

Regime-Switching Asset Volatility and Optimal Capital Structure: A Simulation-Based Extension of the Leland Model

AUTHORS: Peter Mutua Malonza

SUMMARY: This paper extends the Leland model by incorporating regime-switching asset volatility, allowing asset volatility to follow a two-state Markov process. The study uses Monte Carlo simulations to solve for the value-maximizing coupon and compares the results to a constant-volatility benchmark.

METHOD: The research employs a Monte Carlo simulation approach, using non-proprietary parameters to model the impact of regime-switching volatility on firm value and default probability. The sample includes a range of coupon levels, and the data is generated through simulations.

KEY FINDINGS:

- The value-maximizing coupon, and thus the target leverage, is essentially unchanged between the regime-switching and constant-volatility specifications for well-capitalized firms.
- For highly leveraged firms already in financial distress, regime-switching volatility reduces firm value and increases default probability.
- Regime-switching volatility modestly increases firm value and decreases default probability for coupon levels at or below the value-maximizing level.

IMPLICATION FOR TRADING: Practitioners should consider the impact of volatility clustering on capital structure decisions, especially for firms that are already over-leveraged or in financial distress. This research suggests that volatility clustering could be a significant risk factor for such firms, potentially necessitating more frequent monitoring of volatility regime risk.

Regime-Switching Asset Volatility and Optimal Capital Structure: A Simulation-Based Extension of the Leland Model

Peter Mutua Malonza¹

¹Independent Researcher, Kitengela, Kenya

September 2026

Abstract

The classical structural trade-off theory of capital structure, following [Leland \(1994\)](#), characterises optimal leverage as the coupon level that balances the tax benefit of debt against the deadweight costs of bankruptcy, under the assumption of constant asset volatility. This assumption sits uneasily alongside the well-documented tendency of asset and equity volatility to cluster and spike during macroeconomic stress. This paper extends the Leland framework by allowing asset volatility to follow a two-state Markov regime-switching process, and solves for the value-maximising coupon numerically via Monte Carlo simulation, since no closed-form solution is available once volatility is regime-dependent. We compare the resulting firm-value and default-probability profiles to a constant-volatility benchmark calibrated to match the same long-run average variance. We find, first, that the value-maximising coupon (and hence target leverage, approximately 40% of firm value) is essentially unchanged between the two specifications at the resolution of our coupon grid – an approximate invariance result for well-capitalised firms. Second, and more strikingly, we find a statistically significant crossover in the sign of the effect: for coupon levels at or below the value-maximising level, regime-switching volatility modestly *increases* firm value and *decreases* default probability relative to the constant-volatility benchmark (differences of up to 1.8% of firm value, $t > 19$); but for highly levered firms already deep in financial distress (leverage above approximately 90%), the effect reverses, and regime-switching volatility *reduces* firm value and *increases* default probability relative to the constant-volatility benchmark ($t < -11$ at the point of maximal divergence). We interpret this as evidence that volatility clustering is a second-order consideration for firms operating near a prudently chosen target leverage, but a materially under-appreciated risk driver for firms that are already over-levered or in financial distress – with direct implications for distressed-debt risk assessment and going-concern evaluation.

Keywords: capital structure, trade-off theory, Leland model, regime-switching volatility, structural credit risk, bankruptcy costs, Monte Carlo simulation

This is a working paper / preprint. It has not been submitted to a peer-reviewed journal. All quantitative results are generated from a Monte Carlo simulation using illustrative, non-proprietary parameters; no firm-specific or proprietary data is used.

1 Introduction

The trade-off theory of capital structure, formalised by [Kraus and Litzenberger \(1973\)](#) and given a fully dynamic, structural treatment by [Leland \(1994\)](#), holds that firms choose leverage to balance the present value of the interest tax shield against the expected costs of financial distress and bankruptcy. In Leland’s model, a firm’s unlevered asset value follows a geometric Brownian motion with constant volatility, the firm issues perpetual debt at a fixed coupon, and default is triggered endogenously (or, in simplified implementations, at an exogenously specified barrier) when asset value falls too low. The coupon level – and hence leverage – that maximises total firm value is the model’s central object of interest.

A large subsequent literature has relaxed various assumptions of the original Leland model, including debt maturity structure ([Leland and Toft, 1996](#)) and, closer to the present paper’s concern, the assumption of constant macroeconomic conditions. [Chen \(2010\)](#) and [Bhamra \(Kuehn and Strebulaev\)](#) embed structural default models in economies with regime-switching macroeconomic states, showing that countercyclical risk premia and default costs can jointly help explain the levels of credit spreads and the pattern of observed leverage across firms. That literature is primarily concerned with asset pricing implications (credit spreads, risk premia) at a given, typically calibrated, capital structure. The present paper instead asks a more narrowly corporate-finance question: holding the tax and bankruptcy-cost parameters fixed, how – if at all – does introducing regime-switching *asset volatility* change the coupon level (and hence leverage ratio) that maximises firm value, and how does it change the value and default-risk consequences of leverage choices away from that optimum?

This question matters practically for two reasons. First, corporate treasurers and CFOs setting target leverage ratios typically rely on average or through-the-cycle estimates of asset or equity volatility; if volatility clustering has first-order effects on optimal leverage, such averaging could lead to systematically mis-set capital structure targets. Second, from a credit risk management and monitoring perspective – the author’s own applied background – understanding whether volatility clustering is more dangerous for well-capitalised or for already-distressed firms has direct implications for how frequently, and for which segment of a loan book or bond portfolio, volatility-regime risk should be actively monitored.

This paper makes three contributions. First, it specifies a tractable simulation-based extension of the Leland model with two-state Markov regime-switching asset volatility. Second, it numerically solves for the value-maximising coupon under both the regime-switching specification and a constant-volatility benchmark matched to the same long-run average variance, and documents that the optimal leverage ratio is, to the resolution achievable here, essentially unchanged. Third, and most substantively, it documents and explains a statistically significant sign-reversal in the effect of regime-switching volatility on firm value and default probability, occurring at a leverage level well above the value-maximising optimum – a finding with direct relevance to distressed-firm risk assessment.

2 Literature Review

2.1 Trade-off theory and the Leland model

Kraus and Litzenberger (1973) first formalised the trade-off between the tax advantage of debt and the expected costs of bankruptcy as determinants of an optimal capital structure. Leland (1994) embeds this trade-off in a continuous-time structural model in which unlevered asset value follows a geometric Brownian motion, the firm issues a perpetual bond, and the coupon level determines both the tax shield (proportional to the coupon) and, via an endogenously or exogenously determined default barrier, the probability and cost of bankruptcy. Leland and Toft (1996) extend the framework to finite debt maturity, showing that the tension between the tax benefit and distress costs also interacts with rollover risk.

2.2 Regime-switching structural models

Chen (2010) develops a structural model in which firms' cash flows and the pricing kernel are both driven by a regime-switching macroeconomic state (expansion vs. recession), showing that countercyclical risk premia, default probabilities, and default losses jointly help resolve the credit-spread puzzle and generate more realistic dynamic capital structure choices than constant-condition models. Bhamra (Kuehn and Strebulaev) similarly develop a unified regime-switching framework linking the equity risk premium and credit spreads to capital structure, again with an emphasis on time-varying macro states, generalising Leland-type structural credit risk models to a business-cycle setting. These papers focus primarily on matching asset-pricing moments (credit spread levels, equity premia) at empirically calibrated leverage. The present paper instead isolates the corporate-finance question of *where* the value-maximising leverage sits under regime-switching volatility, and how the value and default-risk consequences of over-leverage differ from a constant-volatility model matched on average variance – a comparison that, to the author's knowledge, has received less direct attention in the applied capital-structure literature.

3 Model Specification

3.1 Unlevered asset value process

Let $V(t)$ denote the (risk-neutral) unlevered asset value of the firm, evolving as

$$d \log V(t) = \left(r - \frac{1}{2}\sigma(t)^2\right) dt + \sigma(t) dW(t), \quad (1)$$

where r is the risk-free rate and $\sigma(t)$ is the (possibly regime-dependent) instantaneous asset volatility. In the regime-switching specification, $\sigma(t) \in \{\sigma_N, \sigma_S\}$ follows a two-state Markov chain (Normal/Stress) as in Section 3.2 of the companion paper on portfolio credit risk, with transition intensities calibrated to an average time between stress episodes of 5 years and an average stress-episode duration of 1 year. In the constant-volatility benchmark, $\sigma(t) \equiv \bar{\sigma}$, where $\bar{\sigma}$ is chosen so that $\bar{\sigma}^2$ equals the empirical time-average of $\sigma(t)^2$ realised along the

simulated regime path, ensuring the two specifications are compared on a matched long-run average variance basis.

3.2 Debt, default, and firm value

The firm issues perpetual debt with continuous coupon C , generating a tax shield at rate τC (where τ is the corporate tax rate) for as long as the firm remains solvent. Default is triggered the first time $V(t)$ falls to a barrier

$$V_B(C) = \kappa \frac{C}{r}, \quad (2)$$

i.e. a fixed proportion $\kappa \in (0, 1)$ of the capitalised value of the coupon, reflecting the practical requirement that a firm's asset base cover its permanent debt-service obligation by some margin before distress is triggered. At default, a fraction α of asset value is destroyed as a deadweight bankruptcy cost; the remaining $(1 - \alpha)V_B(C)$ is recovered by debtholders under absolute priority.

Total levered firm value is

$$v(C) = V(0) + \mathbb{E} \left[\int_0^{\tau_B} e^{-rt} \tau C dt \right] - \mathbb{E} [e^{-r\tau_B} \alpha V_B(C) \mathbf{1}\{\tau_B < \infty\}], \quad (3)$$

where $\tau_B = \inf\{t : V(t) \leq V_B(C)\}$ is the default time (approximated over a finite horizon in the simulation). Debt value is

$$D(C) = \mathbb{E} \left[\int_0^{\tau_B} e^{-rt} C dt \right] + \mathbb{E} [e^{-r\tau_B} (1 - \alpha)V_B(C) \mathbf{1}\{\tau_B < \infty\}], \quad (4)$$

and the leverage ratio is $D(C)/v(C)$. The value-maximising coupon is $C^* = \arg \max_C v(C)$.

4 Simulation Design

Equations (3)–(4) have a closed-form solution under constant volatility (Leland, 1994) but not under regime-switching volatility, since the default barrier is no longer hit according to a time-homogeneous first-passage law. We therefore evaluate both specifications by Monte Carlo simulation, using 60,000 independent asset-value paths, a 50-year horizon (long enough that all remaining unresolved paths carry negligible present value), and a monthly discretisation. To ensure the two specifications are compared under identical idiosyncratic randomness, the same underlying standard normal shocks are used for both the regime-switching and constant-volatility runs, with only the volatility applied to those shocks differing (*common random numbers*, a standard variance-reduction technique). Table 1 reports the parameterisation.

Given the calibrated transition intensities, the regime-switching chain spends approximately 16.4% of simulated time in the Stress state, consistent with the assumed average frequency and duration of stress spells.

Table 1: Model parameters used in the simulation

Parameter	Symbol	Value
Risk-free rate	r	5.0% p.a.
Corporate tax rate	τ	25.0%
Bankruptcy cost fraction	α	25.0%
Initial unlevered asset value	$V(0)$	100
Normal-regime volatility	σ_N	20.0% p.a.
Stress-regime volatility	σ_S	45.0% p.a.
Expected time between stress spells	–	5 years
Expected stress-spell duration	–	1 year
Empirical time-average volatility	$\bar{\sigma}$	25.8% p.a.
Default barrier proportionality	κ	0.80
Horizon / discretisation	$T, \Delta t$	50 years, monthly
Monte Carlo paths	–	60,000

5 Results

5.1 The value-maximising coupon is essentially unchanged

Table 2 reports the value-maximising coupon and associated leverage ratio under both specifications. Both specifications select the same coupon on our grid ($C^* = 2.50$), corresponding to a target leverage ratio of approximately 40% of total firm value. At this optimum, total firm value is modestly but statistically significantly higher under the regime-switching specification (107.02 vs. 106.83; a difference of 0.19, with $t = 5.37$ against the Monte Carlo standard error of the difference), driven by a combination of a slightly higher realised tax shield (9.14 vs. 9.05) and a slightly lower bankruptcy cost (2.13 vs. 2.23).

Table 2: Value-maximising coupon and outcomes at the optimum

	Regime-switching	Constant volatility
Value-maximising coupon, C^*	2.50	2.50
Total firm value at C^*	107.02	106.83
Leverage ratio at C^*	40.13%	40.16%
50-year default probability at C^*	44.64%	45.48%
PV of tax shield at C^*	9.14	9.05
PV of bankruptcy cost at C^*	2.13	2.23

5.2 A statistically significant sign-reversal away from the optimum

Figure 1 plots total firm value and 50-year cumulative default probability against the coupon level for both specifications. Two regions are apparent. For coupons at or below the value-maximising level, the regime-switching specification’s value curve lies *above* the constant-volatility curve, and its default-probability curve lies *below*: firms in this region benefit, on net, from volatility clustering, since the extended periods spent in the low-volatility Normal regime (which, by calibration, occupy the large majority of time) confer a survival benefit that more than offsets the occasional excursion into the high-volatility Stress regime. This effect peaks around

$C = 4.75$ – 5.00 , where firm value under regime-switching exceeds the constant-volatility benchmark by 1.71–1.76 (of firm value normalised to 100), with t -statistics of approximately 19.3–19.4 – among the most statistically decisive differences in the simulation.

Beyond a coupon of approximately $C \approx 6.15$ (leverage of roughly 90–94%, corresponding to a 50-year default probability in excess of 93%), the sign of the effect reverses: the regime-switching specification’s firm value falls *below* the constant-volatility benchmark, and its default probability rises *above* it. The divergence is largest around $C = 6.75$, where regime-switching firm value is 0.55 lower than the constant-volatility benchmark ($t = -16.0$).




Figure 1: Left: total firm value as a function of the annual coupon C , for the regime-switching and matched constant-volatility specifications (markers denote each specification’s value-maximising coupon). Right: 50-year cumulative default probability as a function of coupon. The vertical dotted line marks the approximate crossover point ($C \approx 6.15$) at which the sign of the regime-switching effect on both firm value and default probability reverses.

Table 3 summarises the pattern at selected coupon levels spanning both regions, together with the Monte Carlo t -statistic for the difference in firm value between specifications.

Table 3: Firm value and default probability across the leverage spectrum

C	Total firm value		Diff.	t -stat	Default probability	
	Regime	Constant			Regime	Constant
2.00	106.63	106.58	+0.06	+2.3	36.4%	36.9%
2.50 (C^*)	107.02	106.83	+0.19	+5.4	44.6%	45.5%
4.75	99.03	97.32	+1.71	+19.3	76.6%	78.7%
6.00	84.07	83.52	+0.55	+6.9	93.2%	93.7%
6.25	80.41	80.53	−0.12	−1.9	96.2%	96.2%
6.75	76.06	76.61	−0.55	−16.0	99.4%	99.1%

6 Discussion

Two practically relevant conclusions follow from these results, one reassuring and one cautionary.

The reassuring conclusion is that, for firms operating at or near a prudently chosen target leverage ratio (in this simulation, roughly 40% of firm value), volatility clustering appears

to be a second-order concern: not only does it fail to shift the value-maximising leverage ratio detectably, it modestly *helps* rather than hurts, because the low-volatility regime that dominates the time average provides a compounding survival benefit that is not fully captured by only matching average variance. This is consistent with the broader intuition that well-capitalised firms have a buffer against transient volatility spikes.

The cautionary conclusion concerns firms that are already highly levered – whether by original design or, more realistically, because asset values have fallen and covenant or rollover pressure has pushed effective leverage up. For such firms, our results indicate that regime-switching volatility is unambiguously more dangerous than a constant-volatility model calibrated to the same long-run average variance would suggest: default probability is higher, and firm (and by extension, recovery) value is lower, at every coupon level we examine above the crossover point. This has a direct bearing on the author’s applied area of interest, prudential credit risk monitoring: a distressed-debt or watchlist credit assessment that relies on a through-the-cycle or average-volatility view of a stressed borrower’s risk is, on this evidence, systematically under-stating that borrower’s default risk and over-stating likely recovery value, precisely in the segment of a loan book where the consequences of that understatement are most costly. This suggests that volatility-regime monitoring (e.g. via credit spread indices, implied equity volatility, or sector-level distress indicators) is likely to have the highest marginal risk-management value when applied to already-stressed obligors, rather than applied uniformly across an entire portfolio.

7 Conclusion and Extensions

This paper extends the Leland trade-off model of capital structure to allow for regime-switching asset volatility, solved via Monte Carlo simulation. We find that the value-maximising leverage ratio is essentially unchanged relative to a constant-volatility benchmark matched on long-run average variance, but that the sign of the effect of regime-switching on firm value and default probability reverses as leverage rises: modestly beneficial near the optimum, but materially harmful once a firm is deep in financial distress.

Several extensions would usefully generalise these findings. First, the default barrier here is exogenously specified as proportional to the coupon; an endogenous barrier derived from a smooth-pasting condition, generalised to the regime-switching setting, would remove a modelling degree of freedom and is a natural next step, likely requiring a numerical (rather than closed-form) solution to a regime-dependent free-boundary problem. Second, finite debt maturity and rollover risk, in the spirit of [Leland and Toft \(1996\)](#), would let the model speak to refinancing risk during stress regimes specifically – a channel this perpetual-debt specification cannot capture. Third, calibrating σ_N , σ_S , and the regime transition intensities to observable data (e.g. equity or credit-spread volatility indices, or a fitted hidden Markov model on historical asset returns) would convert the qualitative crossover result documented here into a quantitative, firm- or sector-specific risk assessment tool.

Disclosure: This is a working paper prepared as an independent research exercise. All quantitative results are generated from a Monte Carlo simulation using illustrative, non-proprietary

parameters, and are intended to demonstrate a qualitative mechanism rather than to represent a calibrated estimate for any specific firm or sector.

References

- Kraus, A., & Litzenberger, R.H. (1973). A State-Preference Model of Optimal Financial Leverage. *Journal of Finance*, 28(4), 911–922.
- Leland, H.E. (1994). Corporate Debt Value, Bond Covenants, and Optimal Capital Structure. *Journal of Finance*, 49(4), 1213–1252.
- Leland, H.E., & Toft, K.B. (1996). Optimal Capital Structure, Endogenous Bankruptcy, and the Term Structure of Credit Spreads. *Journal of Finance*, 51(3), 987–1019.
- Chen, H. (2010). Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure. *Journal of Finance*, 65(6), 2171–2212.
- Bhamra, H.S., Kuehn, L.-A., & Strebulaev, I.A. (2010). The Levered Equity Risk Premium and Credit Spreads: A Unified Framework. *Review of Financial Studies*, 23(2), 645–703.