# Generative Scenario Modeling for Stress Testing: Develop generative ML models (e.g., GANs, VAEs) to Simulate Realistic Stress Scenarios Reflecting Private Market Outcomes (Cash Flows, Valuations, Distributions), Informing Stress Testing and Contingency Planning

**Wanqing Feng**

Spider Management Company LLC, Richmond, VA 23226, USA

**Abstract:** Alternative investments can deliver diversification and risk premia, but for LPs the most dangerous feature in stress periods is not mark-to-market valuation drawdowns, but "same-direction cash-flow squeeze": distributions fall sharply or even freeze, while capital calls become more concentrated and more frequent, making liquidity management passive. The latent-factor cash flow model proposed by Cao and van Beek introduces systematic latent factors and heavy-tailed idiosyncratic terms into the TA framework, and links call rates, distribution rates, and growth rates to macro variables, making stress testing estimable, transmissible, and reproducible. Building on this, this paper proposes Generative Scenario Modeling (GSM): using the latent-factor model as an interpretable backbone, and under the cash flow–NAV identity and business-boundary constraints, introducing conditional GANs/conditional VAEs to learn tail dependence, nonlinear regime shifts, and scenario-cluster structure, generating a stress-scenario library of quarterly net cash flows, NAV paths, valuation adjustments, and distribution changes, for shortfall distribution estimation, commitment pacing and rebalancing evaluation, and tiered trigger-based contingency planning. The paper concludes by engaging the discussion of "whether LPs should use AI to assist decision-making," emphasizing definition consistency, confidentiality and compliance, model governance, and human review, to avoid treating AI as an automatic adjudicator. Keywords: private markets; cash flow risk; stress testing; generative models; scenario library; LP liquidity management.

***Keywords***: private markets; LP liquidity risk; latent-factor cash flow model; conditional generative models; scenario library

## 1. Introduction: Why LP stress tests often "get the direction right, but miss the cash"

### 1.1 Why cash-flow risk is more binding than valuation drawdowns

Secondary-market risk management mostly centers on price volatility because assets are tradable and hedgeable; private investing is a long chain of "commitment–call–valuation evolution–distribution return," and cash-flow timing mismatches are amplified in stress periods: exit windows contract and distributions dry up, refinancing tightens and project rescues raise follow-on funding needs, and the denominator effect lifts private-market weights, so LPs can slip from allocation management into shortfall response. Therefore, stress testing must answer "how much is missing in the worst case, for how long, and with what probability," and provide actionable guidance on buffer sizing.

### 1.2 Limitations of traditional methods

The experience approach overlays average call/distribution curves on a J-curve and is often too smooth; the rules approach relies on subjective proportional assumptions and is hard to validate; scenario replay is constrained by scarce samples and structural differences. More importantly, private-market stress is nonlinear: distributions may jump from decline to freeze, calls may jump from accelerating to concentrated bursts, and cross-fund resonance is pronounced; delayed disclosure and inconsistent definitions further raise information costs. This paper therefore argues to use an interpretable latent-factor framework to preserve business consistency, and then use generative models to expand the tail scenario space, upgrading stress testing into an executable scenario library and contingency system.

## 2. Baseline framework: Insights from the latent-factor cash flow model

### 2.1 Identities and boundaries

The model starts from the cash-flow generation mechanism: calls come from uncalled commitments, distributions come from distributable NAV, and NAV is updated under the joint effects of growth, calls, and distributions, forming the cash flow–NAV identity (e.g., $V_a = V_{a-1}(1+RG_a)+C_a - D_a$) and explicit boundaries. If scenarios violate these constraints, "impossible paths" emerge, such as distributions exceeding distributable value or calls exceeding remaining commitments; therefore, the identity should be treated as a hard rule for generation and validation.

### 2.2 Latent factors and heavy tails

At quarterly frequency, call rates and distribution rates often take 0 or boundary values, which can be modeled with a censored structure and mapped back to [0,1] through latent variables; systematic latent factors describe common drivers within an asset class and are persistent, and AR(1) is often used to capture "low distributions persist"; fund-level idiosyncratic terms are age-related and heavy-tailed, and Student-t is often used to capture the high frequency of extreme calls/extreme distributions, bringing "resonance" and "extremes" into an estimable structure and outputting percentile bands and tail probabilities that provide levers for commitment pacing and liquidity buffers.

## 3. Generative Scenario Modeling (GSM): Expanding scenario space on an interpretable backbone

### 3.1 Why generative methods are needed

For contingency planning, LPs care more about mechanistic bifurcations, such as exit freezes leading to prolonged distribution weakness, valuation lags causing risk to be realized in a concentrated way at sale or refinancing, and secondary discounts, rescue-driven follow-on capital, and distribution mechanics (fees, return-of-capital priority, carry deferral, etc.) jointly shaping net return paths; moreover, different vintages may exhibit different cadences under the same shock. A small set of parametric scenarios cannot cover these combinations, whereas generating many interpretable paths under constraints is closer to reality.

### 3.2 Three-layer architecture

The first layer is an interpretable backbone that specifies macro states and maps latent-factor shocks, while hard-coding identities and boundaries; the second layer is a generative engine that, conditional on asset class, vintage, fund age, uncalled commitment ratio, historical cadence, and macro factor paths, uses conditional GANs/conditional VAEs to generate "residuals/perturbations under backbone constraints," focusing on tail dependence and discrete regime shifts; the third layer is constraint validation, clipping call and distribution rates to boundaries and recomputing NAV via the identity, then validating from three angles—distribution coverage, correlation structure, and explanatory consistency—so that scenarios both resemble stress periods and remain feasible and conserving.

## 4. Scenario library and contingency implementation: Making stress testing executable

### 4.1 Organizing the scenario library

Macro-shock scenarios are anchored to public indicators to facilitate consistent communication; private-mechanism scenarios highlight exit freezes, valuation lags, secondary discounts, and follow-on capital; combined scenarios cover chain reactions, such as equity drawdowns→distribution drought→denominator effect lift→forced rebalancing or secondary sales→higher liquidity-event probability. The richer the scenario clusters, the more thresholds and buffer sizes can be based on probability distributions rather than subjective guesses.

### 4.2 Metrics and triggers

It is recommended to output both outcome metrics (percentile bands of net cash flows over coming quarters, the distribution of cumulative shortfalls, NAV paths and private-market weight paths) and trigger metrics (shortfalls exceeding a share of deployable liquid assets, probability of elevated calls over the next two quarters, duration of distribution weakness, etc.), and to define a three-tier trigger action list—"early warning–stress–event": the early-warning tier emphasizes increasing forecast frequency and controlling new commitments; the stress tier emphasizes activating buffers and rule-based rebalancing; the event tier emphasizes compliance-first and loss minimization. Contingency effectiveness depends on triggers being reviewable, actions being executable, and outcomes feeding back into threshold updates. It is also recommended to link scenario outputs with cash planning: translate the net outflow distribution over the next 8–12 quarters into an interval for "minimum available cash + standby financing capacity," and track deviations with monthly rolling updates; for options such as secondary sales, delaying investments, increasing liquid assets, and adjusting commitment pacing, predefine discount tolerance, approval paths, and execution order, and include trigger records in a review table to avoid ad hoc decisions and unclear accountability when stress hits. Technically, align generated residual scenarios with historical crisis windows for validation, ensuring that extreme correlation structures, zero-distribution probabilities, and call-peak frequencies are not overly optimistic, while retaining manual “red-team” checks. Mark key assumptions with versions and keep audit trails to enable traceability and verification.

## 5. The role of AI in LP decision-making: An acceleration tool, not an automatic adjudicator

AI is suitable for information structuring and scenario expansion: extracting features from GP quarterly reports, post-investment minutes, and term disclosures to reduce definition noise; and batch-generating scenarios under backbone constraints to form a searchable scenario library. AI should not replace investment conclusions when governance is absent, nor should generative outputs be treated as facts. Implementation should adhere to “controls first, expansion later,” with clear permissions and audit trails, and retain human review and version management.

## 6. Conclusions and outlook

GSM uses the latent-factor cash flow framework as an interpretable backbone, uses conditional GANs/conditional VAEs to learn tail resonance and nonlinear regime shifts, and uses identities and business boundaries as hard constraints, forming a stress-scenario library that is generable, verifiable, and reproducible, supporting LPs in quantifying shortfalls and liquidity-event probabilities and designing tiered contingency plans. Future work can expand data and unify definitions under compliance constraints, strengthen constraints on fees and distribution mechanics, and improve model risk management and audit processes to achieve traceability across data, models, scenarios, and decision records.

## References


[1] CAO W, VAN BEEK M. A latent factor cash flow model for alternative investment funds[J]. Financial Analysts Journal, 2025, 81(3): 60-75. DOI:10.1080/0015198X.2025.2489923.

[2] KUBIAK S, WEYDE T, GALKIN O, et al. MacroVAE: Counterfactual financial scenario generation via variational autoencoders[C]//Proceedings of the 6th ACM International Conference on AI in Finance (ICAIF ‘25). 2025. DOI:10.1145/3768292.3770360.

[3] LOOS A, KRAUSE A, et al. Diffusion models for correlation matrices[EB/OL]. arXiv:2410.20153, 2024-10-26[2026-01-18]. https://arxiv.org/abs/2410.20153.

[4] FLAIG S, JUNIKE G. Scenario generation for market risk models using generative adversarial networks[J]. Risks, 2022, 10(11): 199. DOI:10.3390/risks10110199.

[5] LIU Y, DEMOND A. Has liquidity dried up in private equity?[EB/OL]. (2022-07-15)[2026-01-18]. https://www.msci.com/our-solutions/alternative-investment-solutions/private-assets-research-center/has-liquidity-dried-up-in-private-equity.

[6] PAPURA D, LIU J, STAGER B, et al. US PE/VC benchmark commentary: calendar year 2024[EB/OL]. (2025-08-04)[2026-01-18]. https://www.cambridgeassociates.com/insight/us-pe-vc-benchmark-commentary-calendar-year-2024/.

[7] MACARTHUR H, BURACK R, ROSE G, et al. Private equity outlook 2024: the liquidity imperative[EB/OL]. (2024-03)[2026-01-18]. https://www.bain.com/insights/private-equity-outlook-liquidity-imperative-global-private-equity-report-2024/.

[8] MCKINSEY & COMPANY. Global private markets report 2025: braced for shifting weather[EB/OL]. (2025-05-20)[2026-01-18].

[9] INSTITUTIONAL LIMITED PARTNERS ASSOCIATION (ILPA). Capital call & distribution reporting template guidance[R/OL]. (2025)[2026-01-18].

[10] HARBOURVEST PARTNERS. 2025 midyear market outlook[EB/OL]. (2025-06-10)[2026-01-18]. https://www.harbourvest.com/insights-news/insights/2025-midyear-market-outlook/.