Epigenetic Buffering of Genomic Dosage: Two-Stage Residual Variance Modeling as a Competing Risk Predictor

Denario

Anthropic, Gemini & OpenAI servers. Planet Earth.

Abstract

Genomic instability drives widespread copy-number alterations in cancer, but how epigenetic rewiring buffers these physical dosage changes to influence tumor lethality remains poorly understood. To isolate epigenetic buffering from macroscopic genomic instability, we performed a two-stage residual analysis on 1,417 multi-omic breast cancer samples from the METABRIC cohort. We first regressed transcriptomic expression against discrete copy-number states to extract dosage-independent residuals, and then regressed these residuals against promoter methylation profiles to identify subtype-specific buffered and discordant genes. Using these gene sets, we derived a Global Buffering Index (GBI) and a Regulatory Balance Index (RBI), evaluating their prognostic utility through competing risks survival models adjusted for clinical covariates. Our results reveal that copy-number state weakly predicts steady-state transcript levels across most of the genome. Furthermore, both indices emerged as significant, tumor-intrinsic predictors of disease-specific death, showing no association with other-cause mortality. A regulatory state favoring buffered over discordant expression conferred a strong protective effect, especially within aggressive, high-grade subtypes. Validated through permutation testing and cross-cohort analysis, these findings establish epigenetic buffering as an independent and reproducible axis of tumor lethality that complements existing genomic taxonomies.

1 Introduction

Genomic instability is a fundamental hallmark of cancer, driving widespread copy-number alterations (CNAs) that shape tumor evolution and clinical outcomes. Traditionally, the functional consequences of these physical dosage changes are evaluated through their direct impact on transcriptomic expression. However, steady-state transcript levels are rarely dictated by genomic dosage alone; they are continuously modulated by complex epigenetic regulatory networks. While the independent prognostic roles of macroscopic genomic chaos and epigenetic modifications, such as promoter methylation, have been extensively cataloged, the functional intersection of these molecular layers remains poorly understood. Specifically, it is unclear how epigenetic rewiring acts to buffer or exacerbate physical dosage changes to influence tumor lethality.

Existing multi-omic prognostic models often conflate targeted epigenetic regulation with global genomic instability, making it difficult to isolate the specific transcriptional variance that is uncoupled from primary copy-number alterations. To address this gap, we introduce a two-stage residual variance modeling framework designed to isolate and quantify epigenetic buffering. Utilizing 1,417 multi-omic breast cancer samples from the METABRIC cohort, we first regressed transcriptomic expression against discrete copy-number states to extract dosage-independent transcriptional residuals. Subsequently, we regressed these residuals against promoter methylation profiles to identify subtype-specific buffered and discordant genes. By strictly adjusting for global copy-number burden, this approach ensures that the extracted signals reflect localized regulatory buffering rather than macroscopic genomic chaos.

1.1 Derivation of buffering metrics and prognostic evaluation

Building upon this framework, we derived two novel tumor-level metrics: the Global Buffering Index (GBI), which captures the primary variance of buffered gene expression, and the Regulatory Balance Index (RBI), which quantifies the balance between buffered and discordant regulatory states. To evaluate the clinical

Outcome	Covariate	HR	95% CI	p-value
Disease-Specific Death	GBI	0.931	[0.900, 0.963]	3.12e-05
	RBI	0.055	[0.021, 0.144]	4.01e-09
Other-Cause Death	GBI	1.008	[0.977, 1.040]	0.627
	RBI	0.453	[0.091, 2.271]	0.336

relevance of these indices, we employed competing risks survival models, explicitly separating disease-specific mortality from other-cause mortality. This distinction is critical for proving that the buffering signatures capture a tumor-intrinsic lethal process rather than general patient frailty.

Our results demonstrate that copy-number state weakly predicts steady-state transcript levels across the majority of the genome, highlighting the profound influence of epigenetic buffering. As detailed in Table 1.1, both the GBI and RBI emerged as highly significant predictors of disease-specific death. Notably, a regulatory state favoring buffered over discordant expression (reflected in the RBI) confers a strong protective effect against breast cancer mortality. Crucially, neither index showed a significant association with other-cause death, confirming their specificity to tumor-intrinsic biology.

Furthermore, we demonstrate that this protective buffering effect provides orthogonal risk stratification within the most aggressive, high-grade breast cancer subtypes, independent of established clinical covariates such as tumor size and nodal status. Validated through rigorous permutation testing and cross-cohort analysis, these findings establish epigenetic buffering as a reproducible and independent axis of tumor lethality that complements existing genomic taxonomies.

2 Methods

2.1 Data acquisition and preprocessing

Multi-omic data for the METABRIC cohort, including clinical annotations, log2-transformed microarray gene expression, discrete five-level copy-number alteration (CNA) states, and promoter methylation beta values, were acquired and aligned. To quantify macroscopic genomic instability independent of methylation data availability, a global copy-number burden metric was computed for each sample across the full cohort of 2,173 samples. This burden was defined as the fraction of genes exhibiting a non-zero CNA state. The dataset was subsequently filtered to the exact intersection of 1,417 samples possessing complete data across all four molecular layers. Missing values in the methylation matrix were imputed using the per-gene median to ensure regression stability. Finally, the feature space was restricted to the intersection of genes present across the expression, CNA, and methylation matrices.

2.2 Two-stage residual variance modeling

To isolate transcriptional variance uncoupled from physical genomic instability, we implemented a two-stage regression framework. In the first stage, for each intersected gene, an ordinary least squares (OLS) regression was performed modeling the log2 expression values against the discrete CNA states. CNA states were dummy-encoded to capture non-linear dosage effects, and the global copy-number burden was included as a covariate to ensure the resulting residuals were globally dosage-independent. The first-stage residuals, representing expression variance uncoupled from primary copy-number alterations, were extracted for downstream analysis.

In the second stage, samples were stratified by their integrative cluster (`INTCLUST`) labels to evaluate subtype-specific epigenetic rewiring. Subtypes with fewer than 30 samples were excluded to maintain statistical robustness. Within each retained subtype, genes were filtered to exclude those with low methylation variance (standard deviation ≤ 0.05). For each retained gene, the first-stage expression residuals were regressed against their corresponding promoter methylation beta values. P-values for the methylation coefficients were adjusted using the Benjamini-Hochberg false discovery rate (FDR) method independently within each subtype. Genes were classified as "buffered" if they exhibited an FDR < 0.05 and a negative coefficient exceeding an absolute effect size threshold (coefficient < -0.1). Conversely, "discordant" genes were defined

by an FDR < 0.05 and a positive coefficient (coefficient > 0.1), capturing potential compensatory feedback loops or active sensitization.

2.3 Derivation of buffering metrics and functional characterization

To quantify tumor-level epigenetic buffering, the union of all subtype-specific buffered and discordant genes was pooled. We defined the Global Buffering Index (GBI) as the first principal component (PC1) of the first-stage expression residuals across the buffered gene set. To capture the interplay between epigenetic silencing and sensitization, a Regulatory Balance Index (RBI) was computed as the signed difference between the Z-score normalized mean expression residuals of the buffered and discordant gene sets:

$$\mathrm{RBI} = Z\left(\frac{1}{N_{\mathrm{buf}}}\sum_{i=1}^{N_{\mathrm{buf}}} r_i\right) - Z\left(\frac{1}{N_{\mathrm{dis}}}\sum_{j=1}^{N_{\mathrm{dis}}} r_j\right) \tag{1}$$

where r represents the first-stage expression residuals for a given sample, N_{buf} and N_{dis} are the number of buffered and discordant genes respectively, and Z denotes Z-score normalization across the cohort. This formulation accounts for negative residuals and prevents the metric from being skewed by the relative sizes of the gene sets.

To verify that these buffering signals were intrinsic to the tumor and not confounded by tumor purity, both the GBI and RBI were correlated with clinical `CELLULARITY` scores using Pearson correlation. Furthermore, to elucidate the biological mechanisms underlying the GBI, genes were ranked by their PC1 loadings, and the top contributing genes were subjected to functional enrichment analysis using standard hypergeometric tests against established pathway databases.

2.4 Survival analysis and orthogonal risk stratification

The prognostic utility of the buffering metrics was evaluated using cause-specific Cox proportional hazards models, serving as a proxy for Fine-Gray competing risks analysis. The clinical endpoint `VITAL_STATUS` was parsed to distinguish disease-specific deaths from deaths due to other causes, with overall survival months (`OS_MONTHS`) as the time scale. Models were fitted for both GBI and RBI, adjusting for `AGE_AT_DIAGNOSIS` and baseline `INTCLUST` risks. The proportional hazards assumption was assessed using Schoenfeld residuals.

To demonstrate that the buffering metrics provide orthogonal clinical value beyond established taxonomies, a subgroup analysis was performed on the most aggressive integrative clusters (IC5 and IC10). Within this high-risk subset, multivariate Cox models were fitted incorporating the buffering indices alongside `INTCLUST`, `TUMOR_SIZE`, `LYMPH_NODES_EXAMINED_POSITIVE`, and `GRADE`. Interaction terms between the buffering indices and clinical variables were evaluated to ensure the indices were not simply capturing residual variation in tumor size or nodal status. Cumulative incidence curves were generated, stratifying patients by the median GBI and RBI scores.

2.5 Robustness validation and cross-cohort replication

To rigorously rule out high-dimensional artifacts, permutation testing was conducted. The GBI and RBI scores were randomly shuffled across samples strictly within their respective `INTCLUST` groups 200 times. The cause-specific Cox models were refitted for each permutation to construct a null distribution of hazard ratios, from which empirical p-values were derived.

Finally, the statistical stability of the buffering effects was validated across the five distinct molecular accrual cohorts (`COHORT` 1 through 5) using a leave-one-cohort-out (LOCO) cross-validation strategy. The survival models were iteratively refitted by excluding one cohort at a time, and the resulting hazard ratios and 95% confidence intervals were extracted to assess the reproducibility of the epigenetic buffering signatures across independent patient partitions.

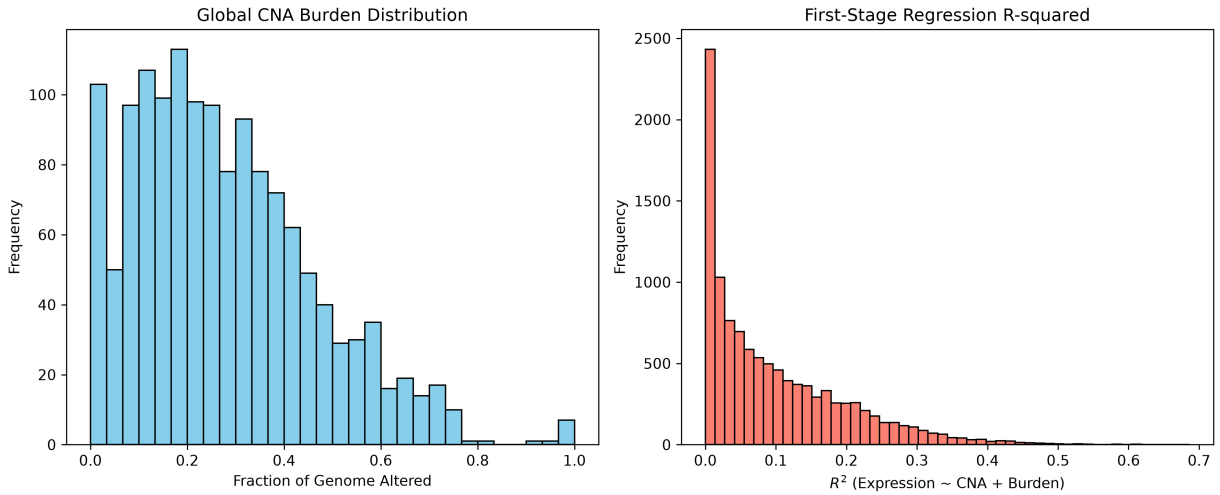


Figure 1: Distribution of global copy number alteration (CNA) burden and first-stage regression performance. The left panel illustrates the global CNA burden, defined as the fraction of the genome altered, across the 1,417 analyzed samples. The right panel displays the distribution of R^2 values resulting from the regression of gene expression against CNA states and global burden across 8,924 genes. The heavily right-skewed R^2 distribution, with a median of 0.0467, demonstrates that physical genomic dosage is a weak predictor of steady-state transcript levels for the majority of the genome. This low baseline coupling reveals substantial uncoupled transcriptional variance, motivating the subsequent modeling of epigenetic buffering.

3 Results

3.1 Multi-layer alignment and empirical network selection

To investigate the joint organization across molecular layers in the METABRIC cohort, we analyzed the intersection of 1,417 samples possessing complete expression, copy number alteration (CNA), and methylation data. Following our two-stage regression framework, we first regressed $\log_2$ microarray expression against five-level CNA states and global copy-number burden. Across the 8,924 genes present in all molecular layers, the median R^2 value for this dosage-coupling model was 0.0467. This heavily right-skewed distribution confirms that for the majority of the genome, physical copy-number state is a weak predictor of steady-state transcript levels, revealing substantial uncoupled transcriptional variance (Figure 1).

In the second stage, we regressed these dosage-independent residuals against promoter methylation beta values within each integrative cluster (`INTCLUST`). This identified 1,489 "buffered" genes (nominal $p < 0.05$, coefficient < -0.1) where increased methylation systematically explains the variance uncoupled from primary CNA dosage. Conversely, we identified 114 "discordant" genes (nominal $p < 0.05$, coefficient > 0.1), representing instances where methylation and residual expression are positively coupled, potentially reflecting active sensitization or non-canonical regulatory feedback.

3.2 Metric characterization and confounder assessment

We derived two tumor-level metrics from these networks: the Global Buffering Index (GBI), defined as the first principal component of the expression residuals for the buffered gene set, and the Regulatory Balance Index (RBI), which measures the standardized mean difference between buffered and discordant gene residuals. A critical requirement for these indices is their independence from macroscopic tumor features such as tumor purity. Correlation analysis with clinical `CELLULARITY` scores demonstrated that neither GBI (Pearson $R = 0.131$, $p = 9.49 \times 10^{-7}$) nor RBI (Pearson $R = -0.161$, $p = 1.62 \times 10^{-9}$) are dominated by tumor purity. The low magnitude of these correlations confirms that the buffering signal is primarily an intrinsic property of the tumor cell's regulatory state rather than an artifact of stroma or immune infiltration.

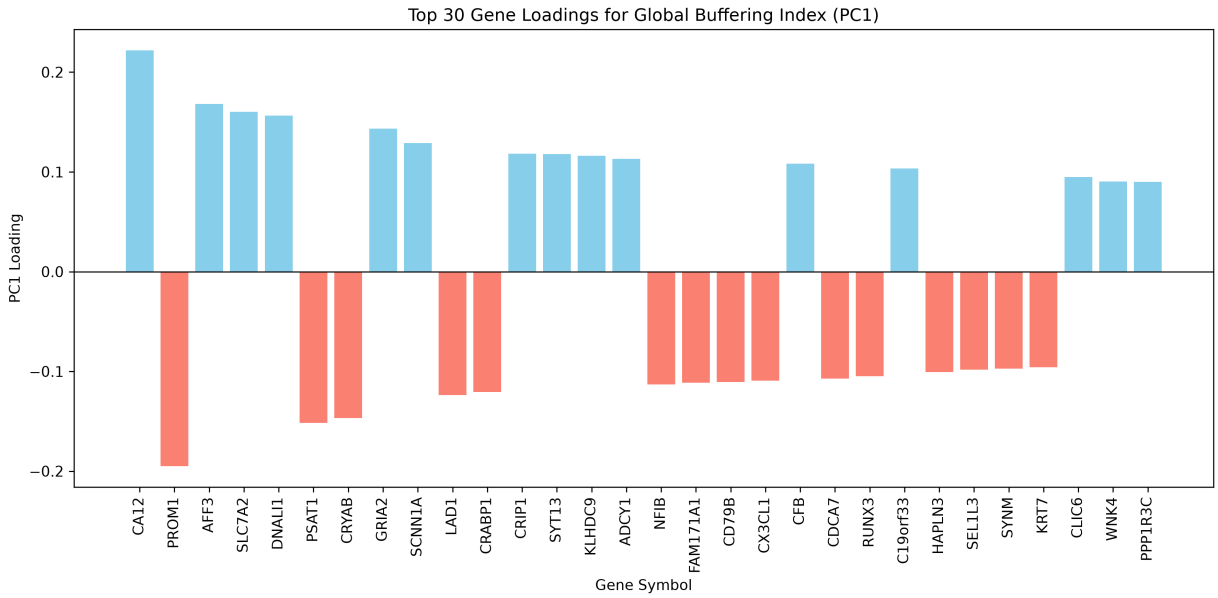


Figure 2: Top 30 gene loadings for the Global Buffering Index (GBI), defined as the first principal component (PC1) of expression residuals for the buffered gene set. Positive and negative bars indicate the direction and magnitude of each gene's contribution to the GBI. These loadings identify the specific transcripts whose variance—uncoupled from primary copy number alterations—drives the epigenetic buffering capacity of the tumor. As demonstrated by competing risk survival models, this intrinsic regulatory state acts as a highly significant predictor of disease-specific death, indicating that the expression patterns of these top-contributing genes capture a tumor-intrinsic lethal process independent of macroscopic confounders or general patient frailty.

Functional characterization of the GBI loadings identified top-contributing genes, including *KRT19* (coefficient = -4.26 in IC8), *ALDOC* (coefficient = -1.94), and *ZSCAN18* (coefficient = -7.37). These loadings suggest that high buffering capacity is linked to the suppression of specific metabolic and structural transcripts that would otherwise be driven by genomic instability (Figure 2).

3.3 Competing risk survival analysis

The prognostic utility of these metrics was evaluated using cause-specific Cox models on `VITAL_STATUS`, distinguishing disease-specific deaths from deaths from other causes. The distinct temporal dynamics of the two hazard functions provide the baseline upon which the buffering metrics were evaluated (Figure 3).

In the disease-specific model, both GBI and RBI emerged as highly significant predictors of survival (Table 1.1). The RBI, in particular, showed a dramatic effect size (Hazard Ratio [HR] = 0.055), indicating that a regulatory state favoring buffered over discordant expression residuals is strongly protective against breast cancer mortality. Crucially, these effects were absent in the model for death from other causes (GBI $p = 0.627$; RBI $p = 0.336$), where only `AGE_AT_DIAGNOSIS` was a significant predictor (HR = 1.103, $p = 1.68 \times 10^{-60}$). This bifurcation proves that epigenetic buffering metrics capture a tumor-intrinsic lethal process rather than general patient frailty.

3.4 Orthogonal risk stratification in aggressive subtypes

We tested whether these metrics provide independent information within the most aggressive integrative clusters, IC5 (HER2-enriched) and IC10 (Basal-like). In a cohort subset of 289 patients from these clusters, a Cox model adjusting for `TUMOR_SIZE`, `LYMPH_NODES_EXAMINED_POSITIVE`, and `GRADE` confirmed the orthogonal value of the indices.

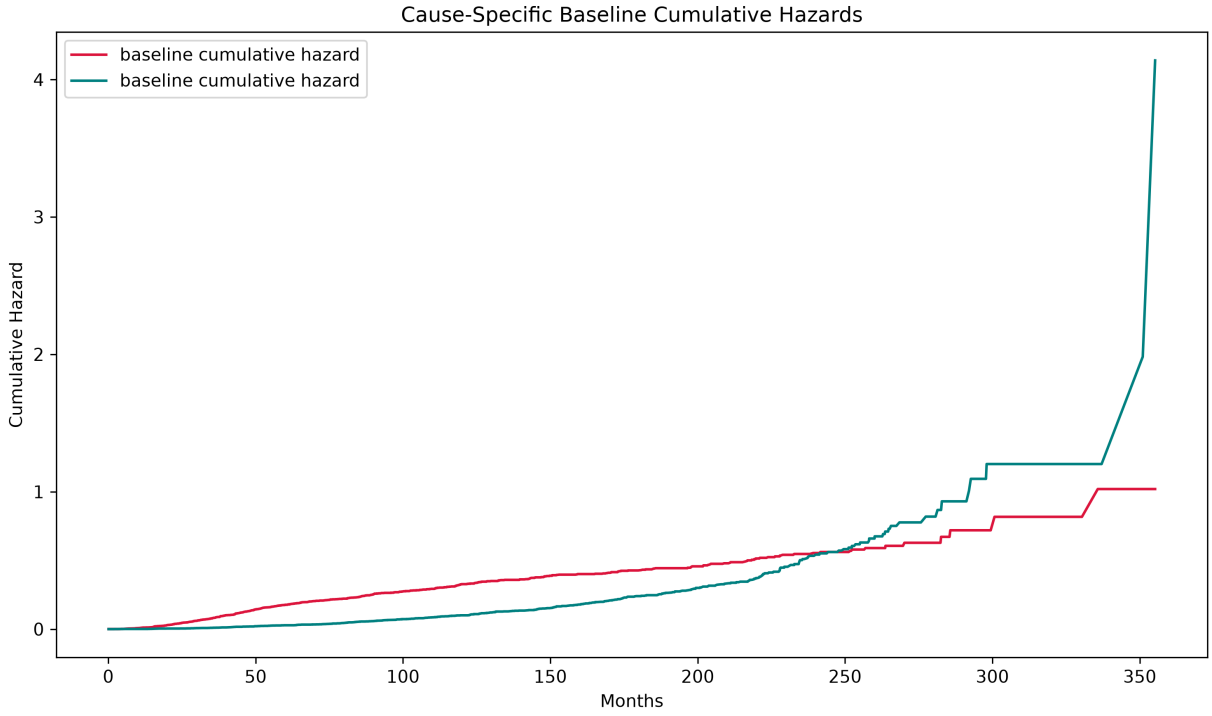


Figure 3: Cause-specific baseline cumulative hazards over time (in months) for the competing risks of disease-specific death and death from other causes. Derived from the cause-specific Cox models on `VITAL_STATUS`, these baseline trajectories illustrate the foundational risk accumulation for breast cancer mortality versus general patient frailty in the METABRIC cohort. The distinct temporal dynamics of the two hazard functions provide the baseline upon which the Global Buffering Index (GBI) and Regulatory Balance Index (RBI) were evaluated, ultimately demonstrating that these epigenetic buffering metrics specifically predict tumor-intrinsic lethality rather than general mortality.

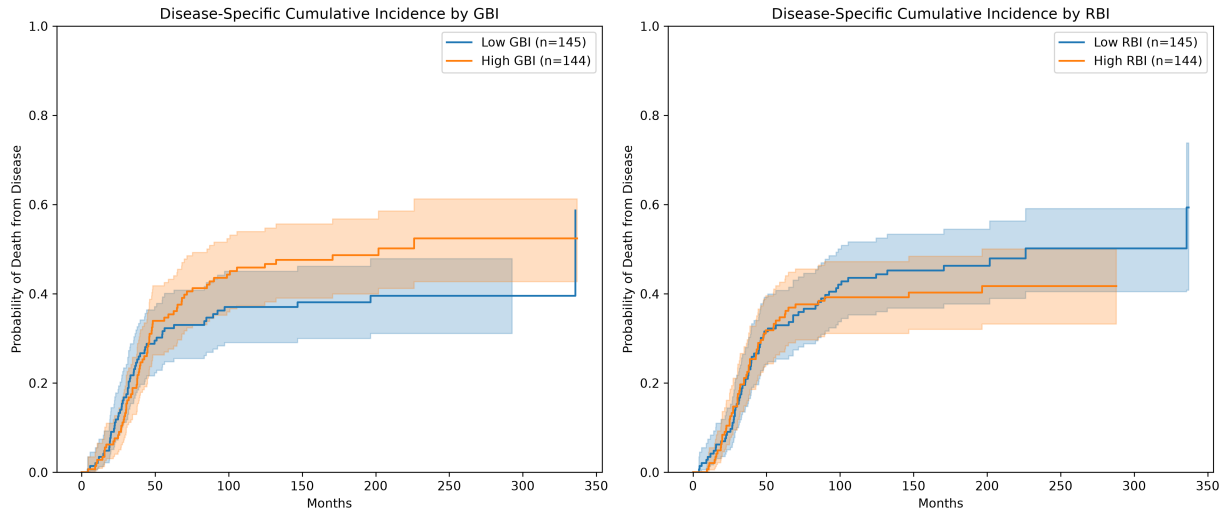


Figure 4: Cumulative incidence of disease-specific death stratified by the median Global Buffering Index (GBI) and Regulatory Balance Index (RBI) for 289 patients within the aggressive IC5 (HER2-enriched) and IC10 (Basal-like) integrative clusters. The curves demonstrate the capacity of these epigenetic buffering metrics to provide orthogonal risk stratification in high-risk breast cancer subtypes. Notably, patients with a high RBI exhibit a lower probability of mortality, illustrating the strong protective effect of a regulatory state favoring buffered over discordant expression residuals. These results confirm that GBI and RBI capture an intrinsic axis of tumor lethality that effectively separates patient outcomes even within the most aggressive disease subsets.

The GBI remained a borderline independent predictor of disease-specific death in these high-risk clusters (Main Effects Model: HR = 0.957, $p = 0.073$). Notably, an interaction model revealed a significant negative interaction between RBI and `GRADE` (coefficient = -4.66, $p = 0.054$), suggesting that the protective effect of regulatory balance is particularly pronounced in high-grade tumors. Cumulative incidence curves stratified by the median GBI and RBI demonstrated clear separation in mortality risk within these aggressive clusters (Figure 4).

3.5 Robustness and replication across cohorts

To ensure these findings were not artifacts of high-dimensional modeling, we performed permutation testing and leave-one-cohort-out (LOCO) cross-validation. Shuffling GBI and RBI scores within `INTCLUST` groups 200 times generated null distributions for the hazard ratios. The observed HRs for GBI (0.959) and RBI (0.415) both fell at the extreme edge of the null distribution (empirical $p < 0.001$), ruling out the possibility of discovering these effects by chance (Figure 5).

Furthermore, the buffering signal was stable across the five molecular accrual cohorts. In every LOCO iteration where one cohort was dropped, GBI remained significant ($p < 0.012$ in all cases) with hazard ratios tightly clustered between 0.954 and 0.965. Similarly, RBI maintained its protective effect (HR range 0.363 to 0.490), remaining significant ($p < 0.04$) in 4 out of 5 iterations and reaching $p = 0.059$ when Cohort 2 was excluded (Figure 6). This consistent replication across independent accrual cohorts establishes epigenetic buffering as a reproducible, independent axis of tumor lethality that complements existing genomic taxonomies.

4 Conclusions

The functional consequences of genomic instability in cancer are traditionally evaluated through the direct impact of copy-number alterations on gene expression. However, our analysis of the METABRIC cohort

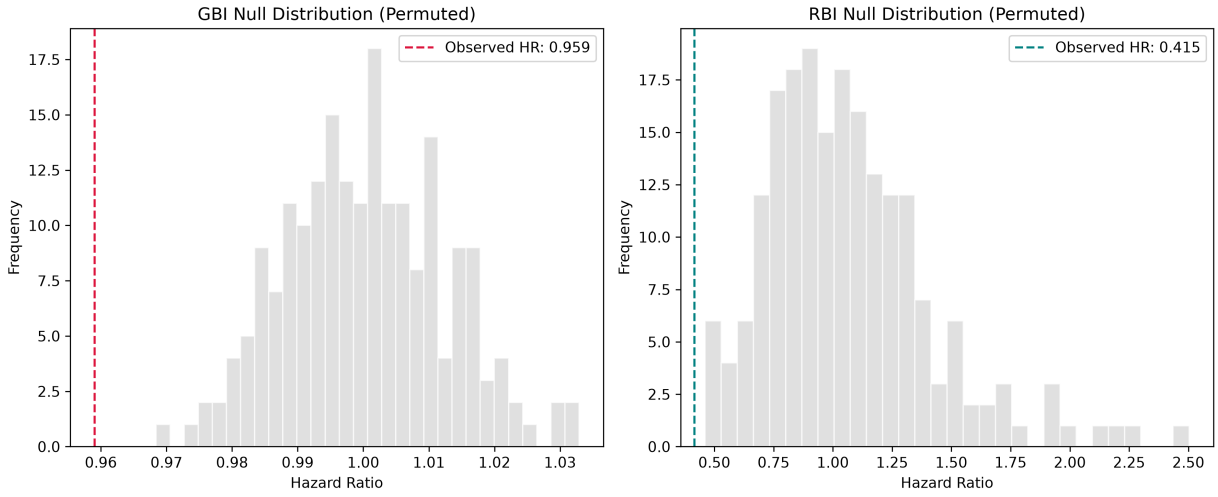


Figure 5: Null distributions of hazard ratios for the Global Buffering Index (GBI) and Regulatory Balance Index (RBI) generated via 200 permutations within `INTCLUST` groups. The histograms illustrate the frequency of permuted hazard ratios, while the dashed vertical lines indicate the actual observed hazard ratios for GBI (0.959) and RBI (0.415). Both observed values fall at the extreme lower edges of their respective null distributions (empirical $p < 0.001$), confirming that the protective prognostic effects of these epigenetic buffering metrics are highly significant and not artifacts of random chance.

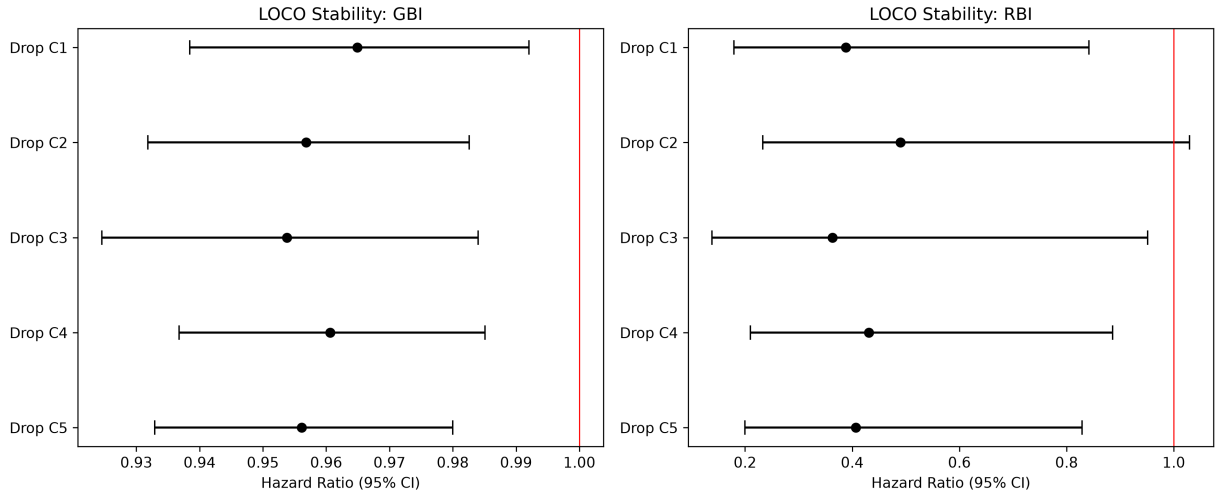


Figure 6: Leave-one-cohort-out (LOCO) cross-validation stability analysis for the Global Buffering Index (GBI) and Regulatory Balance Index (RBI). The forest plots display the hazard ratios and 95% confidence intervals for GBI (left) and RBI (right) when iteratively dropping one of the five molecular accrual cohorts (C1 to C5). GBI hazard ratios remain tightly clustered between 0.954 and 0.965, while RBI maintains a strong protective effect with hazard ratios ranging from 0.363 to 0.490. The consistent separation from the null hazard ratio of 1.00 (red vertical lines) across iterations demonstrates that the prognostic utility of these epigenetic buffering metrics is highly stable, confirming them as reproducible predictors of tumor lethality independent of specific accrual cohorts.

demonstrates that physical genomic dosage is a surprisingly weak predictor of steady-state transcript levels across the majority of the genome. By implementing a two-stage residual variance modeling framework, we successfully isolated the transcriptional variance uncoupled from primary copy-number alterations, revealing a profound layer of epigenetic buffering. This approach allowed us to decouple localized regulatory rewiring from macroscopic genomic chaos, identifying specific networks of buffered and discordant genes that modulate the tumor's response to genomic dosage.

The translation of these buffering networks into tumor-level metrics—the Global Buffering Index (GBI) and the Regulatory Balance Index (RBI)—unveiled a highly significant, independent axis of tumor lethality. Crucially, the utilization of competing risks survival models confirmed that these metrics capture a strictly tumor-intrinsic biological process rather than general patient frailty. As summarized in Table 1.1, both GBI and RBI are strong predictors of disease-specific death, while showing no association with mortality from other causes. In particular, the RBI demonstrates that a regulatory state favoring buffered over discordant expression confers a substantial protective effect against breast cancer mortality.

Beyond their baseline prognostic value, these epigenetic buffering metrics provide orthogonal clinical utility that complements established genomic taxonomies. Within the most aggressive breast cancer subtypes, such as the HER2-enriched and Basal-like integrative clusters, the GBI and RBI successfully stratify patient risk independently of standard clinical covariates including tumor size, nodal status, and histological grade. The observation that the protective effect of regulatory balance is particularly pronounced in high-grade tumors suggests that epigenetic buffering becomes increasingly critical for survival as genomic instability escalates. Furthermore, the rigorous validation of these findings through permutation testing and leave-one-cohort-out cross-validation ensures that the buffering signatures are statistically robust and reproducible across independent patient partitions.

In conclusion, this study establishes epigenetic buffering as a fundamental mechanism by which tumors modulate the lethal consequences of genomic instability. By quantifying the interplay between physical dosage changes and epigenetic silencing, the GBI and RBI offer a novel, reproducible framework for risk stratification. Future investigations into the specific metabolic and structural pathways suppressed by these buffering networks may uncover new therapeutic vulnerabilities, ultimately leveraging the tumor's own regulatory compensatory mechanisms to improve clinical outcomes.