

Quantitative Reverse Stress Testing, Bottom Up

Claudio Albanese¹, Stéphane Crépey² and Stefano Iabichino³

October 25, 2022

Abstract

We propose a bottom-up quantitative reverse stress testing framework that identifies forward-looking fragilities tailored to a bank's portfolio, credit and funding strategies, models, and calibration constraints. Thus, instead of relying on historical events, we run a Monte Carlo simulation, and we mine those future states that contribute the most to a bank's cost of capital expressed in terms of scenario differential. We find that such an approach allows identifying both the systemic and idiosyncratic weaknesses of the bank's portfolio, with applications that include solvency risk, extreme events hedging, liquidity risk management, trading and credit limits, model validation and model risk management.

Keywords: quantitative reverse stress testing, cost of capital (KVA), model validation, model risk, trading limits, PFE.

JEL Classification: D81, G13, G28, G32.

Mathematics Subject Classification: 91B30, 91G20, 91G30, 91G40.

1 Introduction

Synthetic events feeding a model are the primary tool a risk manager has to evaluate the impact the next market event might carry. For example, sensitivity analysis employs artificial scenarios deviating from current markets by a handful of basis points (bps); chronicles inspire the creation of mock events in Historical (or market-based) Stress Testing (HST). The assumption that *the past could repeat as we recorded it* is the main drawback of HST. Even more stringent limitations affect the usage of machine learning methods to assess the fragility of a bank, as (i) difficulties in historical information

¹ Global Valuation, London.

² Université Paris Cité, Laboratoire de Probabilités, Statistique et Modélisation (LPSM), CNRS UMR 8001.

³ JP Morgan. This paper represents the opinions of the authors and it is not meant to represent the position or opinions of JP Morgan or their members.

Acknowledgement: The authors are thankful to two anonymous referees, Alex Miscampbell and Martin Lamb for their useful comments. The research of Stéphane Crépey benefited from the support of the Chair Stress Test, RISK Management and Financial Steering, led by the French Ecole polytechnique and its Foundation and sponsored by BNP Paribas.

gathering and management (e.g., echo and ghost effect) affect their training sets, (ii) historical observations dilute the information current quotes carries, undermining their risk reactivity, and (iii) their calibration routines disregard the probabilities agents assign to future possible market events.

In the aftermath of the 2007-2009 financial crisis, regulatory authorities emphasised the complementary role of *Reverse Stress Test*, which focuses on discovering tailored fragilities synthetically (see, e.g., BCBS (2009) and Fed (2012)).

Grundke and Pliszka (2018) distinguishes between qualitative and quantitative Reverse Stress Tests. A qualitative Reverse Stress Test enhances the *one scenario to stress them all* approach hallmarking HST, as it considers the peculiarities and current state of a firm (e.g., via SWOT and PESTLE analysis), customises a non-quantitative adverse event to it (e.g., a severe operational or climate crisis), and estimates its impact. A quantitative Reverse Stress Test, to which we reserve the RST acronym hereafter, uses models to generate a synthetic pool of scenarios to data-mine. Instead of assessing the loss an artificial scenario (e.g., akin to the Great Financial Crisis) created externally from a model might carry, an RST mines fragility states from the synthetic simulation pool of a model, conditional to a loss threshold set exogenously. In other words, while HST picks an artificial event from registers and estimates its re-occurrence losses by using a model, RST picks the critical loss exogenously and treats the paths a model projects as the mining archive. Although a model can not label its forward-looking projections (e.g., dot-com, euro-debt, covid or Ukrainian Crisis), it calibrates on running quotes, which reflect agents' fair price to exchange protection against one or more future events. Quoting Grundke and Pliszka (2018, Section 1), "*a qualitative approach alone would not work, or, at least, would have to be supported by quantitative elements [...] Papers on quantitative reverse stress tests are also very rare.*"

Despite the interest received from regulators, years after Grundke and Pliszka (2018), the related literature is still scarce. Montesi and Papiro (2018) and Montesi et al. (2020) took the pioneering step of tailoring the notion of (quantitative) RST to a bank's balance sheet by proposing a top-down approach. However, a top-down approach collapses a bank's portfolio, the risk factors driving it, its CSA terms, accounting and regulatory requirements into a handful of explanatory variables. We find further limitations in current RST literature, as the models used to describe the bank's loss process rely on simplistic distributions only (e.g., elliptical, beta, logistic), for which expert judgment calls (and not market observable) fix their parameters. To counterbalance their limitations, Montesi and Papiro (2018, page 23) propose a supra-entity delegated to "*run a stochastic simulation for every single financial institution with a common standard methodological paradigm*", i.e., a supra-entity devoted to RST consistency. As a result, a bank is left with one action only if it overshoots its survival threshold, i.e., to reserve cash: an expensive and sub-optimal strategy that does not protect against any particular scenario, as its usage is invariant across them all.

This paper is a contribution to quantitative Reverse Stress Testing (RST). It founds on the intuition that a coherent financial model used for business purposes expresses views around the realm of market possibilities, to which a calibration routine assigns a

probability number. We propose a *bottom-up* RST framework that identifies forward-looking fragilities tailored consistently to a bank’s credit and funding strategy, its model inventory and the prevailing market sentiment around future possibilities. Whereas current literature abstracts RST from a bank’s view on which future economic states are possible, in line with ECB (2019), we propose a theoretical framework and a computational methodology sensitive to the projected states of (i) all the market factors driving a bank’s exposure and their correlations, (ii) the risk-free mark-to-market (MtM) value of all the claims collected in its portfolio, (iii) each netting set’s collateral obligations, as expressed in CSA terms, (iv) all clients’ hazard rates, (v) their default losses, and (vi) their credit and funding valuation adjustment (CVA and FVA respectively) volatilities. We account for CVA and FVA volatilities as the 2008 global financial, 2020 Covid and 2022 Ukrainian crises proved their material impact on a bank’s P&L. A key step in building an RST is the specification of a suitable objective function. Cost of capital (KVA) sizes the cost of the lack of internal capital buffers needed for a business-as-usual operativity, a property that we believe sets it as merit function naturally.

We show in the paper how such a bottom-up RST framework fosters a “*teoria normativa per la coerenza del comportamento*” [*normative theory for the coherent behaviour*] (de Finetti (1970, Page 226, Note 1)), promoting coherent risk management of the bank’s fragilities across Lines of Defence. For example, once a state of fragility emerges, loss-driving counterparties and trades appear by deepening its mining phase, and a bank could increase its capital buffers (the only option in current RST literature), unwind, renegotiate and restructure specific trades or CSA terms, or even transfer the risk externally. By changing the mining direction (from states to counterparties), an RST covering the entire business strategy consistently can ameliorate credit limits, generating a metric that sizes a client’s potential capital exposure (PKE) on the back of the whole portfolio held, capturing credit and funding wayness too.

However, because RST anchors on models, their usage must satisfy a superior requirement to local pricing tasks, i.e., global coherency (see de Finetti (1931)) that must cover specification soundness, global calibration accuracy, and consistent usage of the paths a model generates to propagate a risk-factor among all the business units of a bank. To stress the importance of coherence in RST: let’s fix a risk factor and consider the industry-standard practise of cherry-picking a model based on its local performances in performing a narrow task (e.g., price a fixed-income derivative by matching the claim’s hedges). Performing an RST run on a bank that uses inconsistent models (e.g., Hull and White (1990) and Albanese and Trovato (2008)) to model different products (e.g., swaps and path-dependent derivatives) referencing the same risk factor (e.g., USD IR) compromises output interpretability as incoherent (or model-specific) stress clusters will surface inevitably.

A model can not label its synthetic events, and the stress conditions it generates might not have a relationship to any possible reality. We advocate invalidating the models that deem econometrically unrealistic future scenarios as possible. In other words, econometrically realistic extreme states must take the same importance as a model’s calibration and sensitivity stability performances. A bank must disallow the

entire usage of those models that mingle reality and fiction in the blueprint of their future states, regardless of their performance on local tasks.

Lastly, we argue that a coherent RST could provide a pioneering framework to quantify the impact that model risk carries to bank's survival via its model-interconnectedness channel (MIC; cf. Devasabai (2017)). In studying the impacts that the MIC carries to a bank's survival, we extend the model risk analysis presented in Albanese et al. (2021), which focused on model risk effects on MtM (and hedging) risk only.

To illustrate our methodology, we discuss a case study concerning a portfolio of approximately 2,500 counterparties, 100,000 derivative trades with credit, all G10 currencies and their foreign-exchange rate exposures. Using globally specified and calibrated models only (e.g., cf. Albanese and Vidler (2007); Albanese and Trovato (2008); Albanese et al. (2011)), we propagated 20,000 primary scenarios for market and credit factors, over 200-time points, covering 50 years. We relied on nested simulations to compute future conditional CVA and FVA P&L distributions, branching off 1,000 additional nested paths with yearly observation frequency and covering the portfolio's residual lifespan. We combine default, CVA and FVA distributions to heavy tailing the P&L distribution further. We share all the realisations of our risk factors among the entire portfolio coherently, using a single server. The computation reached a calculation time of 3 hours circa.

We organised the paper as follows: Section 2 details our theoretical framework, Section 3 unfolds the panoply of actionable insights a coherent RST carries to a bank's Lines of Defence (expanding on extreme scenario hedging in Section 3.1, credit limit amelioration in 3.2, and model validation in 3.3), and Section 4 concludes.

2 Theory of a Coherent Reverse Stress Test

2.1 Executive Summary

We start from a pool of individual scenarios ω and evaluation time-points t_i used to characterize the following processes:

- $X_{t_i}^M$, collecting the realisations of all the market risk-factors that drive the bank's expose. For examples, $X_{t_i}^M$ collects the realisation of interest rates, foreign rates, credit, and the clients' conditional probability of default (cf. Albanese and Trovato (2008), Albanese et al. (2011), and Albanese and Vidler (2007) for global models specifications and calibration routines; Albanese et al. (2002) explains the re-interpretation of de Finetti (1931) under our framework);
- $X_{t_i}^T$, tracking the realisations of the MtM process of all trades a bank has with its clients, their collateral posting obligations, exposure at default and CVA fair value (see (3)). The realisation of $X_{t_i}^T$ is conditional to the realisation of $X_{t_i}^M$;
- $X_{t_i}^P$, gathering those random variables tightened to the bank's business strategy. For example, $X_{t_i}^P$ collects FVA and cost of capital (see (4), and (5)-(8) respectively). The realisation of $X_{t_i}^P$ is conditional to the realisation of $X_{t_i}^T$.

A consistent simulation engine is only the first building block for coherent RST, as its implementation necessitates a landmark metric that can be differentiated and ranked to express preferences concerning each scenario and a loss threshold (or quantile) set exogenously. Our ranking proposal revolves around a bank's cost of capital (KVA). We chose KVA as the root vertex of our RST framework, because it allows to quantify the contribution that an individual projection carries to the bank's demand of shareholders' capital, which is sensitive to all the above risk factors.

In a nutshell, we denote the scenario-differential KVA as $\delta\text{KVA}^{(-\omega)}$, and mine fragilities by applying the following framework:

1. Compute the initial KVA value (KVA_0) over all the scenario pool, and a set of $\text{KVA}_0^{(-\omega)}$ numbers, computed using a “*scenario extraction with reintroduction*” procedure. In other words, we compute a $\text{KVA}_0^{(-\omega)}$ entry by excluding a single scenario ω iteratively;
2. Select the scenarios for which their $\delta\text{KVA}^{(-\omega)} = \text{KVA}_0 - \text{KVA}_0^{(-\omega)}$ contribution falls above a threshold (or inside a confidence interval (CI)) set exogenously, and label this set as *Stress Scenario Pool* (Ω^*);
3. Perform an event-based temporal analysis for each scenarios belonging to the Stress Scenario Pool to locate the temporal peak (t_{i*}) of the crisis (see Figure 2);
4. Data-mine the peak of a crisis to gather the simulated synthetic market condition (see Figure 3).

2.2 Economic Capital and its Cost

Hereafter, for notation simplicity, we describe our theoretical RST framework using a continuous-time and space formalism (while our approach relies on a finite specification only).

Consider a filtered probability space $(\Omega, \mathcal{A}, \mathfrak{F}, \mathbb{R})$, in which Ω is the state space, $\mathfrak{F} = (\mathfrak{F}_t)$ is a filtration over the σ -algebra \mathcal{A} , and \mathbb{R} is our reference probability measure. \mathbb{R} is a blend between physical (\mathbb{P}) and risk-neutral (\mathbb{Q}) measures, having the following properties: (i) \mathbb{R} equals \mathbb{Q} on the financial sub- σ -algebra of the entire σ -algebra \mathcal{A} and (ii) \mathbb{R} equals \mathbb{P} conditionally on the financial sub- σ -algebra (cf. Artzner et al. (2020, Proposition 2.1) for proof of \mathbb{R} 's well definiteness on \mathcal{A}). We use the risk-free asset growing at the OIS rate as a numéraire, and we present formulae assuming no collateral exchanges between the bank and its clients, an assumption relaxed in our framework's implementation. We denote by \mathbb{E}_t the $(\mathfrak{F}_t, \mathbb{R})$ conditional expectation and, for a generic loss random variable ℓ , the expected shortfall ($\mathbb{ES}_t(\ell)$) at the quantile level α as:

$$\mathbb{ES}_t(\ell) = \mathbb{E}_t[\ell \mid \ell \geq \text{VaR}_t(\ell)], \quad (1)$$

in which $\text{VaR}_t(\ell)$ is the corresponding value-at-risk, i.e.:

$$\text{VaR}_t(\ell) = \inf\{y; \mathbb{E}_t[\mathbb{1}_{\{\ell \leq y\}}] \geq \alpha\}. \quad (2)$$

We emphasise that both (1) and (2) are conditional on \mathfrak{F}_t . For $t = 0$ we just write \mathbb{E} , $\mathbb{V}a\mathbb{R}$ and $\mathbb{E}\mathbb{S}$.

We start by defining CVA as:

$$\begin{aligned} \text{CVA}_t &= \mathbb{E}_t \left[\sum_c \mathbb{1}_{\{t < \tau_c < \infty\}} (1 - R_c) (\text{MtM}_{\tau_c}^c)^+ \right] \\ &= \mathbb{E}_t \int_t^\infty \underbrace{\left[\sum_c (1 - R_c) (\text{MtM}_s^c)^+ \delta_{\tau_c}(ds) \right]}_{dD_s} \\ &= \mathbb{E}_t \int_t^\infty \left[\sum_c \mathbb{1}_{\{t < \tau_c\}} e^{-\int_t^s \lambda_\zeta^c d\zeta} \lambda_s^c (1 - R_c) (\text{MtM}_s^c)^+ ds \right], \end{aligned} \quad (3)$$

in which: MtM_s^c is the time- s risk-free MtM value of a client's netting set c , λ_s^c is the client's default intensity, τ_c his default-time, R_c his recovery rate, and δ_{τ_c} denotes a Dirac mass at time τ_c . We point out that the last identity assumes some conditional independence between τ_c and the reminder of the filtration (cf. Crépey et al. (2014, Section 13.7)).

Concerning FVA, we follow the entity-level FVA definition initiated by Albanese et al. (2015), i.e.:

$$\begin{aligned} \text{FVA}_t &= \mathbb{E}_t \int_t^\infty \underbrace{\left[\mathbb{1}_{\{s < \tau_b\}} \lambda_s^b \left(\sum_c \text{MtM}_s^c \mathbb{1}_{\{s < \tau_c\}} - \text{CVA}_s - \text{FVA}_s - \text{EC}_s \right)^+ ds \right]}_{dF_s}, \end{aligned} \quad (4)$$

in which: λ_s^b is the funding spread of the bank at time s , τ_b is the bank's default arrival time and EC is defined in (6).

We denote by L_t the bank's total accumulated negative P&L process (a positive dL_t signals a loss). To derive the loss process L , we combine default, CVA and FVA losses, assuming the bank hedges its portfolio's MtM risk actively and hedges its counterparty credit risk (CCR) exposures passively, i.e.:

$$L_t = \int_0^t (dD_s + dF_s) + \text{CVA}_t - \text{CVA}_0 + \text{FVA}_t - \text{FVA}_0, \quad (5)$$

in which dD_s and dF_s represent a shorthand notation for future default and funding losses, which terms are defined in line with (3) and (4) respectively. CVA_t and FVA_t terms in (5) express the conditional future CVA and FVA levels, which we capture via nested simulations (cf. Albanese et al. (2017)).

Economic Capital (EC) sizes the bank's capital buffer dedicated to absorbing unexpected losses and, following the fundamental review of the trading book (FRTB; cf.

BCBS (2013)), we define EC as the Expected Shortfall of the 1-year portfolio losses at the quantile level α , i.e., as:

$$EC_t = \mathbb{E}_t(L_{t+1} - L_t). \quad (6)$$

Inspired by Solvency II's risk margin, we chose KVA as the root vertex of our RST as it monetises the bank's shareholder capital at risk. Technically, following Albanese et al. (2016, 2021); Crépey (2022), we define KVA as:

$$\begin{aligned} KVA_t &= \mathbb{E}_t \int_t^\infty h(EC_s - KVA_s)^+ ds = \mathbb{E}_t \int_t^\infty h(\max(EC_s, KVA_s) - KVA_s) ds \\ &= \mathbb{E}_t \int_t^\infty h e^{-h(s-t)} \max(EC_s, KVA_s) ds, \quad t \geq 0, \end{aligned} \quad (7)$$

in which h is the (inter-temporal) hurdle rate, and $(EC_s - KVA_s)^+$ the bank's shareholder's capital at risk (also dubbed excess capital demand). Regarding the amount of capital at risk, although we could interpret EC as such directly, we note that similar to CVA and FVA, the KVA that EC generates is a loss-absorbing buffer too. Thus, at time s , a bank's capital at risk is $\max(EC_s, KVA_s)$, and *shareholder's* capital at risk is $\max(EC_s, KVA_s) - KVA_s = (EC_s - KVA_s)^+$ (as, similarly to CVA and FVA, a bank can transfer KVA to its clients). Concerning h , it synthesises the target rate of return to remunerate shareholders for their capital at risk, and the bank's management sets its level. Because a higher (resp. lower) level of h reflects a less (resp. more) conservative dividend management policy, it impacts the temporal location of the stress cluster we will mine.

From an implementation viewpoint, we highlight that (see Albanese et al. (2017)): (i) path-wise EC computations as per (6) would require *multiple* nested simulations, which we approximate with an *unconditional* version of EC, defined by replacing \mathbb{E}_t with \mathbb{E} in (1)-(2) and hereafter denoted by $EC(t)$ (instead of EC_t for its real process), (ii) Picard iterations can then be used to disentangle the FVA-EC dependency in (4)–(6), and (iii) the *max* term in (7) is not binding typically. To summarise: we evolve default and CVA/FVA losses, by branching a second (conditional) layer of Monte-Carlo paths. Default, CVA, and FVA losses mingle in EC projections, which we synthesise in the KVA number

$$KVA_0 \approx \mathbb{E} \int_0^\infty h e^{-hs} EC(s) ds. \quad (8)$$

Figure 1 depicts the histograms of the bank's yearly losses $(L_{t+1} - L_t)$, in line with (5), in which the white pins demark the (unconditional) $EC(t)$ term structure.

2.3 Scenario Differentiation

By studying EC in (6), we note that we could design our RST framework with various scenario ranking methods. Naturally, by zeroing-out different variables in (5), our RST would generate another merit function, exposing a different risk flavour. For example,




Figure 1: Histograms of the bank's yearly losses ($L_{t+1} - L_t$) and corresponding EC projections (see (5) and (6)).

by muting the dD_s and CVA terms in (5), KVA and the RST mining phase will only reflect funding risk; by silencing funding terms, only credit risk losses will surface. Hereafter, we develop our coherent RST turning-on to all the terms comprising (6).

For each scenario belonging to our finite and pre-generated simulation pool, we compute a scenario-differential KVA term ($\delta KVA^{(-\omega)}$). This is simply the difference between the initial KVA value (KVA_0), computed over the entire projection set, and another KVA number ($KVA_0^{(-\omega)}$), estimated referring to the same pool except for a single scenario ω that we nullified iteratively. In view of Section 2.2, $\delta KVA^{(-\omega)} = KVA_0 - KVA_0^{(-\omega)}$ represents the dividend accrual stream paid by the bank to raise capital buffers along the scenario ω . Because $\delta KVA^{(-\omega)}$ monetises the bank's excess capital demand related to the scenario ω , it enables a scenario-selection exercise based on preferences. A scenario ω belongs to the *Stress Scenario Pool* Ω^* iff its marginal KVA contribution falls inside a pre-set CI:

$$\delta KVA^{(-\omega)} \in [lb, ub] \Rightarrow \omega = \omega^* \in \Omega^*, \quad (9)$$

in which lb and ub are model-exogenous lower and upper bounds that haircut scenario-differential KVAs.

Lastly, we temporally investigate each scenario comprising Ω^* to scout the time-point that maximizes capital consumption, i.e., the peak of a crisis (t_{i^*} ; see Figure 2). As the peak of a crisis is identified, the RST mining phase enters into play:

$$\Omega^* \ni \omega^* \Rightarrow (t_{i^*}, X_M(\omega^*, t_{i^*})). \quad (10)$$

In the end, an RST is a map that returns the economic state linked to the realization of the variable of interest (see Figure 3).

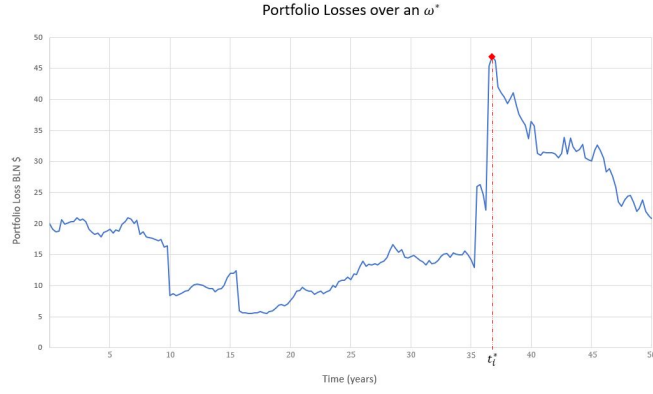


Figure 2: P&L stress scenario (the red rhomboid indicates the peak of a crisis (t_i^*)).




Figure 3: USD OIS and CAD CORRA projection for the P&L stress scenario depicted in Figure 2 (the red rhomboid indicates the peak of a crisis (t_i^*)).

2.4 Extreme Scenario Mining: A Case Study

We consider our portfolio (see Section 1), run our multi-dimensional simulation engine, set an EC quantile of 97.5% and a [99%, 99.9%] CI. Because we use a set of 20,000 primary scenarios, 180 elements comprise our Stress Scenario Pool.

Figure 4 depicts the portfolio’s fragility states exposed by our RST engine for the USD IR, from which we note a pivotal fragility cluster located in regimes of low-interest rates. Generally, we observe low-interest rates post stress regimes, reflecting established monetary policies, a cyclicity that emerges from Figure 4 (as losses materialises in low IR regimes). We also note a second cluster of losses located in higher interest rate environments from Figure 4. Central Banks tight their monetary policies in regimes of material inflation and funding costs typically, and the fragilities in these regimes signal an idiosyncratic weakness in a bank’s funding strategy, a material convexity or exposure concentration.

z

Lastly, we highlight that although we focused on IR risk, studying the uni-variate or bi-variate extreme states of other risk drivers (e.g., Equity, Commodities, Credit) does not represent a challenge.




Figure 4: Extreme USD OIS states (in bps) using the SD model to model interest rates (see Section 3.3). Blue and green lines represent the forward curve and the forward curve conditional to stress scenarios (average of the dots), respectively.

3 Coherent Risk Management

In this section, we discuss the portfolio of actionable insights that will open to a bank’s Lines of Defence if it can surface its forward-looking critical states coherently.




Figure 5: USD OIS vs. GBP SONIA overnight rates (in bps) for stress scenarios using the SD model. Connected dots represent the USD-GBP forward curve.




Figure 6: USD OIS vs. CAD CORRA overnight rates (in bps) for stress scenarios using the SD model. Connected dots represent the USD-CAD forward curve.




Figure 7: USD OIS vs. JPY JSCC overnight rates (in bps) for stress scenarios using the SD model. Connected dots represent the USD-JPY forward curve.




Figure 8: EUR EONIA vs. GBP SONIA overnight rates (in bps) for stress scenarios using the SD model. Connected dots represent the EUR-GBP forward curve.

3.1 Hedge Rare Events

Having identified the Stress Scenario Pool, a coherent RST can deepen its mining phase and reveal counterparties and trades that exacerbate excess capital demand. For example, a bank can examine its peak of a crisis (ω^*, t_{i^*}) , identifying *loss-driving netting sets* $(C^*(\omega^*, t_{i^*}))$. A netting set c is the loss-driving one ($c = c^* \in C^*(\omega^*, t_{i^*})$) if its impact by removing it from the portfolio is the highest

$$c^* = \max_{c \in C} \left[L_{t_{i^*}+1}(\omega^*) - L_{t_{i^*}}(\omega^*) - (L_{t_{i^*}+1}^{(-c)}(\omega^*) - L_{t_{i^*}}^{(-c)}(\omega^*)) \right], \quad (11)$$

with C the bank's client pool and $L^{(-c)}$ the loss process of the bank portfolio unwinding all its trades with the client c .

In Figure 9, we re-consider our portfolio and show the impact of the top five loss-driving counterparties for one of the peak of a crisis we identified in Figure 4.

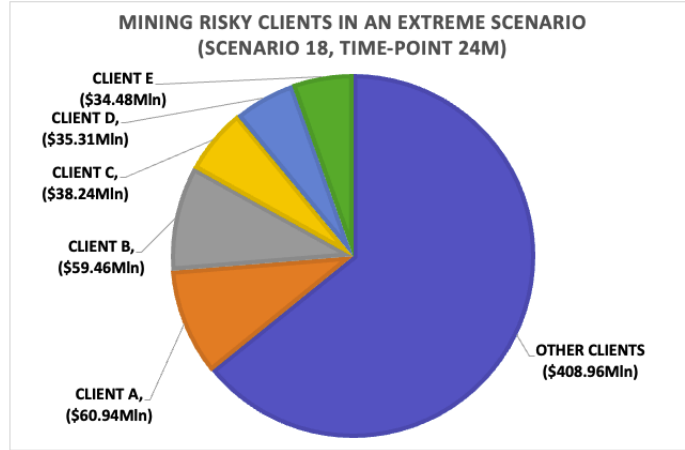


Figure 9: Top five loss-driving counterparties to a peak of a crisis.

After having identified a loss-driving client, a similar scheme repeats to pinpoint those trades that exacerbate loss concentration on the back of the portfolio held. If more than one trade emerges to be a loss-driving, then the bank identified a risk-concentration conditional to a fragility state. To hedge a risk-concentration, the bank could unwind, renegotiate or restructure specific transactions or CSA terms, or even build static hedges.

Assume our bank (B) identifies a selling bank (S) for a trade (ϵ) that might hedge its losses against one of its states of fragility. The two banks will execute the transaction iff they are able to fix a price p such that

$$\begin{aligned} \text{MtM}_0^{(\epsilon)} + (\Delta^S \text{CVA}_0^{(-\epsilon)} + \Delta^S \text{FVA}_0^{(-\epsilon)} + \Delta^S \text{KVA}_0^{(-\epsilon)}) \leq p \leq \\ \text{MtM}_0^{(\epsilon)} - (\Delta^B \text{CVA}_0^{(\epsilon)} + \Delta^B \text{FVA}_0^{(\epsilon)} + \Delta^B \text{KVA}_0^{(\epsilon)}), \end{aligned} \quad (12)$$

in which $\text{MtM}_0^{(\epsilon)}$ is the risk-free valuation of the trade (seen from B 's viewpoint), $\Delta^{B/S} \text{XVA}_0^{(\pm\epsilon)}$ is the incremental impact of the trade on the XVA metrics of banks B

and S , and we assumed no Margin Valuation Adjustment (MVA) charges between the two. Because the price p can be decoupled in the sum of the risk less MtM of the hedge and its Funds Transfer Price $\text{FTP} := \text{MtM}_0^{(\epsilon)} - p$, we have, equivalently to (12):

$$\begin{aligned} \Delta^B \text{CVA}_0^{(\epsilon)} + \Delta^B \text{FVA}_0^{(\epsilon)} + \Delta^B \text{KVA}_0^{(\epsilon)} &\leq \text{FTP} \leq \\ &- (\Delta^S \text{CVA}_0^{(-\epsilon)} + \Delta^S \text{FVA}_0^{(-\epsilon)} + \Delta^S \text{KVA}_0^{(-\epsilon)}). \end{aligned} \quad (13)$$

Assuming

$$\Delta^B \text{CVA}_0^{(\epsilon)} + \Delta^B \text{FVA}_0^{(\epsilon)} + \Delta^B \text{KVA}_0^{(\epsilon)} + \Delta^S \text{CVA}_0^{(-\epsilon)} + \Delta^S \text{FVA}_0^{(-\epsilon)} + \Delta^S \text{KVA}_0^{(-\epsilon)} \leq 0,$$

a Pareto exchange can locate in any segment of the FTP no-arbitrage price range.

From Figure 4, we identified two fragility clusters, i.e., an idiosyncratic and a systemic one. A bank could expect a wider a non-arbitrage price range for an hedge that counterbalances an idiosyncratic fragility than a systemic one and higher chances to drive FTP closer to $\text{FTP}^B := \Delta^B \text{CVA}_0^{(\epsilon)} + \Delta^B \text{FVA}_0^{(\epsilon)} + \Delta^B \text{KVA}_0^{(\epsilon)}$, given the lower market appetite.

3.2 Enrich Limits: from PFE to PKE

By changing the mining direction (from states to counterparties), an RST tailored to the entire exposure of the portfolio held can enrich Credit Limits, generating a metric that sizes a client's potential capital exposure (PKE) on the back of the entire portfolio held, capturing credit and funding wayness too.

Current risk appetite management practices revolve around metrics that already rely on extreme sub-pools of projections. For example, following BCBS (2013), banks often define PFE as an expected shortfall of a client's (c) projected MtM losses at the time horizon t . The metric credit officers typically employ to monitor the credit exposure of a client c is the maximum PFE, defined as

$$\text{PFE}^c = \max_{t \in [0, T]} \mathbb{E}S(\text{MtM}_t^c - \text{MtM}_0^c), \quad (14)$$

in which T is either the residual life of the netting set or its first year of credit exposure, MtM^c is client c 's exposure (equal to the sum of each trade's MtM comprising the client's netting set in which the ISDA does not have a CSA), and with a pre-defined quantile level α , in the [95%, 99%] range typically.

Nonetheless its usage, the industry heavily criticises PFE-like metrics as structural limitations characterise them. For example, PFE^c considers a client's MtM exposure only, which is computed under the assumption of lack of correlation between a client's and his bank's exposure.

We note that also KVA synthesises in a number an ES term structure, with the difference that its input is a client's impact on his bank's capital at risk. We define *Potential Capital Exposure* (PKE^c) as the difference between KVA_0 and another KVA number ($\text{KVA}_0^{(-c)}$), computed assuming a complete unwind of the client c . Thus, PKE^c

shifts the question from “*how much is the maximum MtM loss a client can carry to his bank?*” to “*how much would it cost the bank to sustain a client’s unexpected losses over the residual life of his portfolio?*”.

For each client c , Figure 10 compares PFE^c and PKE^c , from which we observe virtually unrelated risk profiles generated by different risk sensitivities. For example, counterparties with immaterial PKE^c can show material PFE^c , and vice-versa. The sign of the difference between a client’s PFE^c and PKE^c hints a risk-offsetting property of the netting set under consideration: A positive (resp. negative) $\text{PFE}^c - \text{PKE}^c$ difference identifies a capital offsetting (resp. absorbing) client to the reminder of his bank’s exposure that is penalised (resp. praised) by PFE^c (see Figure 10).




Figure 10: PKE^c vs. PFE^c using the SD model (see (16)). The orange lines denote the area within which the intermediary 50% of the data lie.

3.3 Invalidate Counterfeiting Reality Models

In finance, the profitability in running a model is the driving force that validates model usage; in natural science, an event that contradicts the reality described by a theory invalidates it (e.g., cf. Popper (1934)). In Albanese et al. (2021), we analysed the exchange between short-term profits and medium-term risk econometrically unsound models carry, focusing on MtM risk only. However, a bank’s model-interconnectedness channel (MIC) represents a paramount source of model risk, which could severely veil the risk/profitability analysis of a trade. We elaborate on turning a coherent RST framework into a model validation test-bench, that covers all the usages of a model as id customises to its interconnectedness.

We conclude that a bank must disallow the entire usage of those models that mingle reality and fiction in allocating their future states, regardless of their performance on

local tasks. In other words, a model validator could use a coherent RST framework to study the econometric soundness of the blueprint of the future a model embeds.

Because we could imagine that our bank doubles its positions, choosing different IR models to describe the USD IR dynamic in the new and old portfolio, we will shed light on the practice of having a kit of models that differ in their specification and end-usage and not on the risk factor they model. To study the MIC effects, we re-consider the coherent risk-management actions discussed above, considering two perfect substitute IR models that calibrate the same volatility surface globally and accurately but differ on the risk factor domain, as only one uses a cap to force it to verify econometric bounds.

Two IR Models, One Reality In the counterparty credit risk space, the Hull and White (1990) 2-factor (HW) model is a well-accepted industry standard thanks to its implementation ease and low computational resources used to calibrate its parameters on vanilla fixed-income products. To ease implementation by gaining traceability, Hull and White (1990) chose to stylise the IR randomness via an unbounded Gaussian process.

The stochastic differential equation (SDE) driving the HW model for a short-rate process r has the following form:

$$\begin{cases} dr_t = k(\theta_t - r_t)dt + \sigma(t)dW_t^{(1)} \\ d\theta_t = \kappa(a - \theta_t)dt + \nu dW_t^{(2)}, \end{cases} \quad (15)$$

in which k and κ are mean-reversion rates, θ_t and a mean-reversion levels, $\sigma(t)$ and ν volatilities, $W^{(1)}$ and $W^{(2)}$ correlated Brownian motions such that $d\langle W^{(1)}, W^{(2)} \rangle_t = \gamma dt$. To speed-up computational performances, banks often chose to constrain the calibration of $\sigma(t)$ to the backbone of the reference IR volatility surface only.

A model derived from Hull and White (1990) is Albanese and Trovato (2008), a.k.a., the stochastic drift (SD) model. The SD model starts from the same two mean-reverting factors characterising (15), augmenting it as

$$\begin{cases} dr_t = \kappa(\theta_t - r_t)dt + \sigma(t)r_t^\beta dW^{(1)} \\ d\theta_t = k(a - \theta_t)dt + \nu dW^{(2)} \\ \phi(t) = \min(0, f(t) - \psi) \\ r_t = \phi(t) + \lambda(t)r_t, \end{cases} \quad (16)$$

in which $\lambda(t)$ is a drift adjustment factor, $f(t)$ is the (infinitesimal) forward rate, and ψ is a bound calibrated econometrically.

If we compare the HW and SD models (15) and (16), we observe that Albanese and Trovato (2008) bounds the IR process to take non-negative values only, apart from the handful of basis points allowed by ψ . Moreover, a shifted exponential constrains the functional form of $\sigma(t)$, which calibrates to the entire IR volatility surface and β can be adjusted to enhance accuracy.




Figure 11: The HW model's extreme USD OIS states (in bps).




Figure 12: The SD model's extreme USD OIS states (in bps).

Distorted Risk Appetites We start our study from Figures 11-12, which put vis-à-vis the extreme scenarios mined by using HW and SD. The blue and green lines represent the forward curve and the forward curve conditional on stressing scenarios (average of the dots), respectively. Figure 11 shows that HW locates the main fragility cluster both in the short time-horizons and in conjunction with extremely low-interest-rate regimes, in the unrealistic range between -100 and -350 bps.

As visible on Figure 12, the SD model renegades the short-run deep fragilities signalled by the HW model, located in extremely low-interest-rate regimes, and disperses the risk more homogeneously over the lifetime of the book while skewing only slightly it on shorter maturities. Although the SD model recognises that the USD rate could approach a negative territory in the downturn of an economic cycle, it manages the deepen of its access by ψ , fixed at the level of -20bps in our case study.

The difficulties to interpret the RST output if a bank models a risk-factor incoherently (e.g., picking models based on the local task at hand) are well visible by comparing Figures 11 and 12.

The Hull and White (1990) model played a pivotal role in bypassing the technological boundaries that were constraining financial growth, thanks to its underneath Gaussian hypothesis. However, the financial growth it promoted turned the model's biggest strength in its Achille's heel. HW's Achille's heel affect a bank's survival by biasing the effectiveness of his Lines of Defence as follows:

- **From IR models to XVA Billing and Fair-Value:** Table 1 compares the (time-0) XVA metrics computed with the HW and SD model, from which we observe the bias that a model that is deemed as sound to perform constrained tasks could generate to the fair value of the XVA. In other words, the model risk embedded in the blueprint of the future HW embeds propagates to an XVA price unfairness for either the client or the bank's shareholders via the MIC.

The material impact of a model's state space is well visible from Figure 13, which depicts the EC projected by the two models. The different view the models have on the UDS IR domain affects materially the bank's EC profile (see (6) and Figure 13). Because HW does not constrain the realm of possibilities in its blueprint of

future possibilities, its calibration routine could potentially assign a probability to the most unrealistic future state, exacerbating capital demand.

- **Hedge Rare Events:** One of the most trivial arbitrage opportunity is the situation in which *a client is willing to buy a ticket for a lottery that does not exist* (cf. de Finetti (1931)).

By only deeming a future impossible state as possible, a model biases the natural domain of the target risk factor, jeopardising risk appetite and optimal hedging strategies. For example, an ill-intentioned agent can easily exploit the distorted risk appetite of an HW model-user if he receives the biased signal of trading protection against fictitious states (e.g., against the cluster of short-run fragility that materialises in conjunction with a catastrophically deep negative rate regime in Figure 11).

- **Credit Limits:** To analyse the distortion a model’s imagined reality generates to risk appetites, we reconsider the credit limit analysis proposed in Section 3.3, chaining the IR model from SD to HW, generating Figure 14 and 15. Although PFE relies on the extreme MtM losses haircut performed by (14), it has low sensitivity to the model of choice (as its realisations scatter roughly along the diagonal). Contrarily, PKE is highly sensitive to the field of possibilities the underneath IR model deems as possible. The different sensitivity to the radically different nature of rare events confirms the thesis that a client could generate losses that are more material than those originated from the MtM value of his portfolio, an eventuality that credit limit metrics built on canonical PFEs can not size.
- **Model Validation:** The narrow economy often stylised by a model conflicts with the growing sophistication of MIC and clients’ risk profiles, veiling the effects the thin demarcation line between reality and fiction a model draws has on a bank’s survival. For example, the real applicability of a model that behaves well under bespoke usages (e.g., price a fixed-income derivative by matching the claim’s hedges accurately), but mingles reality and fiction in the blueprint of its states, could cascade to different metrics and to the bank’s effectiveness of his Lines of Defence unforeseeably.

Thus, we hint to invalidate the models that deem econometrically unrealistic future states as possible directly, instead of relying on a model reviewer’s knowledge and expertise to identify the hidden structural limitations of a model. A coherent RST framework can elevate current model validation practices, complementing the pillars of ensuring a good fit to the claim’s hedges and a stability in its sensitivities with a model’s extreme events soundness.

Remark. *The combination of sound model specification and accurate calibration to the entire volatility surface do not determine realistic models uniquely. If a bank deems as econometrically sound more than one model, the residual model risk can be accounted*

	HW Model	SD Model	Model Risk
CVA	242	248	2.5%
FVA (RHO)	126	109	-13%
FVA (RHO + EC + CVA + FVA)	62	45	-27%
KVA	275	388	41%

Table 1: Comparison of the time-0 XVA metrics (in million USD) between the SD and HW models. FVA (RHO) accounts only for the re-hypothecation option (i.e., ignores the “-” terms in (4)), where FVA (RHO + EC + CVA + FVA) deems reserve and economic capital as a fungible source of funding.




Figure 13: Comparison of the economic capital projections computed with the HW (blue) and the SD (orange) models.




Figure 14: PFE_c obtained with the HW model (x-axis) versus the same metric obtained with the SD model (y-axis). Note that the regression line (in orange) has an angular coefficient very close to 1, indicating a lack of sensitivity of the PFE metric.

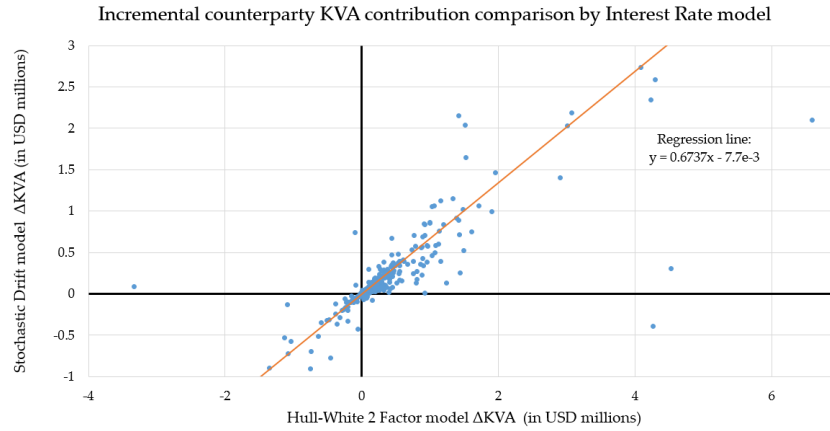


Figure 15: PKE_c obtained with the HW model (x-axis) versus the same metric obtained with the SD model (y-axis). Note that the regression line (in orange) now has an angular coefficient relatively far from 1, indicating a systematic bias due to model risk of the KVA metric.

for through Bayesian averages. A bank can run Bayesian averages in the vein of Black and Litterman (1991, 1992), using EC projections and KVA obtained by mixing (on a path) the realisations obtained using different sound models co-calibrated to the same surface globally (e.g., similarly to the VaR case in Siu et al. (2004)). Because different dynamics could describe the future evolution of a risk factor over an econometrically sound grid, a bank could use a Bayesian average, over the multi-dimensional set of realisations generated by econometrically sound models, to reserve model risk capital buffers soundly.

4 Conclusion

Artificial scenario generation is the main tool a risk manager has to estimate the P&L impact the next market event might carry. Instead of estimating losses by feeding a model with an artificial scenario akin to current markets or historical achieves, we proposed a bottom-up quantitative Reverse Stress Testing framework that identifies forward-looking fragilities tailored consistently to a bank’s credit and funding strategy, its model inventory and the prevailing market sentiment around future possibilities. In summary, the requirements for a coherent Reverse Stress Framework are the possibility to:

1. reflect the risk of the portfolio held faithfully, both in terms of trades, ISDA, and CSA terms;
2. reflect the bank’s modelling choices, which must verify no-arbitrage requirements (cf. de Finetti (1931));
3. bound the subjectivity in the parameter marking of the processes used to model the elementary drivers of a bank’s loss, using market observable and global calibration routines;
4. ensure consistent usage of the paths generated by the model chosen to propagate a risk factor.

Such a bottom-up approach fosters “*a normative theory for the coherent behaviour*” (de Finetti (1970, Page 226, Note 1)), i.e., a coherent risk management across the Lines of Defence of a bank.

References

- Albanese, C., L. Andersen, and S. Iabichino (2015). FVA: Accounting and Risk Management. *Risk Magazine*, 64–68.
- Albanese, C., T. Bellaj, G. Gimonet, and G. Pietronero (2011). Coherent global market simulations and securitization measures for counterparty credit risk. *Quantitative Finance* 11-1, 1–20.

- Albanese, C., S. Caenazzo, and S. Crépey (2016). Capital and funding. *Risk Magazine*, May 71–76.
- Albanese, C., S. Caenazzo, and S. Crépey (2017). Credit, funding, margin, and capital valuation adjustments for bilateral portfolios. *Probability, Uncertainty and Quantitative Risk* 2(7), 26 pages.
- Albanese, C., G. Campolieti, P. Carr, and A. Lipton (2002). Black–Scholes goes hypergeometric. *Risk Magazine*.
- Albanese, C., S. Crépey, R. Hoskinson, and B. Saadeddine (2021). XVA analysis from the balance sheet. *Quantitative Finance* 21(1), 99–123.
- Albanese, C., S. Crépey, and S. Iabichino (2021). A Darwinian theory of model risk. *Risk Magazine*, July pages 72–77.
- Albanese, C. and M. Trovato (2008). A stochastic monetary policy interest rate model. *Computational Methods in Financial Engineering*, 343–392.
- Albanese, C. and A. Vidler (2007). A structural model for credit-equity derivatives and bespoke CDOs. *Wilmott Magazine*, May.
- Artzner, P., K.-T. Eisele, and T. Schmidt (2020). No arbitrage in insurance and the QP-rule. Working paper available as arXiv:2005.11022.
- BCBS (2009). Principles for sound stress testing practices and supervision. <https://www.bis.org/publ/bcbs155.pdf>.
- BCBS (October, 2013). Fundamental review of the trading book: A revised market risk framework. *Basel Committee on Bank Supervision*. Consultative document.
- Black, F. and R. Litterman (1991). Asset allocation: combining investor views with market equilibrium. *The Journal of Fixed Income* 1(2), 7–18.
- Black, F. and R. Litterman (1992). Global portfolio optimization. *Financial Analysis Journal* 48(September), 28–43.
- Crépey, S. (2022). Positive XVAs. *Frontiers of Mathematical Finance* 1(3), 425–465.
- Crépey, S., T. R. Bielecki, and D. Brigo (2014). *Counterparty Risk and Funding: A Tale of Two Puzzles*. Taylor & Francis, New York. Chapman & Hall/CRC Financial Mathematics Series.
- de Finetti, B. (1931). Sul Significato Soggettivo della Probabilità. *Fondamenta Mathematicae* 17(1), 298–329.
- de Finetti, B. (1970). Teoria delle probabilità. *Enaudi*.
- Devasabai, K. (2017). Model risk managers grapple with interconnectedness. *Risk*.

- ECB (2019). ECB guide to internal models. *European Central Bank*.
- Fed (2012). Guidance on stress testing for banking organizations with total consolidated assets of more than \$ 10 billion. <https://www.federalreserve.gov/supervisionreg/srletters/sr1207a1.pdf>.
- Grundke, P. and K. Pliszka (2018). A macroeconomic reverse stress test. *Review of Quantitative Finance and Accounting* 50(4), 1093–1130.
- Hull, J. and A. White (1990). Pricing interest rate derivative securities. *Review of Financial Studies* 4, 573–592.
- Montesi, G. and G. Papiro (2018). Bank stress testing: A stochastic simulation framework to assess banks’ financial fragility. *Risks* 6 (3), 82.
- Montesi, G., G. Papiro, M. Fazzini, and A. Ronga (2020). Stochastic optimization system for bank reverse stress testing. *Journal of Risk and Financial Management* 13(8), 174.
- Popper, K. R. (1934). *The Logic of Scientific Discovery*. Hutchinson.
- Siu, T., H. Tong, and H. Yang (2004). On Bayesian value at risk: From linear to non-linear portfolios. *Asia-Pacific Financial Markets* 11(02), 161–184.