financing that carries a firm to it, rise or fall.

V. Conclusion

The safe harbor was built to move value, and it moves it the same direction through every test the courts have put it to since Lehman. A derivative counterparty keeps its priority against the estate, the harbor reaches the transactions it touches, and the routes a trustee might use to claw value back are closed. The one place the harbor yields returns nothing to the estate, because it yields only where no real stake was ever at issue. Read for distribution rather than doctrine, the arrangement holds its shape across a decade that included a Supreme Court contraction, and that persistence is what gives it the look of something engineered to its result. The paper claims no more than that look: an outcome recoverable from the design and the record behind it, not a claim about any court's intentions.

What doctrine could trace but not resolve, the engine takes up. It answers only the part of the dynamics question that bears on a single estate. Letting the firm's value jump represents the close-out as the discrete loss it is; letting a shared regime drive a set of securities represents the moment when one firm's distress coincides with its neighbors', and making the harbor channel a lever lets its contribution to the downside be recovered by differencing the firm against itself, with close-out switched off, rather than asserted alongside a tail that merely exists. The output is a value distribution whose downside carries that coincidence, and that downside is what plan confirmation and DIP financing both turn on, in opposite directions.

That distribution is where two registers meet: legal, tracing the priority the safe harbor confers, and quantitative, showing how that priority reshapes the value a judge must weigh at confirmation and at DIP approval. Neither register is sufficient alone. The claim rests on three things: that the engineered priority has a shape indirectly traceable rather than merely asserted; that this shape carries consequences a continuous model cannot register, because reality is discontinuous at close-out; and that a model built to trace that discontinuity, switching the priority off, then on, isolates those consequences by difference, at the exact point where the law decides what a failing firm is worth.

Appendix: Stochastic background

This appendix sets out the stochastic processes the engine is built from. A stochastic process W(ω, t) is a random function appearing as the result of a random experiment. If the ω is constant, and the t is variable then we have Brownian motion in function of time. Similarly, if we fix the t and let ω vary, then we have a random variable at the moment t. The stochastic processes are modeled by stochastic differential equations because we cannot apply the classical differential calculus (Riemann) due to the property of the Brownian motion path that has no derivative, integration against this path is carried out instead by the Itô integral, examined closely below.

The jump-diffusion, the multi-security coupling, and the harbor channel that makes the close-out contribution recoverable are added here, each placed after the process it extends.

Brownian motion

A process W_t, $t \geq 0$, is Brownian motion with drift μ and volatility σ if $W_0 = 0$; the increments over disjoint intervals are independent; and $W_s - W_t \sim N(\mu(s-t),\ \sigma^2(s-t))$, so that $\mu(s - t)$ is the expectation of the increment and $\sigma^2(s-t)$ its variance. The variance grows in proportion to elapsed time.

Two features of this definition decide everything that follows. First, the increment over a small interval scales as the square root of its length: from $W_{t+\Delta} - W_t \sim N(0, \sigma^2\Delta)$, a typical increment is of order $\sqrt{\Delta}$, not Δ. Second, and directly because of this, the path is continuous everywhere but differentiable nowhere. The difference quotient

$$\frac{W_{t+\Delta} - W_t}{\Delta} \sim \frac{\sqrt{\Delta}}{\Delta} = \frac{1}{\sqrt{\Delta}} \to \infty \quad (\Delta \to 0)$$

diverges, so the ordinary derivative dW_t/dt exists at no point. The object it would formally name — the derivative of Brownian motion — is white noise: a process uncorrelated across every instant, with no value of its own at a point, defined only under an integral. It is for this reason that the increment dW_t, never a derivative, is what enters the equations below, and that integrating against it requires the construction developed next.

White noise and the Itô integral

Because the Brownian path has no derivative, a dynamic written in the ordinary way, $dY/dt = \dots$, has no meaning: there is no dW/dt to place on the right-hand side. The evolution is written instead in differential form,

$$dY_t = \mu Y_t dt + \sigma Y_t dW_t$$

and understood as shorthand for an integral equation,

$$Y_t = Y_0 + \int_0^t \mu Y_s ds + \int_0^t \sigma Y_s dW_s$$

The first integral is an ordinary Riemann integral in time. The second, taken against Brownian motion, cannot be: the Riemann construction needs its integrator to have bounded variation, and the Brownian path does not — it oscillates so violently that its total length over any interval, however short, is infinite.

The integral against dW_s must therefore be built from scratch, as a limit of sums

$$\int_0^t \sigma Y_s dW_s = \lim_{n\to\infty} \sum_k \sigma Y_{s_k} (W_{s_{k+1}} - W_{s_k})$$

with each integrand evaluated at the *left* endpoint s_k of its subinterval. The choice is not cosmetic: evaluating on the left fixes the integrand before the increment it multiplies is drawn, so the sum never anticipates the future, and the resulting process keeps the martingale property the modeling depends on. This left-endpoint construction is the Itô integral. Its use is not a matter of preference — the path is nowhere differentiable and of unbounded variation, Riemann integration is, for the reasons above, not applicable, and the Itô integral is what stands in its place.

The Itô integral obeys a calculus of its own, and the difference from the ordinary one is the entire point. In ordinary calculus a squared increment is negligible: $(dx)^2$ vanishes faster than dx and is discarded. Under Brownian motion it is not discarded, because the increment scales as √dt:

$$(dW_t)^2 \longrightarrow dt$$

The square of the increment survives, as dt. Any change of variables must carry this surviving term. For a function f(Yt) the ordinary chain rule gains a second-order correction,

$$df(Y_t) = f'(Y_t)\,dY_t + \tfrac{1}{2}\,f''(Y_t)\,(dY_t)^2$$

which is Itô's lemma; the extra half-term is the trace of $(dW)^2 = dt$, absent from ordinary calculus and present here for that reason alone.
The term is not abstract — it is already visible in the solution stated for geometric Brownian motion. Taking f(Y) = ln Y, so f′ = 1/Y and f″ = −1/Y², and using $(dYt)^2 = \sigma^2 Yt^2$ dt, the second-order piece contributes

$$\tfrac{1}{2}\Big(-\frac{1}{Y^2}\Big)(\sigma^2 Y^2)dt = -\tfrac{1}{2}\sigma^2\,dt$$

so that

$$d(\ln Y_t) = \big(\mu - \tfrac{1}{2}\sigma^2\big)dt + \sigma\,dW_t, \qquad Y_t = Y_0\exp\!\big[\big(\mu - \tfrac{1}{2}\sigma^2\big)t + \sigma\,W_t\big]$$

The $-\tfrac{1}{2}\sigma^2$ in the exponent is the Itô correction. It is present for no reason other than that $(dW)^2 = dt$ failed to vanish; a reader applying ordinary calculus would omit it and misprice the process.

That single term is the signature of the whole construction — the point at which the non-differentiability of the path, carried through the integral, reaches the number the engine reports.

The jump term introduced next demands the same care. The Poisson increment dNt, like dWt, is not an object for ordinary calculus — zero at almost every instant, one at the jump times — so the jump-diffusion too is read as an integral equation, its jump integral accumulating discrete contributions at the arrival times rather than any derivative in t.

Markov chains

A discrete process $\{X_t\}$ over a finite set of conditions $\{x_1, \ldots, x_N\}$ is a Markov chain if the condition at time t depends only on the condition at $t-1$. It is specified by the set of conditions, a transition matrix A with entries $a_{ij} = P(x_t = x_j \mid x_{t-1} = x_i)$, and a vector of initial probabilities $\pi = (\pi_1, \ldots, \pi_N)$ **with** $\pi_i = P(x_0 = x_i)$ **and** $\sum_i \pi_i = 1$ **.**

Hidden Markov chains

A hidden Markov chain carries two stochastic processes, one of which is unobserved. Let $\{X_t\}$ be a Markov chain over hidden conditions $\{x_1, \ldots, x_N\}$ — say, the market regime — and let it govern an observed price process $\{Y_t\}$ whose values y_j can be measured. Each observed value is generated by the hidden condition through emission probabilities $b_{ij} = P(Y_t = y_j \mid X_t = x_i)$. The hidden regime drives the observable prices.

HMM features:

1. Set $X = \{x_1, x_2, \ldots\}$ is a Markov chain, not directly observable.

2. $Y = \{Y_1, Y_2, \ldots\}$ is set of observations.

3. N is set dimension of Markov's chain conditions.

4. M is dimension of set of observation sequence (array).

5. $S_X = (s_1, s_2, \ldots, s_N)$ is set of Markov's chain conditions.

6. $O_Y = \{o_1, o_2, \ldots, o_N\}$ is set of observations.

7. Transition probability matrix $\Pi = \{\pi_{ij}\}, \quad i, j = 1, \ldots, N$ defined with:

$\pi_{ij} = P\left(x_{k+1} = s_j \,\middle|\, x_k = s_i\right\}$.

8. Conditional probability matrix $B = \{b_{ij}\}, \ i = 1,2,\dots,N; \ j = 1,2,\dots,M.$

$$b_{ij} = P(Y_k = o_j | x_k = s_i)$$

9. Distribution of initial state

$$\boldsymbol{A} = \begin{matrix} \boldsymbol{a_1} \\ \boldsymbol{a_2} \\ \vdots \\ \boldsymbol{a_n} \end{matrix}$$

Model for HMM is $\boldsymbol{\xi} = (\boldsymbol{\Pi}, \boldsymbol{B}, \boldsymbol{A})$

Three problems of HMM:

1. For given model $\boldsymbol{\xi} = (\boldsymbol{\Pi}, \boldsymbol{B}, \boldsymbol{A})$ and given sequence of observations $\{\boldsymbol{Y_1}, \boldsymbol{Y_2}, \dots\}$ how to effectively calculate probabilities $\boldsymbol{P} = (\boldsymbol{Y} \,|\boldsymbol{\xi})$, that is, when the model is known, what is the probability to get the particular sequence of observations. This problem is solvable applying forward – backward algorithm.

2. For a given model $\boldsymbol{\xi} = (\boldsymbol{\Pi}, \boldsymbol{B}, \boldsymbol{A})$ and sequence of observations $\{\boldsymbol{Y_1}, \boldsymbol{Y_2}, \dots\}$, what is the initial condition of the sequence that best represents observations, in other words, we are seeking best initial sequence $\boldsymbol{x_1}, \boldsymbol{x_2}, \dots$ under the given conditions. This problem can be solved using the Viterbi algorithm.

3. For a given sequence of observations $\boldsymbol{Y} = \{\boldsymbol{Y_1}, \boldsymbol{Y_2}, \dots\}$ and set of conditions $\boldsymbol{S_x} = \{\boldsymbol{s_1}, \boldsymbol{s_2}, \dots \boldsymbol{s_N})\}$, how can we adopt parameters to get the model $\boldsymbol{\xi} = (\boldsymbol{\Pi}, \boldsymbol{B}, \boldsymbol{A})$. That is, we are searching the best parameter estimates of the model $\boldsymbol{\xi} = (\boldsymbol{\Pi}, \boldsymbol{B}, \boldsymbol{A})$, which we can get using Baum – Welch algorithm.

These three algorithms solve the HMM problem.

Geometric Brownian motion

Geometric Brownian motion models a security price and projects it forward. The price process $\{Y_t\}$ satisfies the stochastic differential equation *(Itô)*

$$dY_t = \mu Y_t dt + \sigma Y_t dW_t$$

where W_t is Brownian motion with drift μ and volatility σ. Its solution is

$$Y_t = Y_0 \exp\left[\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W_t\right]$$

Since $W_t = \sqrt{t}\, Z_t$ with $Z_t \sim N(0,1)$, simulating the standard normal simulates $\exp$ and through it the process $\{Y_t\}$.

Jump-diffusion

Geometric Brownian motion has continuous paths and cannot represent a discrete loss event. The minimal extension adds a jump term driven by a Poisson process N_t of intensity λ:

$$dY_t = \mu Y_t dt + \sigma Y_t dW_t + Y_{t^-} dJ_t$$

where $J_t = \sum_{i=1}^{N_t} (\eta_i - 1)$ accumulates the jumps, $N_t \sim \mathrm{Poisson}(\lambda t)$ counts them, and Y_{t^-} is the value an instant before a jump. Each jump multiplies the value by η_i; taking $\ln \eta_i \sim N(\mu_J, \sigma_J^2)$ with $\mu_J < 0$ gives the multiplicative log-normal form, with the negative mean encoding that a close-out lands as a loss. The multiplicative jump keeps Y_t positive, as a firm value must be. Between jumps the process is the geometric Brownian motion above; the solution is that solution scaled by the accumulated jumps,

$$Y_t = \exp\left[\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W_t\right] \prod_{i=1}^{N_t} \eta_i$$

The intensity λ and the loss-skew of η are estimated from data; they are the model's representation of discrete loss events, the footprint of a close-out rather than the decision that produces it.

Hidden Markov chains with geometric Brownian motion

When the hidden regime does not change the parameters of the price process, μ and σ are constant and the logarithmic yield of the security is $r_{k+1} = \ln(Y_{k+1}/Y_k) = \mu - \frac{\sigma^2}{2} + \sigma(W_{k+1} - W_k)$. When the regime does drive the parameters, let the hidden chain $\{X_t\}$ control the price process so that each condition x_k carries its own values μ_k and σ_k. The logarithmic yield at step k is then

$$r_{k+1} = \mu_k - \frac{\sigma_k^2}{2} + \sigma_k(W_{k+1} - W_k) = f(x_k) + \sigma_k Z_{k+1}, \quad Z_k \sim N(\mu_k, \sigma_k^2)$$

For the process one defines the initial probabilities $\pi = (\pi_1, \ldots, \pi_n)$, the transition matrix A, and the regime-indexed vectors $\mu = (\mu_1, \ldots, \mu_n)$ and $\sigma = (\sigma_1, \ldots, \sigma_n)$ of the attendant Brownian motion. The process is then simulated by the Monte Carlo method.

The same regime over several securities

The preceding construction fuses one hidden regime with one price process. The regime, however, is defined over conditions through π and A; the security is an emission meeting those conditions, and nothing in the definition restricts the regime to a single emission. Let a set of securities indexed $m = 1, \ldots, M$ be emissions of one common regime $\{X_t\}$. In condition x_k, each security takes its own parameters

$$\mu_m^{(k)}, \qquad \sigma_m^{(k)}, \qquad \lambda_m^{(k)}$$

so that when the shared regime enters the stressed condition, volatility and jump intensity rise together across every security meeting that condition. The securities' idiosyncratic shocks — their Brownian increments $W_t^{(m)}$ and their own jumps — remain independent given the regime; what they share is the regime itself. This common conditioning is the coupling: the securities move together because they meet a common set of conditions defined by the regime's initial and transition probabilities, which is the model's account of how distress arrives jointly. It does not encode a transmission path from one security's default to another's — that is a different model for a different question.

The harbor channel and the recovered difference

To attribute tail mass to the safe harbor rather than to discrete loss in general, the jump intensity decomposes into a baseline channel and a harbor channel under a lever $\theta \in [0, 1]$:

$$\lambda = \lambda_{\text{base}} + \theta \lambda_{\text{H}}(s)$$

The baseline channel λ_{base} carries idiosyncratic discrete loss; the harbor channel λ_{H} carries close-out, fires through the shared regime so that it is common-mode across coupled securities, and draws the more severe jump size $\ln \eta_{\text{H}} \sim N(\mu_{J,\text{H}}, \sigma^2_{J,\text{H}})$ with $\mu_{J,\text{H}} < \mu_J < 0$. Its intensity activates with accumulated stress s as a sigmoid,

$$\lambda_{\text{H}}(s) = \lambda_{\text{H}}^{\max} \cdot \sigma(k(s - s_0)), \quad \sigma(z) = \frac{1}{1 + e^{-z}}$$

so that s_0 is the activation midpoint, k the onset sharpness, and the lever scales the ceiling $\lambda_{\text{H}}^{\max}$. Setting $\theta = 0$ removes the channel; $\theta = 1$ is full close-out; intermediate values model a partial stay. The parameters s_0 and k are estimated from the acceleration of the loss-arrival rate against a stress proxy, not from losses labeled as close-outs.

The harbor's contribution is the difference in below-floor mass between the two lever settings, holding all else fixed:

$$\Delta_{\text{H}} = P_{\theta=1}(Y_T < F) \; - \; P_{\theta=0}(Y_T < F)$$

with F the liquidation floor and Y_T the simulated value at the horizon. A placebo lever θ_{P} scales a channel matched in intensity and severity but routed idiosyncratically rather than through the regime; $\Delta_{\text{H}} \gg \Delta_{\text{P}}$ confirms the recovered difference is carried by the harbor channel's simultaneity and not by discontinuity alone. The decomposition is structural and the placebo is its validation; the harbor channel is not identified from a labeled counterfactual in the data.

Monte Carlo method

The theoretical ground of the Monte Carlo and historical simulation methods is the Central Limit Theorem: for $X_1, X_2, \ldots$ independent and identically distributed with expectation μ and variance σ^2, the standardized sum $Z_n = \left(\frac{1}{n}\sum_{i=1}^{n} X_i - \mu\right)/(\sigma/\sqrt{n})$ converges in distribution to $N(0,1)$ as $n \to \infty$.

Generating uniform variates and transforming them yields samples of the target distribution: accumulating many such samples reconstructs it. In the completed engine each simulated step draws the regime transition, then for each security its regime-conditioned diffusion and jumps, including the stress-activated harbor channel at the chosen lever setting. The distribution is generated at both lever settings, and their difference below the floor is the recovered harbor contribution the figures above read.

NOTES

1. Warren, *Chapter 11: Reorganizing American Business*. ↩
2. Skeel, *The New Financial Deal*, at 158–60. ↩
3. Valukas Examiner's Report; Fleming & Sarkar, *The Failure Resolution of Lehman Brothers*, FRBNY Econ. Pol'y Rev. (Dec. 2014), at 186. ↩
4. *In re Lehman Bros. Holdings Inc.*, 422 B.R. 407 (Bankr. S.D.N.Y. 2010) (Peck, J.). ↩
5. *Lehman Bros. Special Financing Inc. v. Branch Banking & Trust Co. (In re Lehman Bros. Holdings Inc.)*, 970 F.3d 91 (2d Cir. 2020). ↩
6. *Merit Management Group, LP v. FTI Consulting, Inc.*, 583 U.S. 366, 370 (2018). ↩
7. *Petr v. BMO Harris Bank N.A.*, 95 F.4th 1090 (7th Cir. 2024). ↩
8. Skeel, *The New Financial Deal*, at 158; Stout, *Derivatives and the Legal Origin of the 2008 Credit Crisis*, 1 Harv. Bus. L. Rev. 1 (2011); Schwarcz & Sharon, *The Bankruptcy-Law Safe Harbor for Derivatives: A Path-Dependence Analysis*, 71 Wash. & Lee L. Rev. 1715 (2014). ↩
9. Roe, *The Derivatives Market's Payment Priorities as Financial Crisis Accelerator*, 63 Stan. L. Rev. 539 (2011); Skeel, *The New Financial Deal*, at 161. ↩
10. GAO, *Financial Company Bankruptcies: Need to Further Consider Proposals' Impact on Systemic Risk*, GAO-13-622 (2013); GAO-15-299 (2015). ↩
11. CRS, *P.L. 119-21 … Provisions Related to CFPB Funding*, IN12578 (2025); CRS, IN12579 (2025) (Title III). ↩
12. SVB Financial Group, Chapter 11 petition, No. 23-10367 (Bankr. S.D.N.Y. 2023); Admin. Office of U.S. Courts, *Dodd-Frank Report* (2025). ↩
13. See Appendix (Brownian motion; geometric Brownian motion). ↩
14. Merton, *Option Pricing When Underlying Stock Returns Are Discontinuous*, 3 J. Fin. Econ. 125 (1976). ↩

15. Haas & Pigorsch, *Financial Economics: Fat-Tailed Distributions*, in *Encyclopedia of Complexity and Systems Science* 46–78 (Robert A. Meyers ed., Springer 2009). ↩
16. See Appendix (hidden Markov chains; the same regime over several securities); Erlwein, Mitra & Roman, *Hidden Markov Model–Based Scenario Generation*. ↩
17. See Appendix (Monte Carlo method). ↩
18. On valuation in plan confirmation, the feasibility standard (11 U.S.C. § 1129(a)(11)) and the best-interests-of-creditors test (§ 1129(a)(7)); on DIP financing and adequate protection, § 364. ↩

LITERATURE

Nassim Nicholas Taleb, *Skin in the Game: Hidden Asymmetries in Daily Life* (2018).

Nassim Nicholas Taleb, *Dynamic Hedging: Managing Vanilla and Exotic Options*, (1997).

Cao, L., B., *Metasynthetic Computing and Engineering of Complex Systems*, (2015)

David Skeel, *The New Financial Deal*, (2011).

Ulrike Heuser – Greipl, *Risiko Management – Beratung für Derivate* (1999).

Hull, J., White, A., *Handbook of the Economics of Finance: Credit Derivatives*, (2013).

Jürgen Kremer, *Portfoliotheorie, Risikomanagement und die Bewertung von Derivaten* (2nd Ed. 2011).

Karen L. Higgins, *Financial Whirlpools*, (2013).

Thomas Eger, Jochen Bigus, Claus Ott, Georg von Wangenheim, *Internationalization of the Law and its Economic Analysis* (2008).

Gil - Lafuente, A., M., *Fuzzy Logic in Financial Analysis*, (2005).

Maik Kleinschmidt, *Venture Capital, Corporate Governance and Firm Value*, (March, 2007).

Merkle, M., *Statistic and Probability implementations and examples* (1995).

Erlwein, C., Mitra, G., Roman, D., *Hidden Markov Model – Based, Scenario Generation for an Investment Optimization Problem*, (2011).

Roman Müller, *Finanzcontrolling* (2008).

Warren Elizabeth, *Chapter 11: Reorganizing American Business* (2008).

Christopher J. Redd, *Treatment of Securities and Derivatives Transactions in Bankruptcy, Part I*, 24-6 ABIJ 36, 37 (2005).

Titman, S., *The Modigliani and Miller Theorem and Market Efficiency*, National Bureau of Economic Research, No. 8641 (2001).

Lynn A. Stout, *Derivatives and the Legal Origin of the 2008 Credit Crisis*, Harvard Business Law Review, (2011).

Steven L. Schwarcz & Ori Sharon, *The Bankruptcy-Law Safe Harbor for Derivatives: A Path-Dependence Analysis*, 71 Wash. & Lee L. Rev (2014).

Moraux, Working paper, Université de Rennes I. (2004); François and Morellec, Journal of Business (2004); Galai, Raviv and Wiener, Journal of Banking & Finance (2007); Broadie and Kaya, Journal of Financial and Quantitative Analysis (2007).

Michael Chui, *Derivatives markets, products and participants: an overview*, IFC Bulletin No 35.

Carly Landon, *Making Assignments For The Benefit Of Creditors As Easy As A-B-C*, Fordham Urban Law Journal, Vol. 41 No. 4 (2015).

Qihe Tang (University of Iowa), *Probability analysis of Chapter 7 and Chapter 11 US Bankruptcy Code*, (March 01, 2013).

ABI: Brown, Glenn, Messersmith, Sabry, *Use of Derivatives and Probabilistic Methods in Bankruptcy, Workouts, and Insolvency Matters*, (March 14, 2016).

Israel Shaked, Paul d'Arezzo, David Plastino, *Capital Adequacy and Debt-Refinancing Assumption*, ABI Journal (Dec. 28, 2014).

Franklin R. Edwards, Edward R. Morrison, *Derivatives and the Bankruptcy Code: Why the Special Treatment?* Yale Journal on Regulation Vol. 22., (2005).