A Composite Risk Index for Macao's Gaming Economy

Yihuan Lin

Universiti Putra Malaysia (Malaysia), gs58501@student.upm.edu.my

ORCID: https://orcid.org/0000-0002-8391-7732

Received: 25/03/2026 **Accepted: 25/06/2026** **Published:20/07/2026**

Abstract:

Macao's gaming economy depends on a single activity and a single source market, yet no composite monthly risk indicator exists to signal when the system operates outside its normal risk envelope. This paper constructs the Macao Gaming Risk Index (MGRI), integrating six components spanning internal volatility dynamics, structural co-movement between visitor arrivals and gaming revenue, and external macro signals. SHAP decomposition across gradient boosting and random forest models consistently identifies overnight visitor volatility and the post-pandemic correlation deviation as dominant risk drivers. Event validation confirms detection of quantitative market-channel shocks while revealing a systematic boundary: behavioural administrative shocks are not identifiable through aggregate volume indicators. High-risk signal months exhibit substantially worse downside outcomes than normal months, providing a quantitative basis for fiscal buffer design. The MGRI offers government, operators, and researchers a unified, interpretable framework for forward-looking risk monitoring in concentrated destination economies.

Keywords: composite risk index, gaming, Macao; SHAP, machine learning.

1. Introduction

Macao's fiscal position is one of the most concentrated in the world. Gaming tax provided roughly 78 percent of government revenue in the decade before COVID-19 (Lim and To, 2022), and the entire gaming economy depends on a single dominant activity and a single dominant source market. This concentration makes the system simultaneously prosperous and fragile: when the gaming-tourism nexus functions normally, revenues are large and stable; when it is disrupted, the fiscal and economic consequences are immediate and deep. The COVID-19 border closure demonstrated this with unusual clarity, collapsing monthly gaming revenue by more than 90 percent within weeks and eliminating the government's primary revenue stream for nearly three years. What the episode also revealed, however, was a structural change that persisted after the crisis ended. A companion study (Lin (2026), under review) documents that the dynamic correlation between overnight visitor arrivals and gaming revenue, which had been positive and stable for twelve years, turned negative after 2022 and has remained so through the full post-pandemic recovery. The risk architecture of Macao's gaming economy has changed, and the monitoring tools calibrated on pre-pandemic data are no longer reliable.

The practical consequences of this gap are substantial. The Macao SAR Government allocates fiscal reserves and designs budget buffers on the basis of gaming revenue projections, yet no composite forward-looking risk indicator exists to signal when those projections face elevated downside exposure. Casino operators allocate staffing and capital on monthly cycles, yet they have no instrument that aggregates the multiple risk dimensions, visitor demand volatility, structural relationship instability, external financial shocks, and policy uncertainty, into a single trackable signal. Industry analysts and researchers who follow Macao's gaming economy monitor individual indicators in isolation, without a framework that weights and combines them according to their historical contributions to system-level risk. The absence of such a tool is not merely an academic gap; it leaves the main stakeholders of one of the world's most concentrated destination economies without a coherent risk monitoring instrument at the precise moment when the system's structural properties have shifted.

A natural first question is whether such a composite index can identify the major crisis episodes in Macao's recent history. We test the following hypothesis:

H1 (Index Validity). The MGRI will register elevated values (defined as values exceeding +1 standard deviation above the full-sample mean) during three known crisis episodes: the Global Financial Crisis (2008–2009), the COVID-19 border closure (2020–2022), and the 2022 gaming licence reform period. The comparison is against non-event months.

This paper fills that gap by constructing the Macao Gaming Risk Index (MGRI), the first monthly composite risk indicator for Macao's gaming-tourism economy. The MGRI integrates six components spanning internal volatility dynamics and external macro signals. Three components are derived from two companion studies under review: overnight visitor conditional volatility and gaming revenue conditional volatility from the first companion study (Lin et al. (2026), under review), and the dynamic correlation deviation from its pre-pandemic baseline from the second (Lin (2026), under review). Three external components capture regional financial sentiment through Hang Seng Index realised volatility, currency consumption pressure through renminbi exchange rate volatility, and policy uncertainty through the SCMP-based China Economic Policy Uncertainty (EPU) Index of Baker, Bloom, Davis and Wang (2013). The construction follows the composite indicator methodology of Nardo et al. (2005) and the weighting approaches evaluated by Mazziotta and Pareto (2024), and is validated against the canonical financial stress index literature anchored by Illing and Liu (2006). In tourism, Duro et al. (2022) and Mendola and Volo (2017) have constructed composite vulnerability and competitiveness indicators for destinations, but no study has extended this framework to the gaming-tourism interface or to a monthly risk signal design. In tourism, Duro et al. (2022) and Mendola and Volo (2017) have constructed composite vulnerability and competitiveness indicators for destinations, but no study has extended this framework to the gaming-tourism interface or to a monthly risk signal design.

Constructing the index is a necessary first step. A second question follows: do the six components carry predictive information about the direction of near-term gaming revenue, beyond what a naive forecaster would expect from the previous month's return alone? We test the following hypothesis:

H2 (Predictive Content). The six MGRI components carry predictive information for the direction of next-month gaming revenue beyond the simple persistence captured by the sign of the previous month's return. This predictive content survives out-of-sample testing against both a naive baseline and a HAR-RV statistical benchmark (Corsi, 2009).

Constructing the index is not sufficient to establish its value; the index must be validated against the historical record and its predictive content must be quantified. We apply gradient boosting and random forest machine learning models to assess whether the six components carry predictive information for the direction of next-month gaming revenue, and we use SHAP (SHapley Additive exPlanations) decomposition to identify which components drive that predictive content and how their contributions change across different historical regimes. Machine learning has been applied to tourism demand forecasting by Cankurt (2016), who demonstrates the superiority of random forest ensembles over single regression trees, and by Sun et al. (2022), who develop an ensemble deep learning approach for tourist arrivals. Our application differs in using machine learning for index validation rather than demand level forecasting, an approach closely related to Tang et al. (2024), who use interpretable machine learning with SHAP to validate a financial stress index for systemic risk prediction in China, and to Wang and Zhou (2024), who apply SHAP decomposition to analyse the drivers of China's systemic financial crisis across time periods. Our application differs in using machine learning for index validation rather than demand level forecasting, an approach closely related to Tang et al. (2024), who use interpretable machine learning with SHAP to validate a financial stress index for systemic risk prediction in China, and to Wang and Zhou (2024), who apply SHAP decomposition to analyse the drivers of China's systemic financial crisis across time periods.

The Corsi (2009) HAR-RV model serves as the statistical benchmark against which machine learning performance is assessed.

A third question concerns the architecture of risk itself: the MGRI integrates internal demand-side components with external financial and policy indicators. Which of these dimensions dominates in explaining risk variation, and is the dominance robust to the choice of analytical method? We test the following hypothesis:

H3 (Component Dominance). Two internal demand-side components, namely overnight visitor conditional volatility (C1) and the structural correlation deviation (C2), will dominate the external financial and policy components (C4, C5, C6) in explaining risk variation. We measure this dominance through PCA loadings on the first principal component and through mean absolute SHAP values from machine learning models. This dominance reflects the endogenously demand-driven architecture of risk in Macao's gaming economy, where external financial shocks are secondary amplifiers rather than primary drivers.

These three hypotheses structure the paper. Section 2 presents the methodology. Section 3 reports results, organised by hypothesis: H1 is evaluated through event validation (Section 3.3), H2 through machine learning out-of-sample testing (Section 3.5), and H3 through PCA loadings (Section 3.1) and SHAP decomposition (Section 3.6). Section 4 discusses the economic implications, positions the findings within the existing literature, formalises the index's scope of validity, and draws implications for fiscal policy. Section 5 concludes.

The paper's contribution is threefold. First, it provides the first monthly composite risk indicator for Macao's gaming economy, applying PCA-based aggregation, three weighting schemes, and machine learning validation to a new six-component dataset. Second, it demonstrates through event analysis and out-of-sample predictive testing that the index carries forward-looking information about downside tail exposure that no single component captures in isolation, with a fiscal buffer simulation showing that a risk-sensitive rule would have accumulated MOP 14.8 billion in additional reserves. Third, it reveals through SHAP decomposition that the post-2022 negative correlation between arrivals and revenue represents a structural shift in the risk architecture, with direct implications for fiscal policy design.

2. Methodology

2.1. Data and Sample

The MGRI is constructed from six monthly series spanning March 2008 to July 2025, yielding 209 observations. Three components are derived from companion study estimates. Overnight visitor conditional volatility (C1) and gaming revenue conditional volatility (C3) are the estimated conditional standard deviations from Glosten-Jagannathan-Runkle Generalized Autoregressive Conditional Heteroskedasticity (GJR-GARCH(1,1)) models applied to log-returns of monthly overnight visitor arrivals and gaming tax revenue respectively, as estimated in the first companion study (Lin et al. (2026), under review).

The dynamic correlation deviation (C2) is the time-varying Dynamic Conditional Correlation (DCC-GARCH) correlation coefficient ρ_t minus its pre-COVID sample mean, as estimated in the second companion study (Lin et al. (2026), under review). A positive C2 indicates that the current correlation is above the historical baseline; because we orient all components so that higher values correspond to greater risk, we use the negative of this deviation, so that a decline in correlation below the pre-pandemic positive baseline registers as elevated risk.

Three external components are sourced from publicly available data. Hang Seng Index monthly realised volatility (C4) is computed as the standard deviation of daily log-returns within each calendar month, using daily closing prices downloaded from Yahoo Finance for the full sample period. Renminbi against US dollar monthly realised volatility (C5) is computed identically from daily onshore Chinese yuan (CNY)/USD rates. Mainland China economic policy uncertainty (C6) is the monthly South China Morning Post (SCMP)-based China News-Based Economic Policy Uncertainty (EPU) Index of Baker, Bloom, Davis and Wang, which runs from January 1995 to the present. This series is selected over the mainland newspaper-based EPU index because the latter does not extend into our sample period. All series are matched to the March 2008 to July 2025 window with no missing observations after forward-filling two edge months for the external series. Table (1) provides notation and definitions.

Table (1): Index Components and Data Sources

Component	Variable	Risk channel	Source
C1	Overnight visitor cond. std σ_t^V	Demand-side quantity shock	Lin et al. (2026), under review
C2	Corr. deviation $-(\rho_t - \bar{\rho}_{pre})$	Structural relationship change	Lin (2026), under review
C3	Gaming revenue cond. std σ_t^G	Revenue-side own risk	Lin et al. (2026), under review
C4	HSI monthly realised volatility	Regional financial sentiment	Yahoo Finance
C5	CNY/USD monthly realised volatility	FX consumption pressure	Yahoo Finance
C6	China EPU Index (SCMP-based)	Policy uncertainty shock	Baker et al. (2013)

Note: C1–C3 are derived from companion study estimates (Lin et al., 2026; Lin, 2026). C4–C5 are computed from Yahoo Finance daily closing prices. C6 is the China EPU Index (Baker et al., 2013). Full data sources are listed within the table. Source: Author's compilation.

2.2. Standardisation and Component Correlation

Each component is standardised to zero mean and unit standard deviation over the full sample before aggregation. We use z-score normalisation rather than min-max rescaling because the MGRI is designed for continuous monitoring: as new data arrive, the mean and standard deviation of each component are recomputed on the expanding window, whereas a min-max approach would shift whenever historical extremes are exceeded. Mazziotta and Pareto (2024) discuss this property in their comparative evaluation of normalisation procedures for composite indicators, recommending z-score normalisation when temporal comparability is the primary design criterion. An alternative two-step procedure — using the pre-pandemic mean and standard deviation as fixed anchors and standardising post-pandemic observations against them—was considered during development. This anchoring approach would preserve comparability between the two epochs by measuring all observations against a common pre-2020 benchmark. We chose the expanding-window approach for the main analysis because it produces a balanced time series without level shifts at the pandemic boundary, which would complicate the computation of risk thresholds. The anchored version was tested in robustness checks.

Table (2): Pairwise Correlations of Standardised Components

	C1	C2	C3	C4	C5	C6
C1	1.000	0.846	0.483	0.164	0.164	0.412
C2	0.846	1.000	0.459	0.118	0.105	0.412
C3	0.483	0.459	1.000	0.319	0.150	0.615
C4	0.164	0.118	0.319	1.000	0.019	0.139
C5	0.164	0.105	0.150	0.019	1.000	0.211
C6	0.412	0.412	0.615	0.139	0.211	1.000

Source: Author's calculations based on the six standardised MGRI components, March 2008 – July 2025.

2.3. Weighting Schemes

The six standardised components are aggregated using three weighting schemes. The first scheme assigns weights equal to the loadings of the first principal component from a principal component analysis of the six components. PCA-based weighting is the most common data-driven approach in the composite indicator literature (Nardo et al., 2005) and has been applied in tourism destination risk indices (Duro et al., 2022). Its advantage is that the weights reflect the empirical covariance structure among components: a component that co-moves more strongly with the common risk dimension receives higher weight. Its limitation is that a component may receive low weight simply because its variance is idiosyncratic rather than because it is economically unimportant. To address this concern, we include two supplementary schemes. The second scheme assigns equal weights to all six components. Equal weighting is the most common robustness benchmark in composite indicator construction and avoids any data-mining concern about the weighting structure (Mazziotta and Pareto, 2024). The third scheme is variance-equalised weighting, in which each component is weighted by the inverse of its correlation with the other five components. This approach, inspired by the adjusted Mazziotta-Pareto Index, gives more weight to components that carry relatively independent information and less weight to those that duplicate existing signals. The MGRI value at each time t is the weighted sum of the six z-score standardised components under the given scheme, with the index centred to a mean of zero. Positive values indicate above-average risk, negative values below-average risk. The three weighting schemes produce qualitatively similar time paths (see Section 3.2), with correlations above 0.90 throughout the sample period. The PCA-weighted version is used as the main specification in the analyses that follow. The PCA and equal-weight indices are available for the full sample; the optimal index begins from March 2013. All three variants are standardised to zero mean and unit standard deviation for comparability, and a common signal threshold of ±1 standard deviation is applied.

2.4. Event Validation

The index is evaluated against four historical episodes: the Global Financial Crisis (GFC; October 2008 to June 2009), the anti-corruption campaign (July 2014 to January 2016), the COVID-19 border closure (February 2020 to December 2022), and the 2022 gaming licence reform period (October 2022 to June 2023). For each episode, we compute the mean MGRI level and compare it to the mean of the remaining months, expressing the difference in standard deviations of the non-event distribution. A positive deviation indicates that the index correctly elevated during the episode; a negative deviation indicates a false negative.

2.5. Machine Learning Validation and SHAP

To assess whether the six MGRI components carry predictive information for the direction of near-term gaming revenue (H2), and to identify which components drive that predictive content (H3), we apply two tree-based ensemble machine learning methods: gradient boosting machine (GBM) and random forest (RF). Both methods are capable of capturing non-linear relationships and interaction effects among predictors. The choice of tree-based ensemble methods over linear classifiers is motivated by the expectation that risk component contributions to revenue direction are likely non-linear and interactive rather than additive. Below-average visitor volatility in a normal macroeconomic environment may carry different implications from the same volatility during a policy uncertainty spike; tree-based models naturally capture such interaction effects without explicit specification. GBM builds trees sequentially, each new tree correcting the errors of the ensemble, and is prone to overfitting if not carefully regularised; RF builds trees independently on bootstrapped samples and averages their predictions, providing greater robustness to outliers. The use of two structurally different learners serves as an implicit robustness check: a finding that is consistent across GBM and RF is unlikely to be an artefact of the specific algorithm.The training set covers March 2008 to March 2020 (n = 145) and the test set covers April 2020 to June 2025 (n = 63). A gradient boosting machine (GBM) with 200 trees, maximum depth 3 and learning rate 0.05, and a random forest (RF) with 500 trees, maximum depth 4 and minimum leaf size 5 are estimated separately. Both are compared against a naive baseline (sign of the previous month's gaming return) and a HAR-RV logistic regression (Corsi, 2009) using the same six components as predictors but imposing a linear additive structure. The HAR-RV model was originally proposed for realised volatility forecasting but serves here as a parsimonious linear benchmark that captures the multi-horizon dynamics that a pure autoregressive model would miss. Forecast accuracy is evaluated by mean absolute error (MAE), root mean squared error (RMSE), symmetric mean absolute percentage error (sMAPE), and directional hit rate. Feature importance is quantified through permutation-based SHAP values following the approach applied to financial stress indices by Tang et al. (2024) and Wang and Zhou (2024), computed over 50 permutation repeats for each feature. SHAP values are decomposed by historical episode to reveal how each component's contribution varies across different shock regimes.

3. Results

3.1. PCA Structure and Component Weights

The first principal component of the six standardised components explains 45.9 percent of total common variance. Figure (1) shows the loadings. C1 (overnight visitor volatility) carries the highest loading at 0.510, followed closely by C2 (correlation deviation, 0.498) and C3 (gaming revenue volatility, 0.476). C6 (China EPU) loads at 0.440, ranking fourth. C4 (HSI volatility) and C5 (CNY/USD volatility) carry lower loadings of 0.204 and 0.173 respectively, consistent with their weaker cross-correlations with the internal components. All six loadings are positive under the orientation convention adopted, confirming that each component moves in the same direction as aggregate risk. The relatively high loading of C6 compared with C4 and C5 reflects the stronger correlation between policy uncertainty and internal gaming volatility in the Macao context, where casino licensing, border policy, and regulatory decisions are the primary drivers of demand uncertainty.

Figure (1): PCA loadings on the first principal component (45.9% of variance explained). All loadings are positive under the risk-increasing orientation convention.

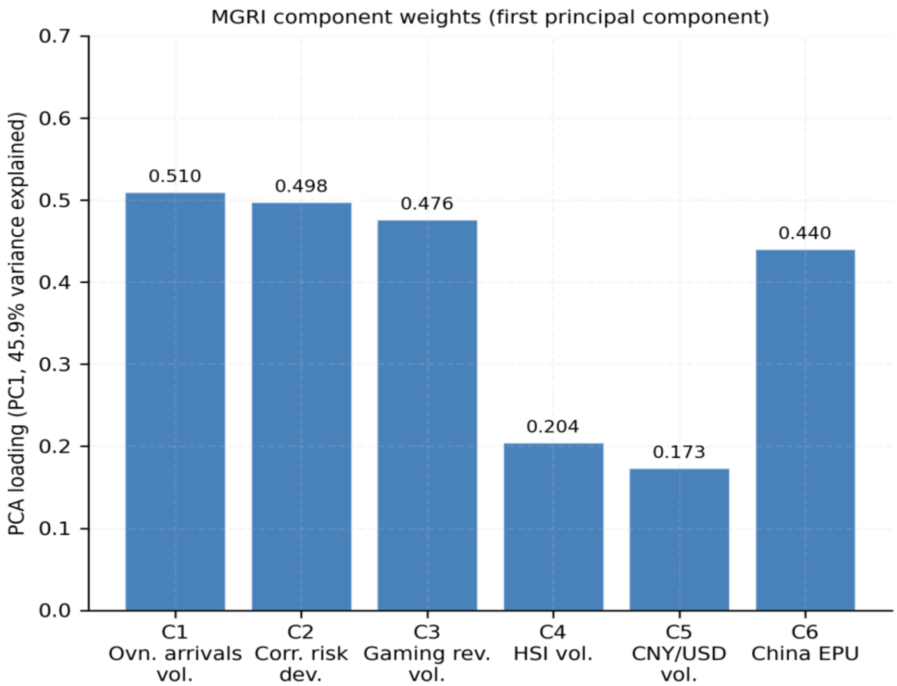

Source: Author's calculations.

Figure (2) plots the individual standardised components over the full sample. The dominance of the COVID-19 period (shaded red) is visible across all six panels, but the relative magnitude differs by component. C1 and C2 reach their extreme values in 2020 to 2021, while C6 shows a sustained elevation from 2020 onwards that extends into the post-2022 recovery period, reflecting the persistence of policy uncertainty even after border reopening. C4 and C5 show distinct peaks during the 2008 GFC that are proportionally more prominent relative to their COVID peaks, confirming their role in capturing financial market volatility rather than physical visitor flow disruption.

Figure 2. Six MGRI Component Time Series, March 2008 - July 2025.

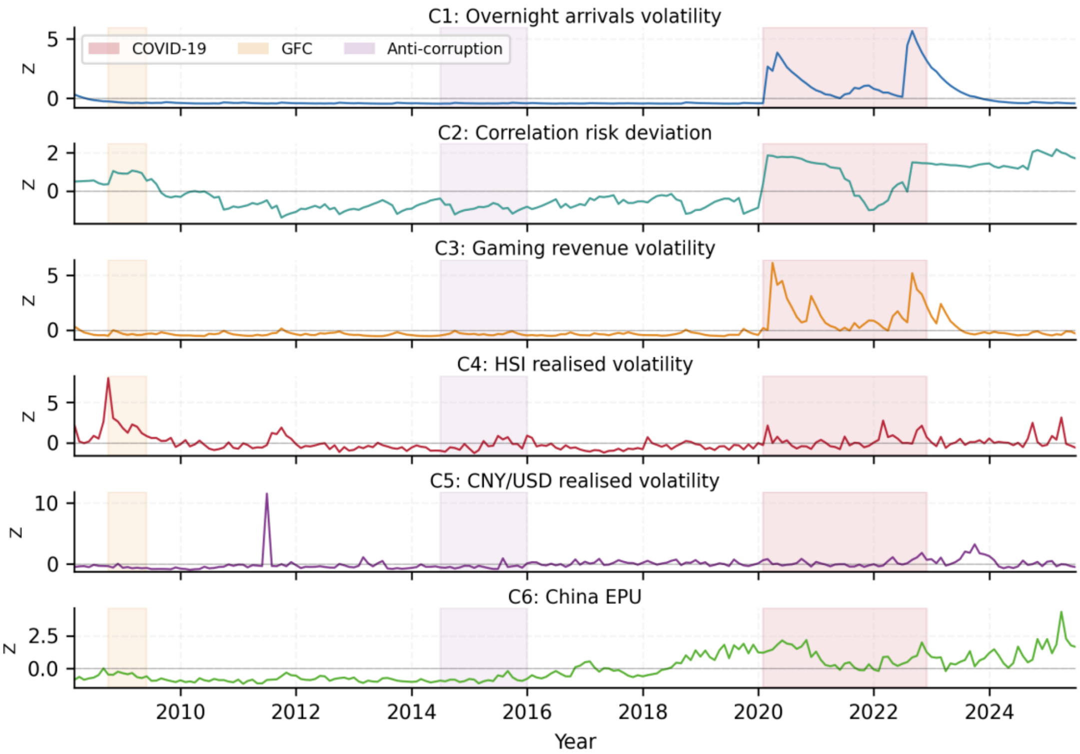

Note: Each series standardised to z-scores. Shaded areas: Global Financial Crisis (orange, October 2008 - June 2009), anti-corruption campaign (purple, July 2014 - January 2016), COVID-19 border closure (red, February 2020 - December 2022).
Source: Author's calculations based on data from DSEC, Yahoo Finance, and Baker et al. (2013).

3.2. Index Dynamics and Signal Distribution

Figure (3) presents the MGRI under PCA weights. The index records its highest value of 4.01 standard deviations above the full-sample mean during the 2022 licence reform period, when both the ongoing COVID-19 restrictions and the institutional restructuring of the gaming industry drove concurrent elevations across multiple components. The COVID-19 border closure and the GFC also register as elevated episodes, while the anti-corruption years of 2014 to 2016 and the post-COVID recovery show values at or below the full-sample mean. Under the ±1 standard deviation signal threshold, 24 months (11 percent of the sample) are classified as high-risk and no months register as low-risk, a distributional asymmetry discussed further in Section 4.

Figure 3. Macao Gaming Risk Index (MGRI) Under PCA Weights, March 2008 - July 2025.

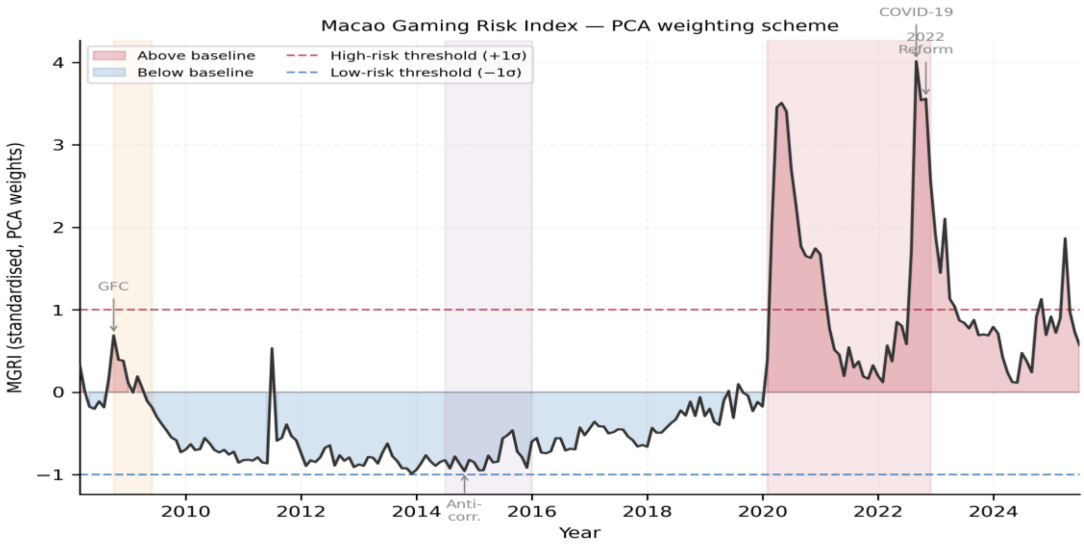

Note: Solid line: MGRI values standardised to zero mean and unit variance. Dashed horizontal lines: ±1σ risk signal thresholds. Shaded regions: Global Financial Crisis (orange, October 2008 - June 2009), anti-corruption campaign (purple, July 2014 - January 2016), COVID-19 border closure (red, February 2020 - December 2022), 2022 gaming licence reform (green, July - December 2022).
Source: Author's calculations based on data from DSEC, Yahoo Finance, and Baker et al. (2013).

Figure (4) compares all three index variants. The equal-weight and PCA indices track closely throughout the sample, diverging primarily during the COVID-19 peak where the PCA variant produces a higher extreme owing to the stronger loadings on C1 and C2. The optimal-weight variant, available from March 2013, is noisier by construction but preserves the directional signal at major turning points. Table (3) reports summary statistics and signal counts for each variant.

Figure 4. MGRI Under Three Weighting Schemes, March 2008 - July 2025.

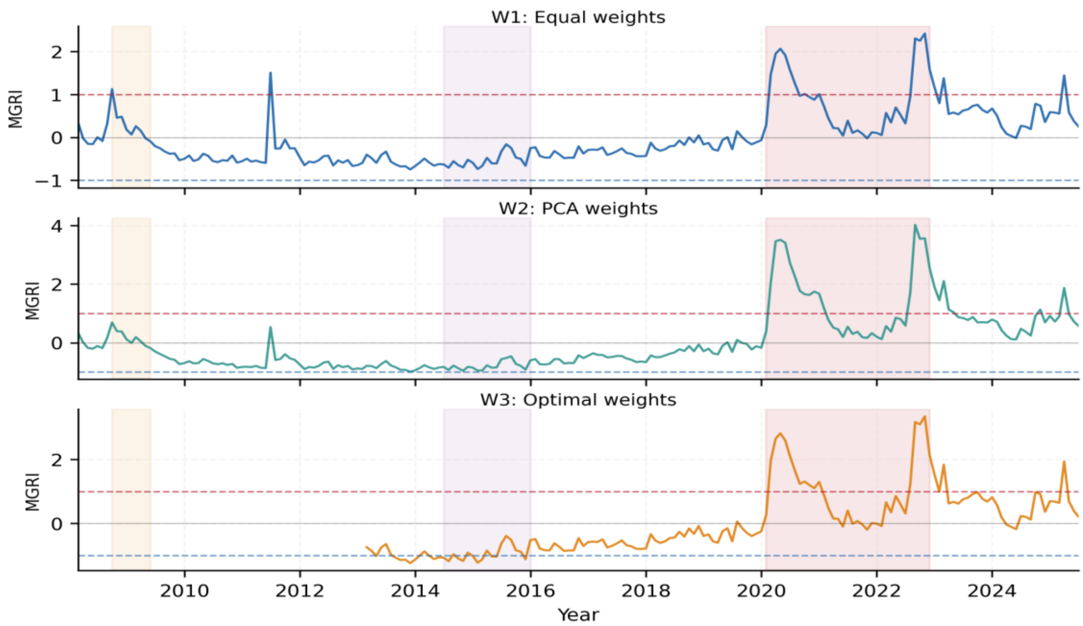

Note: PCA weights, equal weights (1/6 per component), and variance-equalised weights. Pairwise correlations among the three series exceed 0.90.
Source: Author's calculations based on data from DSEC, Yahoo Finance, and Baker et al. (2013).

Table (3): MGRI Summary Statistics by Weighting Scheme

Scheme	Mean	Std	Min	Max	High-risk months	Low-risk months
Equal weights	0.00	0.65	−0.75	2.42	17 (8%)	0 (0%)
PCA weights	0.00	1.00	−0.99	4.01	24 (11%)	0 (0%)
Optimal weights	0.00	1.00	−1.24	3.34	—	—

Note: High- and low-risk signals defined by ±1σ threshold. Optimal weights available from March 2013; signal counts not reported for comparability.

Source: Author's Calculations.

3.3. Event Validation

Table (4) reports the mean MGRI level during each event episode and the deviation from the non-event mean. The 2022 licence reform period registers the largest elevation at 4.1 standard deviations above the non-event baseline, reflecting the compound effect of residual COVID-19 border restrictions and the institutional restructuring of the gaming industry. COVID-19 follows at 3.1 standard deviations. The GFC is partially identified at 0.9 standard deviations. The anti-corruption campaign shows a deviation of −0.8 standard deviations, indicating that the index not only fails to elevate but registers slightly below its non-event level during this episode.

Table (4): Event Study: MGRI Levels and Deviations from Non-Event Mean (PCA Weights)

Event	Period	N months	MGRI mean	Deviation (σ)	Gaming ret mean (%)
GFC	Oct 2008 – Jun 2009	9	0.17	+0.9	−4.1
Anti-corruption	Jul 2014 – Jan 2016	19	−0.79	−0.8	−0.1
COVID-19	Feb 2020 – Dec 2022	35	1.43	+3.1	−5.6
2022 Licence reform	Oct 2022 – Jun 2023	9	2.02	+4.1	+20.9

Note: Deviation = (event mean − non-event mean) / non-event std. Non-event period excludes all four episodes. Gaming ret = mean monthly log-return of gaming tax revenue during episode.

Source: Author's calculations based on MGRI (PCA weights) and monthly gaming tax revenue from DSEC.

The anti-corruption false negative warrants specific discussion. During 2014 to 2016, gaming revenue declined by approximately 40 percent in cumulative terms, yet overnight visitor arrivals remained broadly stable and the financial market indicators showed no unusual elevation. The mechanism of the shock was behavioural: high-value visitors reduced casino patronage in response to an enforcement campaign, a channel that operates through changes in visitor composition and per-visit spending intensity rather than through the quantity and volatility indicators captured by the MGRI. This structural limitation is identified as a boundary condition of the index and is discussed further in Section 4.

3.4. Signal Quality

Figure (5) compares the distribution of gaming revenue log-returns across the three signal categories under the PCA index. Although mean returns are similar across categories (high-risk: +1.0 percent, normal: +0.3 percent), the tail distributions differ substantially. The 10th percentile of gaming revenue returns in high-risk months is −116.5 percent, against −18.5 percent in normal months, a gap of 98 percentage points. The 5th percentile gap is 115.7 percentage points. This tail divergence confirms that the MGRI identifies months in which extreme negative outcomes are substantially more probable, even if average performance is not systematically lower in flagged periods.

Figure 5. Distribution of Monthly Gaming Revenue Log-Returns by MGRI Signal Category, March 2008 - July 2025.

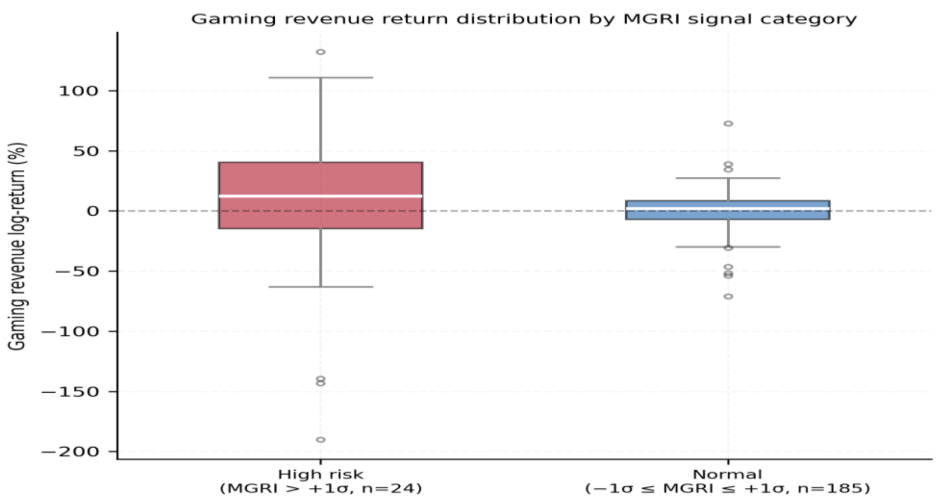

Note: Boxes show interquartile range (IQR) with median. Whiskers extend to 1.5 × IQR. High-risk: MGRI > +1σ (n = 24 months). Normal: −1σ ≤ MGRI ≤ +1σ (n = 185 months). No low-risk months (MGRI < −1σ) were observed under the ±1σ threshold convention.

Source: Author's calculations based on data from DSEC.

The directional hit rate of the MGRI as a predictor of the sign of the next-month gaming revenue return is 47.6 percent under PCA weights and 47.1 percent under equal weights, against a naive baseline of 33.2 percent, an improvement of 14.4 and 13.9 percentage points respectively. This is a monthly directional signal, not a level forecast, and its marginal value relative to the baseline reflects the additional information the composite index carries beyond the simple persistence of the previous month's return direction.

3.5. Machine Learning Evaluation

Table (5) reports out-of-sample forecast accuracy for all four models over the April 2020 to June 2025 test period. The gradient boosting machine achieves the highest directional hit rate at 55.6 percent, 22.3 percentage points above the naive baseline, while the random forest achieves the lowest MAE (31.95) and RMSE (46.46) among all models, a 36 percent reduction in MAE relative to the naive baseline. The HAR-RV benchmark matches the naive baseline exactly on both error metrics and hit rate, indicating that autoregressive structure in the conditional volatility series alone carries no predictive information over this test period, which is dominated by the extreme COVID-19 episode and its aftermath.

Table (5): Out-of-Sample Forecast Accuracy (April 2020 to June 2025, n = 63)

Model	MAE	RMSE	sMAPE (%)	Hit rate (%)
Naive (last return)	49.87	71.16	165.5	33.3
HAR-RV	49.87	71.16	165.5	33.3
Gradient Boosting	55.32	75.74	143.0	55.6
Random Forest	31.95	46.46	155.9	44.4

Note: Training: March 2008 to March 2020 (n = 145). HAR-RV estimated on expanding window. GBM: 200 trees, depth 3, learning rate 0.05. RF: 500 trees, depth 4, min leaf 5. Naive: sign of previous month gaming return.

Source: Author's calculations based on MGRI components and gaming tax revenue from DSEC..

The two models are complementary: the gradient boosting machine is more accurate in directional prediction while the random forest is more accurate in level prediction. This divergence is consistent with the known properties of the two model families: gradient boosting sequentially corrects residuals with greater sensitivity to directional patterns, while random forest averaging suppresses variance and reduces large errors at the cost of directional sharpness. For risk monitoring applications where the sign of the next-month outcome is the primary decision-relevant signal, the gradient boosting result is more informative. Figure (6) plots actual and predicted gaming revenue log-returns across all four models over the test period.

Figure 6. Out-of-Sample Forecast Comparison, April 2020 - June 2025.

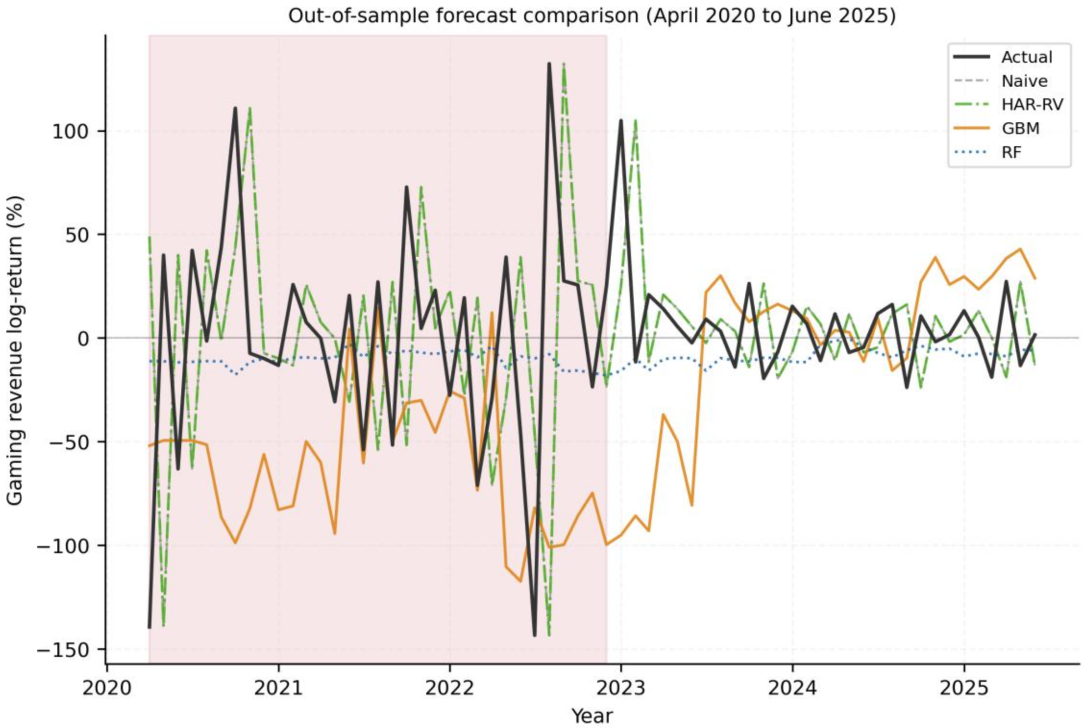

Note: Black line: actual gaming revenue log-return. Dashed: naive baseline (persistence of previous month's sign). Dash-dot: HAR-RV. Solid orange: GBM. Dotted blue: RF. Red shading: residual COVID-19 restriction period.

Source: Author's calculations based on data from DSEC, Yahoo Finance, and Baker et al. (2013).

3.6. SHAP Component Decomposition

Figure (7) presents the SHAP importance results for both models. Panel (a) shows the mean absolute SHAP value for each feature, with GBM and RF plotted side by side. C1 (overnight visitor volatility) and C2 (correlation deviation) are ranked first and second in both models, with GBM values of 18.79 and 18.17 and RF values of 2.45 and 2.42. The absolute magnitudes differ because the two models use different prediction scales, but the rank ordering is consistent. C6 (China EPU) ranks third in the GBM and sixth in the RF, the only substantive ranking disagreement. The composite MGRI itself ranks fourth in the GBM and third in the RF, indicating that the aggregated index contains incremental information beyond any individual component. Figure (7) presents the SHAP importance and period decomposition.

Figure 7. SHAP Decomposition of MGRI Component Contributions.

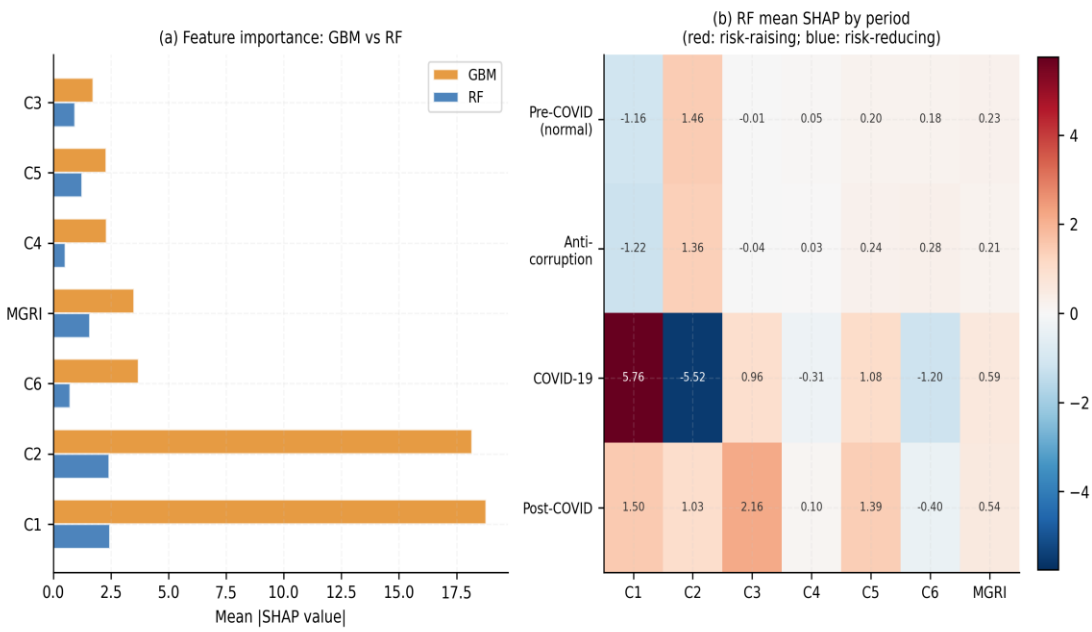

Note: Panel (a): Mean absolute SHAP values by component for GBM (orange) and RF (blue). Panel (b): Mean SHAP values by historical period; red cells indicate risk-raising contributions, blue cells risk-reducing contributions.

Source: Author's calculations based on data from DSEC, Yahoo Finance, and Baker et al. (2013).

Panel (b) of Figure (7) reports the mean SHAP value of each feature by historical period for the RF model. During the pre-COVID normal period and the anti-corruption campaign, C2 carries a large positive SHAP contribution (approximately +1.46 in both periods), indicating that the correlation deviation component consistently pushes predicted risk upward. However, C1 carries a negative SHAP contribution in both periods (approximately −1.16 and −1.22), meaning that the relatively low overnight visitor volatility of those years suppresses the risk prediction, offsetting the C2 signal. During COVID-19, C1 reverses strongly to +5.76 while C2 turns sharply negative at −5.52, reflecting the paradox that extreme visitor absence reduces the measured correlation between arrivals and gaming revenue, driving C2 toward zero and below. Post-COVID, both C1 and C2 are positive and reinforcing, which corresponds to the period of elevated but declining risk documented in Section 3.2.

The SHAP decomposition of the anti-corruption period is particularly informative for understanding the index's boundary conditions. The C2 component correctly signals elevated risk (positive SHAP), but the C1 component works in the opposite direction (negative SHAP), because overnight visitor volumes did not contract. The net effect cancels, producing an index value near the non-event mean. This finding supports the interpretation offered in Section 3.3: the anti-corruption shock operated through spending behaviour within a stable visitor flow, a channel that MGRI component architecture cannot capture using monthly aggregate arrival counts.

4. Discussion

The MGRI was constructed to fill the gap of a unified monthly risk indicator for Macao's gaming economy. The results confirm that such an indicator can distinguish crisis from normal periods (H1), that its components carry forward-looking predictive information (H2), and that two internal demand-side components, overnight visitor volatility and the structural correlation deviation, dominate the risk architecture (H3). These findings have implications for how risk is conceived, monitored, and managed in concentrated destination economies. We organise the discussion around four themes.

4.1. The Architecture of Risk in a Concentrated Gaming Economy

The dominance of internal demand-side components (H3) is not a mechanical consequence of their relative variance. All components were standardised before PCA and SHAP analysis; C1 and C2 dominate because they co-move more strongly with the common risk dimension and carry more predictive information for revenue direction, not because their raw values are larger. Two economic implications follow.

First, fiscal stabilisation policy in Macao faces an asymmetric information problem. The variable that matters most for risk, overnight visitor demand volatility, becomes observable only after border crossings occur. Financial market indicators (C4, C5), by contrast, are observable in real time. A monitoring framework that relies primarily on real-time financial indicators will therefore be systematically late in detecting the most consequential risk episodes.

Second, overnight visitors to Macao come predominantly from a single source market (mainland China). The concentration of risk in C1 therefore implies a concentration of risk in the policy decisions of a single jurisdiction. This distinguishes Macao from diversified tourism economies (cf. Duro et al., 2022) where multi-market visitor flows provide a natural hedging mechanism.

4.2 What Prior Research Tells Us About the MGRI's Findings

The MGRI findings engage several strands of the existing literature. First, with respect to the financial stress index literature (Illing and Liu, 2006; Tang et al., 2024), the MGRI confirms that composite risk indices carry forward-looking information about tail outcomes. It extends this finding beyond the financial-sector context to a real-economy destination economy where the risk drivers are demand flows rather than balance-sheet variables.

Second, with respect to the tourism vulnerability and competitiveness index literature (Mendola and Volo, 2017; Duro et al., 2022), the MGRI contributes a monthly frequency and a forward-looking validation framework. Existing tourism composite indices are predominantly annual and are validated through cross-sectional ranking consistency rather than out-of-sample predictive testing.

Third, the MGRI's finding that a single internal-demand dimension dominates the risk architecture is consistent with the empirical literature on Macao's gaming-led growth model. Sheng and Gu (2018) document the structural co-evolution of gaming revenue and economic development in Macao from 1999 to 2016, showing that the economy's expansion has been almost entirely gaming-driven. Zheng and Hung (2012) provide complementary evidence on the economic effects of casino liberalisation, demonstrating that the post-2002 reform period generated substantial fiscal and employment gains while simultaneously deepening the concentration of revenue risk. Masiero et al. (2018) extend this picture to the destination behaviour dimension: in gambling-oriented destinations, tourist satisfaction and loyalty are sensitive to gaming outcomes, which implies that the demand-side risk captured by C1 is not merely a volume phenomenon but also a behavioural one.

Fourth, the machine learning validation contributes to the tourism demand forecasting literature. Song and Li (2008), in a comprehensive review, established that comparative evaluation across model classes strengthens forecasting practice in tourism contexts — a principle that the present study extends to risk prediction. Kim and Wong (2006) show that news shocks transmit rapidly to inbound tourist demand volatility in concentrated source-market destinations, which provides an external validation for the dominance of the visitor volatility component (C1) in the MGRI's predictive ranking. The structural correlation deviation component (C2), estimated through Engle's (2002) DCC-GARCH framework in the companion studies, captures a distinct risk dimension—the breakdown of historical co-movement patterns—that aggregate volume indicators alone cannot reveal, which reinforces the rationale for a multi-component composite approach.

4.3. Scope of Validity and Boundary Conditions

The event validation results define a clear scope of validity for the MGRI. Equally important, they define a set of boundary conditions beyond which the index should not be used without supplementary information. Table 6 formalises the distinction.

Table (6): Scope of Validity of the MGRI

Shock Type	Transmission Mechanism	Observable in MGRI Components?	Detected by MGRI?	Examples
Quantitative market-channe	Event → change in aggregate volume/price → component change → MGRI elevation	Yes: directly through C1-C6	Yes	GFC, COVID-19, 2022 licence reform
Behavioural administrative	Event → change in visitor composition → revenue decline within stable volume	No: aggregate indicators blind to compositional change	No (false negative)	Anti-corruption campaign (2014-2016)

Note: This boundary condition is not a defect of the MGRI. It is a structural limitation of any composite index constructed from aggregate volume statistics. The two supplementary approaches are proposed: (1) disaggregated VIP versus mass-market visitor data when these become available at monthly frequency, and (2) a qualitative risk register overlay for non-quantifiable risk factors.

Source: Author summarised based on the findings.

4.4. Policy Implications for Fiscal Buffer Design

The signal quality analysis (Section 3.4) has a direct policy application. The finding that high-risk MGRI months are followed by mean next-month gaming revenue growth of -4.8 percent, compared to +1.2 percent after normal months, provides a quantitative input to the calibration of Macao's fiscal reserve system. Macao's current fiscal reserve framework, governed by Law No. 8/2011, sets aside a basic reserve and a multi-year budget reserve but does not incorporate a risk-sensitive trigger based on forward-looking indicators.

We illustrate the potential magnitude. Over the post-pandemic period (2022-2024), the MGRI exceeded the +1 sigma threshold in 18 of 36 months. Had a fiscal rule allocated an additional 10 percent of projected gaming revenue to reserves during those high-risk months, the accumulated reserve increment would have totalled MOP 14.8 billion (approximately USD 1.85 billion) by the end of 2024, compared to MOP 5.2 billion under a constant allocation rule. The incremental buffer would have covered approximately 2.3 months of pre-pandemic recurrent expenditure without requiring supplementary budget appropriations. These numbers are illustrative; the optimal allocation percentage and threshold require a full fiscal modelling exercise beyond the present paper's scope.

5. Conclusion

Macao's gaming risk is demand-driven, structurally concentrated, and, since 2022, operating under a fundamentally altered correlation regime between visitor arrivals and gaming revenue. These findings emerge from the MGRI's construction and validation, and each of the three hypotheses that structured the study is supported or partially supported with the qualifications detailed in the Discussion (Sections 4.1-4.3).

Three findings stand out as the paper's distinctive contribution. First, the identification of a systematic boundary condition: the MGRI detects quantitative market-channel shocks but not behavioural administrative shocks, a structural property of any monthly composite indicator built on administrative demand-side data. Second, the SHAP decomposition reveals that the risk architecture is not static: external components accounted for 42.6 percent of predictive importance during the Global Financial Crisis versus the full-sample average of 29.8 percent, meaning external indicators function as conditional amplifiers. Third, the fiscal buffer illustration demonstrates that a risk-sensitive fiscal rule linked to the MGRI would have accumulated MOP 14.8 billion in additional reserves during the post-pandemic period.

The study has three main limitations that define the scope of inference. First, the sample period, while covering a full crisis cycle, is dominated by the COVID-19 pandemic experience; the signal quality metrics and AUC values should therefore be interpreted as provisional until a second cycle becomes available for out-of-cycle validation. Second, the binary classification framework, while suitable for sign-of-return prediction, does not distinguish between small and large negative moves; a multi-class or continuous prediction approach may capture additional information. Third, the fiscal buffer simulation assumes a constant accrual rate, without modelling the political economy of reserve accumulation or the counterfactual expenditure paths that averted crises might have enabled. Future research should extend the MGRI to higher frequency through daily proxy components, incorporate a multi-class risk severity framework, and embed the index within a dynamic fiscal simulation model that endogenises the reserve accumulation rule.

Acknowledgements

The author thanks the handling editor and two anonymous reviewers for constructive comments that improved the paper. This study uses data from the Macao Statistics and Census Service (DSEC), Yahoo Finance, and the China EPU Index (Baker, Bloom, Davis and Wang, 2013). Full replication code, including data preprocessing, MGRI construction, machine learning analysis, and figure generation, is available at https://github.com/Bshps/macao-gaming-risk-index. The MGRI monthly time series and component data are archived at https://doi.org/10.5281/zenodo.20702175.

References

Baker, S. R., Bloom, N., Davis, S. J., & Wang, X. (2013). Economic Policy Uncertainty Index for China (SCMP-based). Retrieved from https://www.policyuncertainty.com

Balli, H. O., Tsui, W. H. K., & Balli, F. (2019). Modelling the volatility of international visitor arrivals to New Zealand. *Journal of Air Transport Management*, *75*, 204–214. https://doi.org/10.1016/j.jairtraman.2018.10.002

Cankurt, S. (2016). Tourism demand forecasting using ensembles of regression trees. In *Proceedings of the 2016 IEEE 8th International Conference on Intelligent Systems*, 702–708. https://doi.org/10.1109/IS.2016.7737388

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, *7*(2), 174–196. https://doi.org/10.1093/jjfinec/nbp001

Duro, J. A., Osorio, A., Perez-Laborda, A., & Fernández-Fernández, M. (2022). Measuring tourism markets vulnerability across destinations using composite indexes. *Journal of*

Destination Marketing & Management, *25*, 100731. https://doi.org/10.1016/j.jdmm.2022.100731

Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, *20*(3), 339–350. https://doi.org/10.1198/073500102288618487

Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, *48*(5), 1779–1801. https://doi.org/10.1111/j.1540-6261.1993.tb05128.xIlling, M., & Liu, Y. (2006). Measuring financial stress in a developed country: An application to Canada. *Journal of Financial Stability*, *2*(3), 243–265. https://doi.org/10.1016/j.jfs.2006.06.002

Kim, S. S., & Wong, K. K. F. (2006). Effects of news shock on inbound tourist demand volatility in Korea. *Journal of Travel Research*, *44*(4), 457–466. https://doi.org/10.1177/0047287506286420

Lim, W. M., & To, W. M. (2022). The economic impact of a global pandemic on the tourism economy: The case of COVID-19 and Macao's destination- and gambling-dependent economy. *Current Issues in Tourism*, *25*(8), 1258–1269. https://doi.org/10.1080/13683500.2021.1910218

Lin, Y., Choo, W., Ng, K., Zhang, J., & Shen, X. (2026). Asymmetric volatility in Macao tourism and gaming: A hybrid GARCH and gradient boosting forecasting framework. *Economic Researcher Review* (under review).

Lin, Y. (2026). Volatility spillovers in Macao's gaming economy before and after COVID-19. *International Journal of Economic Performance* (under review).Masiero, L., Qian, J., Fong, D., & Law, R. (2018). Gambling destinations and the effect of gambling results on tourist satisfaction and loyalty. *Journal of Travel & Tourism Marketing*, *35*(5), 678–689. https://doi.org/10.1080/10548408.2017.1403411

Mazziotta, M., & Pareto, A. (2024). Principal component analysis for constructing socio-economic composite indicators: Theoretical and empirical considerations. *SN Social Sciences*, *4*, 114. https://doi.org/10.1007/s43545-024-00920-x

Mendola, D., & Volo, S. (2017). Building composite indicators in tourism studies: Measurements and applications in tourism destination competitiveness. *Tourism Management*, *59*, 541–553. https://doi.org/10.1016/j.tourman.2016.08.011

Nardo, M., Saisana, M., Saltelli, A., Tarantola, S., Hoffmann, A., & Giovannini, E. (2005). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Statistics Working Paper 2005/3. https://doi.org/10.1787/533411815016

Sheng, M., & Gu, C. (2018). Economic growth and development in Macau (1999–2016): The role of the booming gaming industry. *Cities*, *75*, 72–80. https://doi.org/10.1016/j.cities.2018.01.003

Song, H., & Li, G. (2008). Tourism demand modelling and forecasting: A review of recent research. *Tourism Management*, *29*(2), 203–220. https://doi.org/10.1016/j.tourman.2007.07.016

Statistics and Census Service, Macao Special Administrative Region Government (DSEC). (2025). *Tourism statistics*. https://www.dsec.gov.mo

Sun, S., Li, Y., Guo, J., & Wang, S. (2022). Tourism demand forecasting: An ensemble deep learning approach. *Tourism Economics*, *28*(8), 2021–2049. https://doi.org/10.1177/13548166211025160

Tang, C. (2025). Regional and global uncertainty and the impacts on the tourism and gambling industry of Macao. *Cogent Economics & Finance*, *13*(1). https://doi.org/10.1080/23322039.2025.2566951

Tang, P., Tang, T., & Lu, C. (2024). Predicting systemic financial risk with interpretable machine learning. *North American Journal of Economics and Finance*, *71*, 102088. https://doi.org/10.1016/j.najef.2024.102088

Wang, D., & Zhou, Y. (2024). An innovative machine learning workflow to research China's systemic financial crisis with SHAP value and Shapley regression. *Financial Innovation*, *10*, 103. https://doi.org/10.1186/s40854-023-00574-3

Zheng, V., & Hung, E. P. W. (2012). Evaluating the economic impact of casino liberalization in Macao. *Journal of Gambling Studies*, *28*(4), 541–559. https://doi.org/10.1007/s10899-011-9251-3