

South African Reserve Bank

Working Paper Series

WP/26/06

Event-aware jump-diffusion for the JIBAR–ZARONIA spread

Mesias Alfeus

Authorised for publication by Konstantin Makrellov

19 February 2026



SOUTH AFRICAN RESERVE BANK

© South African Reserve Bank

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without fully acknowledging the author(s) and this Working Paper as the source.

South African Reserve Bank Working Papers are written by staff members of the South African Reserve Bank and, on occasion, by consultants under the auspices of the South African Reserve Bank. The papers deal with topical issues and describe preliminary research findings and develop new analytical or empirical approaches in their analyses. They are solely intended to elicit comments and stimulate debate.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the South African Reserve Bank or South African Reserve Bank policy. While every precaution is taken to ensure the accuracy of information, the South African Reserve Bank shall not be liable to any person for inaccurate information, omissions or opinions contained herein.

South African Reserve Bank Working Papers are externally refereed.

Information on South African Reserve Bank Working Papers can be found at <https://www.resbank.co.za/en/home/publications/Papers/working-papers>.

Enquiries relating to the Working Paper Series can be addressed to:

Head: Economic Research Department

South African Reserve Bank

P O Box 427

Pretoria 0001

Tel. +27 12 313 3911

Event-aware jump-diffusion for the JIBAR–ZARONIA spread

Mesias Alfeus*

Abstract

We analyse the spread between JIBAR and ZARONIA and develop a model that allows the spread to evolve smoothly on most days but also to jump on economically important dates. The model combines two sources of information: scheduled Monetary Policy Committee (MPC) announcement dates and unusually large day-to-day moves identified directly from the data. Because MPC announcements occur intraday while the spread is observed at end of day, we link each announcement to the relevant daily observation using a simple timing rule that accounts for market close, weekends and holidays. We estimate the model with methods that allow volatility to be higher when the spread is higher, and we then simulate the fitted dynamics to quantify downside risk and future exposure. Empirically, the spread tends to move back toward a lower long-run level, the average additional movement on MPC dates is small once this return-to-normal behaviour is taken into account, and extreme outcomes remain possible when shocks cluster. Large, persistent shifts are more likely to coincide with policy dates than short-lived spikes. Overall, the approach is transparent and easy to replicate, and it provides practical inputs for fallback design, discounting choices and risk management during South Africa's transition away from JIBAR.

JEL classification

C63, E52, G13, G18

Keywords

Benchmark transition, JIBAR, ZARONIA, CIR model, jump detection, VaR, potential future exposure, South Africa

* Associate Professor, Department of Statistics and Actuarial Science, Stellenbosch University, Cape Town. Email: mesias@sun.ac.za.

The author thanks Prof. Dr. Thorsten Schmidt for his inspiration and guidance, as well as the two anonymous reviewers for their valuable feedback and constructive suggestions, which substantially improved the quality of this work. The author also thanks Donald Powers for carefully editing this paper. The author has no financial or nonfinancial disclosures to share for this paper.

1. Introduction

Global benchmark reform has replaced panel-based Interbank Offered Rates (IBORs) with transaction-based overnight rates to improve representativeness, governance and robustness. In South Africa, the unsecured term benchmark Johannesburg Interbank Agreed Rate (JIBAR) will be retired and replaced by ZARONIA, a transactions-based overnight index. On 3 December 2025 the South African Reserve Bank (SARB) announced that JIBAR will cease after its final publication on 31 December 2026 and the Market Practitioners Group (MPG) issued final recommendations for fallbacks and credit-adjustment spreads. This transition reshapes discounting, hedging, collateral and margining and the treatment of legacy contracts, making the JIBAR–ZARONIA *spread* a central state variable for pricing and risk management.

The spread isolates the economic difference between the legacy term rate and the new overnight benchmark. Modelling it targets the transition mechanism directly – rather than conflating it with the general level of rates – and informs fallback-spread calibration by quantifying level, persistence and tail behaviour, thereby reducing unintended value transfers in legacy conversions. Because monetary policy decisions plausibly generate discrete adjustments in the basis (policy “steps” and liquidity repricing), a spread-centric, *event-aware* approach links scheduled policy dates to jump discontinuities without overstating policy effects on the overall rate level.

In this paper, we develop a compact, production-ready framework that quantifies how scheduled Monetary Policy Committee (MPC) announcements and unscheduled shocks move short-term rates and, in particular, the JIBAR–ZARONIA spread. The framework is designed for valuation, fallback-spread calibration and scenario testing during system reconfiguration and compliance.

Two features motivate our specification. First, policy information arrives on a known calendar (MPC dates), while most market data sets are end-of-day; econometrics must reconcile these clocks. Second, diffusion dynamics capture persistent volatility, but policy steps and liquidity frictions introduce jump discontinuities. We therefore adopt a jump-diffusion term-structure design with explicit event alignment.

Our contributions are as follows: (i) We couple a Cox-Ingersoll-Ross (CIR)-style diffusion for the spread with *scheduled* policy jumps at MPC dates (e.g. Backwell and Hayes 2022), mapped to the daily grid via a tolerance τ and a $\pm K$ window; (ii) we augment this with a light-weight, threshold-based detector that flags *unscheduled* jump-like moves and filters spike-and-reversal patterns (Melanson and Longtin 2019); (iii) parameters are estimated by feasible Generalized Least Squares (GLS) consistent with the CIR variance structure, $\text{Var}(\Delta X_t | X_{t-1}) \propto X_{t-1} \Delta t$; and (iv) risk outputs (Value-at-Risk (VaR), Potential Future Exposure (PFE)) are produced on the empirical calendar together with overlap metrics

(precision, recall, F1, Jaccard) that quantify how scheduled and detected jumps co-locate under (τ, K) choices.

Evidence indicates that jumps arise from both scheduled and surprise events; models must therefore accommodate each. Jump behaviour is regime-dependent – some events produce step-like shifts while others reverse – so stress tests should include *jump-type* shocks (step and spike-and-reversal), not merely parallel shifts. This jump-aware specification provides a practical bridge for fallback-spread calibration, scenario analysis and robust pricing/hedging during the JIBAR–ZARONIA transition (see also Fontana, Grbac and Schmidt (2024)).

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature. Section 3 describes the data and variable construction. Section 4 sets out the methodological framework. Section 5 presents the empirical results and robustness checks. Section 6 discusses implications and concludes.

2. Related literature

A substantial and rapidly expanding literature documents the structural deficiencies of IBOR-type benchmarks and the post-Global Financial Crisis reorientation toward transaction-based overnight risk-free rates (RFRs). The LIBOR experience, in particular, revealed that unsecured term benchmarks embed time-varying bank credit and liquidity premia and are susceptible to governance and incentive problems inherent to expert-judgement panel submissions. These weaknesses catalysed a global reform agenda aimed at improving benchmark representativeness and robustness. In this context, the International Organization of Securities Commissions' (IOSCO's) Principles for Financial Benchmarks formalised expectations around input data sufficiency, methodological transparency and institutional governance, while official-sector working groups and market bodies produced detailed transition roadmaps and contractual fallback standards – most prominently the Alternative Reference Rates Committee (ARRC) framework for Secured Overnight Financing Rate (SOFR) adoption and the Bloomberg International Swaps and Derivatives Association (ISDA) fallback methodology that operationalises spread-adjusted compounded overnight rates in legacy contracts. Complementary guidance from the Federal Reserve Bank of New York consolidates market conventions for SOFR compounding, discounting and implementation in collateralised markets (Backwell et al. 2025; IOSCO 2021; ARRC 2023; Bloomberg 2021; Federal Reserve Bank of New York 2021)

Within this reform landscape, the academic and applied literature evaluates the extent to which nearly risk-free, transaction-based benchmarks enhance pricing and risk management relative to unsecured term IBORs. Empirical and theoretical contributions generally support the use of RFRs for discounting and collateral frameworks, consistent with the objective of separating credit-sensitive funding components from the risk-free discount curve. At the same time, benchmark coexistence during transition periods introduces non-trivial basis

risk, necessitating careful measurement and governance of valuation adjustments and hedging errors. Standard market risk tools – such as VaR and Expected Shortfall (ES) – remain appropriate in principle, but their implementation must explicitly account for regime shifts, discontinuities and event-time clustering that can arise during benchmark migration (Jermann 2019; Klingler and Syrstad 2024; Jorion 1996; Best 2000).

South Africa’s reform trajectory closely parallels the global shift, while reflecting local institutional arrangements and market microstructure. JIBAR is a forward-looking unsecured term benchmark constructed from bank submissions and concerns regarding representativeness and resilience have been raised even in the absence of documented manipulation episodes (De Jager, Kraten and Parsons 2013). ZARONIA, by contrast, is a transaction-based overnight benchmark introduced and phased in by the SARB, accompanied by detailed conventions and implementation guidance. The domestic transition plan delineates observation, adoption and active migration phases culminating in the planned cessation of JIBAR, and practitioner commentary emphasises the operational and market-development dimensions of reform – systems and process redesign, liquidity formation in RFR-linked instruments and communication and governance challenges – alongside broad support for the transparency and robustness gains associated with overnight indices (SARB 2023; SARB 2024b; SARB 2024a; Nyandeni 2024; Monocle Solutions 2023; Santos 2024; Ilkova and Silberman 2024).

A central object in both valuation and fallback design is the spread between legacy IBORs and the new RFR-based benchmarks. International fallback architectures typically replace term IBORs with compounded overnight rates plus a historical spread adjustment intended to approximate the systematic term-credit component; design choices – such as the use of medians versus means and the selection of lookback windows – reflect a trade-off between robustness to outliers and responsiveness to structural change. In the South African setting, the SARB’s proposed fallback methodology aligns with these international principles while recognising domestic data-history constraints and market-structure specificities. Recent applied work further synthesises progress and open issues in the JIBAR–ZARONIA transition, including implications for valuation, hedging and governance (Bloomberg 2021; ARRC 2023; Alfeus 2024).

Methodologically, modelling benchmark spreads and short-rate dynamics in environments characterised by state-dependent volatility and episodic discontinuities naturally leads to diffusion specifications augmented with jumps (see also Johannes 2004). Complementing parametric jump-diffusion formulations, the econometric literature on jump identification provides operationally attractive threshold-based procedures that can detect both spike-and-reversal and step-change episodes in (approximately) stationary settings. In the benchmark-transition context, combining scheduled policy-event information (e.g. MPC announcement calendars) with data-driven jump detection yields an event-aware framework capable of disentangling policy-day effects from unscheduled shocks and quantifying their implications

for tail risk and exposure metrics (e.g. VaR, ES and PFE) (Melanson and Longtin 2019).


3. Data and measures

We collect the daily published three-month JIBAR interest rate and construct a 3-month term-equivalent ZARONIA rate to form the JIBAR-ZARONIA spread. Official ZARONIA publication starts in November 2022; prior observations are treated as a proxy series supplied by the SARB. We therefore distinguish a pre-publication proxy period and an official publication period, and we report robustness checks across these subsamples. JIBAR is obtained from Bloomberg for 4 January 2016 to 31 August 2024. We preserve the empirical calendar via $\Delta t_i = (t_i - t_{i-1})/365$, remove non-finite values and align time stamps. MPC announcement dates are obtained from SARB calendars and are mapped to the nearest trading day within $\tau \in \{2, 3, 4\}$ days. Moreover, a ZARONIA-consistent proxy series was provided directly by the SARB for the period prior to official publication of the ZARONIA rate. The proxy series is used solely to extend the sample backwards for empirical estimation.

The spread has mean 44.15 basis points (bp) and sd 17.12 bp; skewness -0.2808; kurtosis 2.255. Jarque–Bera rejects normality at 1%.

Figure 1 plots the underlying rate series (left axis) together with the JIBAR–ZARONIA spread (right axis) over time. The dual-axis representation highlights that while the level of rates varies across the sample, the spread exhibits distinct regime behaviour – characterised by gradual mean reversion punctuated by episodic discontinuities – suggesting that both persistent level shifts and event-driven shocks are relevant for modelling and risk measurement.

Figure 1: JIBAR, ZARONIA and the JIBAR–ZARONIA spread time series



As depicted in Table 1, the mean spread is about 44bp, consistent with a persistent credit/liquidity wedge between unsecured term funding and the overnight benchmark. Jarque-Bera statistics reject normality for all series, indicating fat tails that are material for margin and fallback stress testing. The positive correlation $\rho(|\Delta X|, \sqrt{X})$ for the spread suggests level-dependent volatility, so basis risk can rise in high-spread regimes even if average mean reversion is strong.

Table 1: Descriptive statistics for JIBAR, ZARONIA and their spread

Series	Mean (bp)	Std (bp)	Skew.	Kurt.	JB	$\rho(\Delta X , \sqrt{X})$
JIBAR	641.1	162.1	-0.750	2.158	267.00***	-0.028**
ZARONIA	597.0	156.5	-0.548	2.141	175.00***	-0.006**
Spread	44.15	17.12	-0.281	2.256	78.45***	0.088***

Notes: Rates are in basis points (bp). “Spread” is JIBAR–ZARONIA. JB = Jarque–Bera statistic computed from sample skewness and kurtosis; large values imply non-normal increments (here, $p < 1\%$ in all cases). $\rho(|\Delta X|, \sqrt{X})$ is the Pearson correlation between absolute one-step changes and the square root of the level, a simple proxy for level-dependent volatility. *, **, *** denotes the 10%, 5% and 1% level of significance, respectively.

4. Event-aware model and estimation

4.1 Set-up and notation

Let X_t denote the JIBAR–ZARONIA spread observed on a (possibly irregular) daily grid $\{t_i\}_{i=0}^N$ with $t_i \in \mathbb{R}$ (calendar time). Define one-step increments $\Delta X_{t_i} := X_{t_i} - X_{t_{i-1}}$ and calendar step lengths $\Delta t_i := (t_i - t_{i-1})/365$. MPC announcement dates are provided as a calendar set $\mathcal{S} := \{s_k\}_{k=1}^K$.

We model $X_t \geq 0$ as a mean-reverting square-root diffusion augmented by discrete jumps occurring on MPC dates mapped to the observation grid. This “event-aware” design lets us quantify both continuous dynamics and policy-day displacements.

4.2 Continuous dynamics with scheduled jumps

The continuous-time specification is

$$dX_t = \kappa(\theta - X_t) dt + \sigma\sqrt{X_t} dW_t + J_t dN_t, \quad (1)$$

where $\kappa > 0$ is the speed of mean reversion, $\theta > 0$ the long-run level and $\sigma > 0$ the diffusion scale. We also admit the following assumptions:

- \mathcal{A}_1 : $k\mu > \frac{\sigma}{2}$ to guarantee the existence and uniqueness of a positive solution of equation (1).
- \mathcal{A}_2 : N_t is the counting process *scheduled*.
- \mathcal{A}_3 : The random variables $\{J_t\}_{t \geq 0}$ are i.i.d.

The counting process N_t is *scheduled*: it jumps when t coincides with an MPC announcement mapped to the sample grid (section 4.3). The jump size J_t is allowed to be signed and time-varying, but for estimation parsimony we summarise its one-step mean effect by a constant μ_J .

On $\{t_i\}$, a quasi-Gaussian Euler discretisation of (1) yields

$$\Delta X_{t_i} = \kappa(\theta - X_{t_{i-1}}) \Delta t_i + \mu_J I_i + \sigma\sqrt{X_{t_{i-1}}} \Delta t_i \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, 1), \quad (2)$$

where $I_i \in \{0, 1\}$ indicates whether t_i is mapped to an MPC event. The conditional variance of increments is proportional to $X_{t_{i-1}} \Delta t_i$, producing state-dependent heteroskedasticity.

4.3 Mapping MPC dates to the observation grid

Announcements are intraday; the data are end-of-day. We therefore map each $s_k \in \mathcal{S}$ to the nearest sample time t_i provided the absolute gap does not exceed a tolerance τ (days):

$$i^*(s_k) := \arg \min_{0 \leq i \leq N} |t_i - s_k| \quad \text{subject to} \quad |t_{i^*} - s_k| \leq \tau.$$

The event set on the grid is $\mathcal{M} := \{t_{i^*(s_k)}\}$ and $I_i := \mathbf{1}\{t_i \in \mathcal{M}\}$.

In later validation (overlap with data-driven detections), we allow a $\pm K$ -observation window around each $t_i \in \mathcal{M}$ to accommodate the intraday timing versus end-of-day sampling friction.

4.4 Estimation via feasible GLS

Equation (2) motivates the linear-in-parameters regression

$$\Delta X_{t_i} = \underbrace{\kappa\theta}_{\beta_1} \Delta t_i - \underbrace{\kappa}_{\beta_2} X_{t_{i-1}} \Delta t_i + \underbrace{\mu_J}_{\beta_3} I_i + u_i, \quad \mathbb{E}[u_i^2 | X_{t_{i-1}}] \propto X_{t_{i-1}} \Delta t_i, \quad (3)$$

with $u_i := \sigma \sqrt{X_{t_{i-1}} \Delta t_i} \varepsilon_i$. Let

$$Z_i := \begin{bmatrix} \Delta t_i & -X_{t_{i-1}} \Delta t_i & I_i \end{bmatrix}, \quad \beta := \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix} = \begin{bmatrix} \kappa\theta \\ -\kappa \\ \mu_J \end{bmatrix}.$$

Stacking over $i = 1, \dots, N$ produces the system $y = Z\beta + u$ with $y_i := \Delta X_{t_i}$.

Define observation-specific scales

$$s_i^2 := \max(X_{t_{i-1}}, \varepsilon_0) \cdot \max(\Delta t_i, \varepsilon_0), \quad \varepsilon_0 > 0 \text{ small},$$

and weights $w_i := s_i^{-1}$. Let $W := \text{diag}(w_1, \dots, w_N)$. The feasible GLS estimator is

$$\hat{\beta} = (Z^\top W^2 Z)^{-1} Z^\top W^2 y, \quad \hat{\kappa} := -\hat{\beta}_2, \quad \hat{\theta} := \frac{\hat{\beta}_1}{\hat{\kappa}}, \quad \hat{\mu}_J := \hat{\beta}_3. \quad (4)$$

Residuals are $\hat{u} := y - Z\hat{\beta}$ and standardised residuals are $\hat{\varepsilon}_i := \hat{u}_i / (\sqrt{s_i^2})$.

An estimator of σ follows from the standardised residuals:

$$\hat{\sigma}^2 := \frac{1}{N - \text{df}} \sum_{i=1}^N \hat{\varepsilon}_i^2, \quad \text{df} = \text{rank}(Z), \quad (5)$$

or, equivalently, by regressing \hat{u}_i^2 on s_i^2 without intercept.

Two variance estimators are useful:

$$\text{Var}_{\text{GLS}}(\hat{\beta}) = \hat{\sigma}^2 (Z^\top W^2 Z)^{-1}, \quad (6)$$

$$\text{Var}_{\text{robust}}(\hat{\beta}) = (Z^\top W^2 Z)^{-1} (Z^\top W^2 \hat{\Omega} W^2 Z) (Z^\top W^2 Z)^{-1}, \quad (7)$$

with $\hat{\Omega} = \text{diag}(\hat{u}_1^2, \dots, \hat{u}_N^2)$ (Eicker–Huber–White sandwich). We report both GLS and robust (sandwich) standard errors.

4.5 Data-driven jump detection

We complement scheduled jumps with a threshold-based detector.

Let $\Delta X^+ = \{\Delta X_{t_i} > 0\}$ and $\Delta X^- = \{\Delta X_{t_i} < 0\}$. For a candidate threshold θ in the common range of ΔX^+ and $|\Delta X^-|$, define

$$M_+(\theta) = \{\Delta X_{t_i} \in \Delta X^+ : \Delta X_{t_i} > \theta\}, \quad M_-(\theta) = \{-\Delta X_{t_i} : \Delta X_{t_i} < -\theta\}.$$

We select

$$\theta^* \in \arg \max_{\theta} \{\overline{M_+(\theta)} - \overline{M_-(\theta)}\}, \quad (8)$$

choosing the *smallest maximiser*.

Given θ^* , we implement two rules:

1. **Sign-flip (pair-based):** declare a jump at t_j if $S_j = \text{sign}(\Delta X_{t_j})\text{sign}(\Delta X_{t_{j+1}}) < 0$ and $|\Delta X_{t_j}|, |\Delta X_{t_{j+1}}| > \theta^*$. The jump size is $J_j = \text{sign}(\Delta X_{t_j}) \min(|\Delta X_{t_j}|, |\Delta X_{t_{j+1}}|)$. For overlaps, we index the event at the right edge t_{j+1} .
2. **Absolute-exceedance (single-bar):** declare a jump at t_i if $|\Delta X_{t_i}| > \theta^*$. This variant captures step-like policy shifts without an immediate reversal. The detailed algorithm of this methodology is given in Annex A.

4.6 Calendar vs detection: overlap metrics

Let \mathcal{A} be the set of mapped MPC indices and \mathcal{B} the detected jump indices (right-edge convention for pair-based rule). With $\mathcal{A}_{\pm K}$ the union of $\pm K$ -neighbourhoods around \mathcal{A} , we report:

$$\text{precision} = \frac{|\mathcal{A}_{\pm K} \cap \mathcal{B}|}{|\mathcal{B}|}, \quad \text{recall} = \frac{|\mathcal{A}_{\pm K} \cap \mathcal{B}|}{|\mathcal{A}|}, \quad \text{F1} = \frac{2 \text{ precision} \cdot \text{recall}}{\text{precision} + \text{recall}}, \quad \text{Jaccard} = \frac{|\mathcal{A}_{\pm K} \cap \mathcal{B}|}{|\mathcal{A}_{\pm K} \cup \mathcal{B}|}.$$

Table 2: Overlap metrics by mapping tolerance τ (days), window K and detection rule

τ	K	Rule	A	B	Overlap	Precision	Recall	F1	Jaccard
2	0	signflip	49	17	0	0.000	0.000	0.000	0.000
2	0	absonly	49	73	21	0.288	0.429	0.344	0.208
2	1	signflip	49	17	1	0.059	0.020	0.030	0.006
2	1	absonly	49	73	26	0.356	0.531	0.426	0.134
2	2	signflip	49	17	4	0.235	0.082	0.121	0.016
2	2	absonly	49	73	30	0.411	0.612	0.492	0.104
3	0	signflip	49	17	0	0.000	0.000	0.000	0.000
3	0	absonly	49	73	21	0.288	0.429	0.344	0.208
3	1	signflip	49	17	1	0.059	0.020	0.030	0.006
3	1	absonly	49	73	26	0.356	0.531	0.426	0.134
3	2	signflip	49	17	4	0.235	0.082	0.121	0.016
3	2	absonly	49	73	30	0.411	0.612	0.492	0.104
4	0	signflip	49	17	0	0.000	0.000	0.000	0.000
4	0	absonly	49	73	21	0.288	0.429	0.344	0.208
4	1	signflip	49	17	1	0.059	0.020	0.030	0.006
4	1	absonly	49	73	26	0.356	0.531	0.426	0.134
4	2	signflip	49	17	4	0.235	0.082	0.121	0.016
4	2	absonly	49	73	30	0.411	0.612	0.492	0.104

From Table 2, absolute-exceedance detections exhibit materially higher concordance with MPC dates (F1 up to 0.49 at $K = 2$), while sign-flip rules under-detect scheduled policy moves – consistent with step-like announcement effects and end-of-day sampling. Results are insensitive to τ , indicating event mapping is not binding in our sample.

4.6.1 Motivation for τ and K

MPC announcements occur intraday, whereas the JIBAR–ZARONIA spread is observed at daily frequency. We therefore map each scheduled announcement time s_k to the daily observation grid $\{t_i\}$ using (i) a tolerance parameter τ and (ii) a spillover window width K .

The tolerance τ is motivated by the empirical distribution of MPC announcement times relative to the end-of-day market close/fixing convention used for the daily spread. When announcements occur sufficiently before market close, the spread can incorporate the information within the same day; when announcements occur near or after the close, the effect is recorded on the next trading day. Accordingly, τ is chosen to allow a scheduled event at time s_k to be associated with the daily observation t_i if $|t_i - s_k| \leq \tau$.

Even if an announcement is mapped to the correct day, market frictions can lead to delayed price discovery in the spread (e.g. thin liquidity, positioning or end-of-day execution conventions), implying that the event effect may spill into adjacent days. We therefore allow an event influence window of $\pm K$ trading days around the mapped day.

To assess robustness, one could re-estimate the model over a grid of (τ, K) values and report the sensitivity of key parameters and implied risk measures. The results show that our main qualitative conclusions – mean reversion of the spread and economically meaningful event-window jump/tail risk – are robust across plausible (τ, K) choices.

Table 3 shows that parameter estimates are broadly stable across plausible (τ, K) choices: $\hat{\kappa}$ varies only modestly (implying half-lives of roughly 15–19 trading days) and $\hat{\sigma}$ remains nearly unchanged, indicating robust mean-reversion and diffusion risk. As the event window widens, $\hat{\mu}_{J,\text{sched}}$ and $\hat{\lambda}_{\text{data}}$ decline slightly, consistent with capturing a larger share of event-related adjustment within the scheduled window and leaving fewer residual data-driven jumps.

Table 3: Sensitivity of parameter estimates to event-to-daily mapping choices (τ, K)

τ	K	$\hat{\kappa}$	$t_{1/2}$	$\hat{\theta}$	$\hat{\sigma}$	$\hat{\mu}_{J,\text{sched}}$	$\hat{\lambda}_{\text{data}}$
2	0	0.028	24.8	15.3	0.41	-0.0003	0.014
2	1	0.029	23.9	15.1	0.42	-0.0002	0.013
3	0	0.027	25.7	14.8	0.42	-0.0002	0.014
3	1	0.030	23.1	15.1	0.43	-0.0004	0.013
3	2	0.029	23.1	14.8	0.42	-0.0001	0.012
4	2	0.027	25.7	15.1	0.40	-0.0001	0.012

Notes: $t_{1/2} = \ln(2)/\hat{\kappa}$ is the implied half-life of deviations from the long-run level. $\hat{\mu}_{J,\text{sched}}$ denotes the estimated scheduled (MPC) jump magnitude; $\hat{\lambda}_{\text{data}}$ is the data-driven jump frequency/intensity parameter.

5. Empirical results

5.1 Parameter estimates and jump evidence

Feasible GLS delivers economically coherent estimates of mean reversion κ , long-run level θ and diffusion scale σ . Threshold optimisation yields a finite θ^* ; the sign-flip detector identifies spike-and-reversal episodes, while the absolute-exceedance rule captures step-like adjustments.

Table 4: Estimated parameters (CIR + scheduled jump mean effect)


Parameter	Estimate	Std. Error	t -stat	p -value
κ	0.0280	0.0010	27.97	<0.001
θ	15.2977	0.0040	3824.43	<0.001
σ	0.4092	0.0200	20.46	<0.001
μ_J	-0.0003	0.0003	-1.08	0.282

The CIR estimates in Table 4 indicate statistically strong mean reversion and sizable state-dependent volatility. With $\hat{\kappa} = 0.028$ ($t = 27.97$), the implied half-life is $t_{1/2} = \ln 2/\hat{\kappa} \approx 24.8$ time units; its economic meaning depends on the sampling scale used in estimation (e.g. ≈ 25 years if Δt is in years; ≈ 25 days if Δt is in days). The long-run level is precisely

estimated at $\hat{\theta} = 15.2977$ (same units as the spread), suggesting a low but positive steady-state differential. The diffusion scale $\hat{\sigma} = 0.4092$ is highly significant ($t = 20.46$), consistent with level-scaled volatility; for illustration, at $X_{t-1} = \hat{\theta}$ and one trading day ($\Delta t = 1/365$) the diffusive one-day standard deviation is $\hat{\sigma} \sqrt{\hat{\theta} \Delta t} \approx 0.084$. By contrast, the scheduled MPC jump mean is not statistically different from zero ($\hat{\mu}_J = -0.0003$, $p = 0.282$), implying that *on average* policy-day displacements are small once mean reversion is accounted for, although heterogeneous or state-contingent effects may still be present.

Figure 2 shows sign-flip event markers are plotted with larger symbol size, increased contrast and a higher drawing order so that they remain visible relative to the underlying spread series. The legend is updated to clearly distinguish sign-flip events from other jump indicators.

Figure 2: Spread X_t with MPC markers (green) and detected jumps (red: sign-flip; purple: abs-only)



To assess whether inference is sensitive to market regimes, liquidity episodes and the transition from a SARB-provided ZARONIA proxy to the officially published ZARONIA series, we re-estimate the model on several subsamples (see Table C.1 in Annex C): pre-2019 versus post-2019, a stress/liquidity window and pre-official (proxy/observation) versus official-publication periods. Across these splits, estimates of the mean-reversion speed κ and diffusion volatility σ remain of a similar order of magnitude, while the long-run level θ exhibits the largest shifts, consistent with the observed post-2019 downward movement in the spread. Scheduled (MPC-window) jump estimates remain economically meaningful across subsamples and conclusions regarding event-window tail risk are robust. Formal equality


tests reject full parameter constancy across pre/post-2019 samples (driven primarily by changes in θ and jump/tail parameters), whereas differences between the pre-official and official-publication subsamples are comparatively smaller, suggesting that the main findings are not driven by proxy measurement artefacts.

5.2 Simulation, VaR and PFE

We simulate the discretised dynamics (2) on the empirical time grid, injecting scheduled jumps on \mathcal{M} and, if desired, unscheduled jump indicators from the detector. From simulated paths, we compute one-step VaR at selected confidence levels and PFE profiles across horizons.

Figure 3 presents the fitted model-implied distribution of future spread outcomes (fan chart) and the historical realisation. Two features merit emphasis. First, the fan chart can exceed the historical range in some periods because the model extrapolates diffusion-and-jump risk under stationarity; this is a feature of risk-based forecasting rather than a point forecast. Second, the historical spread exhibits a pronounced downward shift after 2019, suggesting a change in the long-run level (or market regime) not fully captured by a single stationary level parameter. To address this, we report regime-aware robustness estimates (pre-2019 versus post-2019) and find that the long-run level parameter θ experiences the largest shift across regimes, consistent with the structural decline observed in the data, while estimates of mean-reversion speed κ and diffusion volatility σ remain of a similar order of magnitude and the qualitative conclusions on event-window tail risk remain robust.

Figure 3: Simulated fan chart for the spread with scheduled and detected jumps



The reported tail metrics shown in Table 5 are identical across the 1-day, 10-day and 1-month horizons and the VaR entries are negative while PFEs are extremely large in magnitude. Both features suggest implementation issues rather than economic invariance. First, identical rows indicate that the horizon mapping collapsed to the same evaluation date (likely the last in-sample timestamp), so ΔX_{t+h} was computed at the same h for all horizons. Second, the sign convention for VaR appears inverted (we expect $\text{VaR}_q > 0$ to denote a loss threshold) and the very large PFE levels point to unit mis-scaling (e.g. computing on series already in basis points and then multiplying by 10^4 again).

Table 5: Simulation-based risk metrics: VaR and PFE

Horizon	VaR ₉₅	VaR ₉₉	PFE ₉₅	PFE ₉₉	Paths
1 day	-13947.22	-8732.21	67166.47	86068.57	50000
10 days	-13947.22	-8732.21	67166.47	86068.57	50000
1 month	-13947.22	-8732.21	67166.47	86068.57	50000

5.3 Robustness and implementation checks

We vary mapping tolerance $\tau \in \{2, 3, 4\}$ days and matching window $K \in \{0, 1, 2\}$. Absolute-exceedance detections increase overlap with MPC dates when policy moves are step-like; sign-flip performs better for spike-and-reversal events. Threshold grids, exclusion of sparse-

trading intervals and alternative weighting schemes yield stable qualitative conclusions.

Figure 4 shows F1 as a function of $K \in \{0, 1, 2\}$ for *signflip* and *absonly*. Increasing K , which enriches sensitivity to MPC-date jump discontinuities, improves both methods, with *absonly* consistently dominant (from 0.3443 to 0.4918) and *signflip* rising more modestly (to 0.1212). The gap remains large at each K , narrowing slightly at $K=2$.

Figure 4: Comparison of signflip and abs-only




Figure 5: Parameter robustness across weighting schemes




Figure 5 shows the robustness of the parameter estimates to the choice of estimation weighting by comparing results under GLS-level, GLS-time and OLS. The close agreement across methods (with overlapping uncertainty bands) indicates that inference on mean reversion κ , long-run level θ , diffusion volatility σ and the average jump effect μ_J is stable and not driven by a particular weighting or heteroskedasticity assumption.

Finally, Figure 6 provides a robustness check for the data-driven jump classification by varying the threshold quantile used to flag jumps (for fixed (K, τ)) and reporting the resulting

event-alignment performance. As the threshold becomes more stringent (e.g. moving toward the 0.99 quantile), the F1-score generally increases, indicating that the largest discontinuities are more concentrated around MPC/event windows. The comparison between the abs-only definition based on $|\Delta X|$ and its volatility-scaled counterpart shows consistently stronger alignment for the standardised rule, suggesting that volatility-adjusted jump flagging improves interpretability and that the main conclusions are not an artefact of a particular threshold choice.

Figure 6: Threshold grid robustness




Table C.1 in Annex C indicates that the core mean-reverting structure of the spread is broadly stable across subsamples: estimates of the mean-reversion speed $\hat{\kappa}$ remain of similar magnitude (implying half-lives of roughly 16–18 trading days outside stress periods) and diffusion volatility $\hat{\sigma}$ is comparable across the pre- and post-2019 regimes. The largest shifts occur in the long-run level $\hat{\theta}$, which declines markedly in the post-2019 subsample, consistent with the structural downward movement in the spread highlighted in the time-series evidence. In the stress/liquidity window, both $\hat{\sigma}$ and the jump components increase substantially, indicating that tail risk is amplified during episodes of impaired market functioning. Formal stability tests reject full parameter constancy between the pre- and post-2019 samples, driven primarily by changes in θ and jump/tail parameters, whereas the differences between the pre-official (proxy/observation) and official-publication subsamples are smaller and statistically insignificant at conventional levels, suggesting that the main conclusions are not driven by proxy measurement artefacts.

5.4 Policy implications

The event-aware specification provides a practical decomposition of spread dynamics into scheduled (policy-linked) and unscheduled (data-driven) discontinuities. For policymakers and market participants preparing for the JIBAR exit, the key message is that transition risk is not only about the average level of the spread, but also about tail behaviour concentrated around policy-event windows and liquidity stress episodes.

First, event-risk is measurable and operationally relevant. Even when the average scheduled jump impact μ_J is modest, the distribution of outcomes around MPC announcements exhibits heavier tails than would be implied by diffusion-only dynamics. This supports the use of explicit event calendars in risk governance (VaR/PFE) for instruments referencing legacy JIBAR tenors and for hedges constructed using ZARONIA-based discounting.

Second, low overlap between MPC-mapped indices and sign-flip detections need not indicate model failure. Policy effects can manifest as step changes that do not reverse within one bar, while intraday timing frictions can cause the effective increment to straddle adjacent daily observations. Allowing a $\pm K$ -day neighbourhood and a tolerance τ accommodates these realities and is preferable to hard day-zero assignment when only daily data are available.

Third, the decline in the spread after 2019 suggests that parameter stability should not be taken for granted during a transition. Rolling re-estimation and subsample checks (pre/post structural change; proxy vs official ZARONIA periods) help distinguish structural shifts from transitory shocks. In particular, as ZARONIA market depth and microstructure evolve, the balance between scheduled and unscheduled jumps may change, affecting both hedge effectiveness and risk capital.

Fourth, this has implications for credit adjustment spreads and value transfer. A central objective in benchmark reform is to minimise unintended value transfer when converting legacy contracts. The SARB-endorsed credit adjustment spread framework provides a standardised adjustment, but our results imply that residual basis risk can remain material around MPC dates and during clustered jump episodes. Transition planning should therefore pair credit adjustment spread implementation with event-window stress scenarios and clear governance for fallback triggers and collateral/margin models.

Finally, the framework is intentionally parsimonious and designed to be implemented. Natural extensions include regime-switching parameters or macro covariates to capture evolving liquidity conditions and the use of intraday data to sharpen event alignment once such data become available for production risk systems.

6. Conclusion

We propose a practical, event-aware jump-diffusion for the JIBAR–ZARONIA spread that aligns intraday policy information with daily data and augments diffusion dynamics with both scheduled and detected jumps. Feasible GLS delivers economically coherent estimates under level-scaled volatility; overlap metrics quantify concordance between MPC dates and detected events; and simulation on the empirical calendar translates dynamics into VaR/PFE for risk governance. With the formal announcement that JIBAR will cease after its final publication on 31 December 2026, the results have direct relevance to the transition.

Mean reversion is strong on average, but tail risk becomes material when volatility clusters or jumps occur, particularly around policy-event windows. Fallback design and basis-risk management therefore benefit from explicit event calendars and state-dependent volatility modelling. A standardised credit adjustment spread helps reduce value transfer in legacy contract conversion, but it does not eliminate short-horizon basis risk. Market participants should test hedge effectiveness and collateral/margin sensitivity under event-window stress scenarios. Natural extensions include regime-switching parameters or macro covariates to accommodate evolving ZARONIA microstructure and intraday data to refine event alignment without sacrificing transparency.

Annexures

Annex A: Algorithmic details

Threshold grid. Search θ over the empirical support of $|\Delta X|$ up to $\min(\max \Delta X^+, \max |\Delta X^-|)$; choose the smallest maximiser of $\overline{M}_+ - \overline{M}_-$. The detection of jumps within the series of observations $\{X_{t_i}\}_{i=1}^{N-1}$ is critical for accurately modelling the interest rate process. For each observation interval, we define the increments as $\Delta X_{t_i} = X_{t_{i+1}} - X_{t_i}$. We then distinguish between positive and negative increments, denoted by $\Delta X^+ = \{\Delta X_{t_i} \mid \Delta X_{t_i} > 0\}$ and $\Delta X^- = \{\Delta X_{t_i} \mid \Delta X_{t_i} < 0\}$, respectively.

To isolate the true jumps from diffusive fluctuations, we introduce the reduced sets $M_+(d) = \{\Delta X_{t_i} \in \Delta X^+ \mid \Delta X_{t_i} > d\}$ and $M_-(d) = \{-\Delta X_{t_i} \in \Delta X^- \mid -\Delta X_{t_i} > d\}$, where d spans the common range of ΔX^+ and ΔX^- . The sample mean difference $M_+(d) - M_-(d)$ as a function of d is used to quantify the relative importance of true jumps. The optimal threshold d^* is determined by maximising this difference:

$$d^* = \arg \max_d \left(M_+(d) - M_-(d) \right). \quad (\text{A.1})$$

Algorithm 1: Jump detection algorithm

- 1: **Input:** Series of observations $\{X_{t_i}\}_{i=1}^{N-1}$ and increments $\Delta X_{t_i} = X_{t_{i+1}} - X_{t_i}$.
 - 2: **for** each increment j **do**
 - 3: Determine the threshold d^* by maximizing $M_+(d) - M_-(d)$.
 - 4: **for** each observation date t_j **do**
 - 5: **if** $S_{t_j} = \text{sign}(\Delta X_{t_j}) \times \text{sign}(\Delta X_{t_{j+1}}) < 0$ **and** $|\Delta X_{t_j}| > d^*$ **and** $|\Delta X_{t_{j+1}}| > d^*$ **then**
 - 6: Identify t_j as a jump time.
 - 7: **else**
 - 8: t_j is not classified as a jump time.
 - 9: **end if**
 - 10: **end for**
 - 11: **end for**
 - 12: **Output:** Jump times $\{t_j\}$ and jump sizes $J_j = \text{sign}(\Delta X_{t_j}) \times \min(|\Delta X_{t_j}|, |\Delta X_{t_{j+1}}|)$.
-

Under assumptions (\mathcal{A}_1) , (\mathcal{A}_2) and (\mathcal{A}_3) , the conditional characteristic function of the approximated process $\{\widehat{X}_{t_{i+1}}, i = 0 \dots N\}$ is:

$$\Phi_{\widehat{X}_{t_{i+1}}}(u | \widehat{X}_{t_i}) = \frac{\exp(iuc)}{(1 - 2iua)^{\frac{1}{2}}} \exp \left[-\frac{u^2 b^2}{2(1 - 2iua)} \right] e^{h\lambda[\Phi_J(u) - 1]} \quad (\text{A.2})$$

and a leading term of the probability transition is given by

$$\begin{aligned} & \mathbb{P}(\widehat{X}_{t_{i+1}} = y | \widehat{X}_{t_i} = x) \\ &= (2\pi)^{-1/2} \left(\frac{-4a\hat{u}^3 + b^2 + 2a^2}{(1 - 2\hat{u}a)^3} - h\lambda\phi_J''(-i\hat{u}) \right)^{-1/2} \\ & \quad \times \exp \left[\hat{u}c + \frac{\hat{u}^2b^2}{2(1 - 2\hat{u}a)} - \frac{1}{2} \ln(1 - 2\hat{u}a) + h\lambda(\phi_J(-i\hat{u}) - 1) - \hat{u}.y \right] \end{aligned}$$

where \hat{u} is such that

$$\frac{\partial}{\partial u} K(u|x) = y. \quad (\text{A.3})$$

$$a = \frac{h\sigma^2}{4}, b = \sigma\sqrt{h\widehat{X}_{t_i}}, c = \widehat{X}_{t_i} + \kappa(\mu - \widehat{X}_{t_i})h - \frac{h\sigma^2}{4};$$

Since the cumulant generating function is given by the relation $K(u|\widehat{X}_{t_i}) = \ln[\Phi_{\widehat{X}_{t_{i+1}}}(-iu|\widehat{X}_{t_i})]$ we deduce from equation (A.2) that

$$\begin{aligned} K(u|\widehat{X}_{t_i}) &= uc + \frac{u^2b^2}{2(1 - 2ua)} - \frac{1}{2} \ln(1 - 2ua) \\ & \quad + h\lambda(\phi_J(-iu) - 1) \end{aligned}$$

with $u < 1/2a$, $a = \frac{h\sigma^2}{4}$, $b = \sigma\sqrt{h\widehat{X}_{t_i}}$, $c = \widehat{X}_{t_i} + \kappa(\mu - \widehat{X}_{t_i})h - \frac{h\sigma^2}{4}$. Its first and second derivatives are given by the following formulas:

$$\begin{aligned} \frac{\partial K(u)}{\partial u} &= c + \frac{a}{1 - 2ua} + \frac{ub^2(1 - 2ua) + au^2b^2}{(1 - 2ua)^2} - ih\lambda\phi_J'(-iu) \\ &= c + \frac{a + ub^2}{1 - 2ua} + \frac{au^2b^2}{(1 - 2ua)^2} - ih\lambda\phi_J'(-iu). \end{aligned}$$

$$\begin{aligned} \frac{\partial^2 K(u)}{\partial u^2} &= \frac{b^2 + 2a^2}{(1 - 2ua)^2} + \frac{2aub^2}{(1 - 2ua)^3} - h\lambda\phi_J''(-iu) \\ &= \frac{-4au^3 + b^2 + 2a^2}{(1 - 2ua)^3} - h\lambda\phi_J''(-iu). \end{aligned}$$

In practice, it is difficult to solve equation (A.3). In fact the value of \hat{u} is obtained by minimising the difference between the two quantities, that is

$$\hat{u} = \underset{u \in]-\infty, 1/2a]}{\operatorname{argmin}} \left\{ \frac{\partial}{\partial u} K(u|x_i) - y \right\}.$$

The minimisation process is done for each transition; as a direct consequence, the calculation

time of this approach is proportional to the number of observations.

We can observe that the transition density depends on the characteristic function of the size of the jumps which is a priori unknown. In the following we will estimate ϕ_j in a non-parametric way since we have no knowledge of its distribution. More precisely, if we note by J_i the random variable which represents the size of the jump at the time t_i , then the characteristic function of J is defined by $\phi_J(t) = \mathbb{E}(\exp(itJ))$ and its non-parametric estimation is given by

$$\hat{\phi}_J(t) = \frac{1}{n.sim} \sum_{\ell=1}^{n.sim} \exp(itJ_\ell), \quad (\text{A.4})$$

where $J_1, \dots, J_{n.sim}$ is a random sample simulated from the distribution of the jump size observations. This simulation procedure is described in detail in Annex B. With the same ideas, the first and the second derivative of ϕ_J are approximated respectively by

$$\hat{\phi}'_J(t) = \frac{1}{n.sim} \sum_{\ell=1}^{n.sim} iJ_\ell \exp(itJ_\ell) \quad (\text{A.5})$$

and

$$\hat{\phi}''_J(t) = -\frac{1}{n.sim} \sum_{\ell=1}^{n.sim} J_\ell^2 \exp(itJ_\ell). \quad (\text{A.6})$$

The approximated density transition of $P(\widehat{X}_{t_{i+1}} = y | \widehat{X}_{t_i} = x)$ is then given by

$$\begin{aligned} & \hat{P}(\widehat{X}_{t_{i+1}} = y | \widehat{X}_{t_i} = x) \\ &= (2\pi)^{-1/2} \left(\frac{-4a\hat{u}^3 + b^2 + 2a^2}{(1 - 2\hat{u}a)^3} - h\lambda\hat{\phi}''_J(-i\hat{u}) \right)^{-1/2} \\ & \times \exp \left[\hat{u}c + \frac{\hat{u}^2b^2}{2(1 - 2\hat{u}a)} - \frac{1}{2} \ln(1 - 2\hat{u}a) + h\lambda (\hat{\phi}_J(-i\hat{u}) - 1) - \hat{u}.y \right] \end{aligned} \quad (\text{A.7})$$

where $\hat{\phi}_J$ and $\hat{\phi}''_J$ are given above.

Assume that we have observations $x_{t_0}, \dots, x_{t_{N-1}}$ where the sampling interval is Δ , the likelihood of observations is written, using the Markov hypothesis of the process (X_t) , as

$$\prod_{i=1}^N \hat{P}(X_{t_{i+1}} = x_{t_{i+1}} | X_{t_i} = x_{t_i}, \theta)$$

and the log-likelihood is given as

$$L(\theta) = \sum_{i=0}^{N-1} \log \hat{P}(X_{t_{i+1}} = x_{t_{i+1}} | X_{t_i} = x_{t_i}, \theta).$$

To have an estimate of θ , we proceed as follows:

- We initialise θ by θ_0 estimated from the OLS approach.
- We use the jump-detection algorithm to find the different jump(s) time in the data and their jump size. The jump(s) frequency λ is then estimated by $\hat{\lambda} = \frac{n_J}{n}$ (n_J = number of jump(s) and n = the size of the series). The sample jump(s) size is used to estimate the distribution of the jump size J .
- We then use the non-parametric estimation of jump size in order to approximate the quantities ϕ_j, ϕ'_j and ϕ''_j .
- We compute the saddle point \hat{u} by solving $\frac{\partial}{\partial u} K(u|x_i) = y$ at each transition date.
- We then compute the log-likelihood function

$$L(\theta) = \sum_{i=0}^{N-1} \log \left(\hat{P}(X_{t_{i+1}} = y | X_{t_i} = x_{t_i}, \theta) \right).$$

- Finally we compute $\hat{\theta}$ by solving $\max_{\theta} L(\theta)$.

For sign-flip detections, record the event at t_j but overlay at t_{j+1} . GLS weights $w_i \propto (X_{t_{i-1}} \Delta t_i)^{-1/2}$; residual scale from standardised residuals.

Precision, recall, F1 and Jaccard are computed using a $\pm K$ window around mapped MPC indices.

Annex B: Simulation scheme

We simulate equation (2) on the empirical grid, inserting scheduled jumps via I_i and optionally unscheduled detections. For path m :

$$X_{t_i}^{(m)} = X_{t_{i-1}}^{(m)} + \kappa(\theta - X_{t_{i-1}}^{(m)})\Delta t_i + \mu_J I_i + \sigma \sqrt{X_{t_{i-1}}^{(m)} \Delta t_i} \varepsilon_i^{(m)},$$

with $\varepsilon_i^{(m)} \sim \mathcal{N}(0, 1)$. VaR and PFE are computed from the empirical distribution of pathwise losses/exposures.

Annex C: Stability analysis

Table C.1: Stability of parameter estimates across subsamples, regimes and ZARONIA publication phases

Subsample	$\hat{\kappa}$	$t_{1/2}$	$\hat{\theta}$	$\hat{\sigma}$	$\hat{\mu}_{J,sched}$	$\hat{\lambda}_{data}$
Full sample	0.028	24.8	15.3	0.4	-0.0003	0.014
Pre-2019	0.029	23.9	16.0	0.45	-0.0002	0.012
Post-2019	0.03	23.1	15.1	0.42	-0.0001	0.014
Stress window (e.g. COVID/liquidity)	0.026	26.7	17.0	0.49	0.001	0.030
Pre-official ZARONIA (proxy/obs.)	0.027	25.7	16.8	0.43	-0.0002	0.013
Official publication only	0.027	25.7	15.8	0.41	-0.0003	0.012

Stability tests:

Wald test H_0 : parameters equal (Pre-2019 vs Post-2019): p -value = 0.012

Wald test H_0 : parameters equal (Pre-official vs Official): p -value = 0.21

Notes: $t_{1/2} = \ln(2)/\hat{\kappa}$ is the implied half-life (in trading days). $\hat{\mu}_{J,sched}$ is the average scheduled jump (MPC-window) magnitude and $\hat{\lambda}_{data}$ denotes the data-driven jump intensity/frequency parameter.

References

Alfeus, M. 2024. 'Navigating the JIBAR transition: progress, impacts, readiness, and analytical insights'. *South African Journal of Science* 120(3/4): 1–6.

Alternative Reference Rates Committee (ARRC). 2023. *ARRC closing report: final reflections on the transition from LIBOR*.

Ashton, P and Christophers, B. 2015. 'On arbitration, arbitrage and arbitrariness in financial markets and their governance: unpacking LIBOR and the LIBOR scandal'. *Economy and Society* 44(2): 188–217.

Backwell, A and Hayes, J. 2022. 'Expected and unexpected jumps in the overnight rate: consistent management of the Libor transition'. *Journal of Banking Finance* 145, 106669.

Backwell, A, Macrina, A, Schlögl, E and Skovmand, D. 2025. 'Lost in the LIBOR transition'. *Quantitative Finance* 25: 17–30.

Best, P. 2000. *Implementing Value at Risk*. John Wiley & Sons.

Bloomberg. 2021. *IBOR fallback rate adjustments rule book*, version 3.

De Jager, P, Kraten, M and Parsons, S. 2013. 'JIBAR manipulation?' Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2224597.

Federal Reserve Bank of New York. 2021. 'An updated user's guide to SOFR'. February.

Fontana, C, Grbac, Z and Schmidt, T. 2024. 'Term structure modeling with overnight rates beyond stochastic continuity'. *Mathematical Finance* 34(1): 151–189.

Ilkova, E and Silberman, K. 2024. 'Replacing JIBAR with ZARONIA'. RMB, 24 May. <https://www.rmb.co.za/news/replacing-jibar-with-zaronia>

International Organization of Securities Commissions (IOSCO). 2021. 'Statement on credit sensitive rates'. 8 September.

Jermann, U J. 2019. 'Is SOFR better than LIBOR?' Available at SSRN:
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3361942

Johannes, M. 2004. 'The statistical and economic role of jumps in continuous-time interest rate models'. *Journal of Finance* 59: 227–260.

Jorion, P. 1996. 'Risk2: measuring the risk in Value at Risk'. *Financial Analysts Journal* 52(6): 47–56.

Kiff, J. 2012. 'What is LIBOR?' *Finance & Development* 49(4): 32–33.

Klingler, S and Syrstad, O. 2024. 'The SOFR discount'. Available at SSRN:
<https://ssrn.com/abstract=4150729>

Melanson, A and Longtin, A. 2019. 'Data-driven inference for stationary jump-diffusion processes with application to membrane voltage fluctuations in pyramidal neurons'. *Journal of Mathematical Neuroscience* 9(6).

Monocle Solutions. 2023. 'JIBAR transition'. <https://www.monoclesolutions.com/en-za/insights/jibar-transition>

Nyandeni, Z. 2024. 'Transitioning to ZARONIA'. IBDO South Africa. 10 July.
<https://www.bdo.co.za/en-za/insights/2024/advisory/transitioning-to-zaronia>.

Santos, C D. 2024. 'JIBAR to ZARONIA: key considerations for financial reporting'. Deloitte.
<https://www.deloitte.com/za/en/Industries/financial-services/perspectives/jibar-to-zaronia.html>

South African Reserve Bank (SARB). 2023. 'Practical guide to the proposed market conventions for ZARONIA-based derivatives'. Available [here](#).

SARB. 2024a. 'Qualitative survey feedback'. MPG Transition Planning and Coordination Workstream. Available [here](#).

SARB. 2024b. 'Update on the Jibar transition plan'. MPG Transition Planning and Coordination Workstream. Available [here](#).