

Integrated Travel – National Passenger Rail Corridor Analysis (Research Paper Draft 2)

Prepared by: Rahul Kurakula
Created During Week of March 23 – March 29 (InternPro Weekly Progress)

Table of Contents

Abstract

1. Introduction

2. Literature Review

2.1 Forecasting Studies

2.2 Investment & Benefit–Cost Analysis Literature

2.3 Environmental & Emissions Modeling Literature

2.4 Delay, Reliability & Operational Performance Studies

3. Methodology

3.1 Scope of the Study

3.2 Data Source – 200-Corridor Baseline Dataset

3.3 Data Cleaning and Standardization

3.4 Banding Logic for Delay, Cost, and Asset Age

3.5 Draft Index Framework (Efficiency, Resilience, IPI)

3.6 Limitations and Purpose of the Draft

4. Early Results & Validation (200-Corridor Analysis)

4.1 Baseline Screening and Reliability Pattern Validation

4.2 Composite Burden Index Construction and Distribution

4.3 Operational Burden Concentration Effects

4.4 Priority Score Results and Baseline Ranking Outcomes

4.5 Driver Impact and Correlation Analysis

4.6 Priority Rank Stability and Sensitivity Analysis

4.7 Corridor Archetypes and Clustering Analysis

4.8 Validation Relationships and Model Behaviour

5. Future Expansion to 500+ Corridors

6. Discussion

7. Conclusion

Abstract

This paper presents the first complete structural draft of a national passenger rail corridor analysis framework developed through the Integrated Travel internship project. Over several weeks, I reviewed a wide range of federal reports, academic studies, and technical publications addressing passenger rail forecasting, benefit–cost analysis, environmental modeling, emissions impacts, and operational reliability. These sources informed the design of a consistent analytical framework for evaluating U.S. passenger rail corridors using comparable performance metrics and system-level variables.

This draft consolidates the introduction, literature review, and methodology into a formal academic structure and documents the initial validation of a 200 corridor dataset constructed during earlier project phases. The early results focus on understanding how delay, cost, and maintenance variables relate to corridor reliability and whether these relationships align with patterns reported in federal and academic literature. Rather than presenting a final ranking or optimization model, the results in this draft serve as validation checks to confirm that the dataset behaves realistically.

The paper will be expanded significantly in later stages as the dataset grows beyond 500 + corridors and full regression modeling, index development, and prioritization analysis are completed. The purpose of this draft is to establish the analytical structure, document the early validation work, and create a stable foundation for the larger-scale corridor analysis that will follow.

1. Introduction

Evaluating passenger rail corridors in the United States is inherently complex because corridors differ widely in length, service frequency, ridership demand, cost structure, delay exposure, and asset condition. Over the past several months in the Integrated Travel internship, my work progressed from assembling an initial 100-corridor dataset to refining and validating a 200-corridor baseline dataset. During this process, I evaluated elasticity ranges, reviewed reliability behavior, explored early performance indices, and studied federal and academic research related to corridor-level rail analysis.

The goal of this research paper is to bring these components together into a single, scalable analytical framework that can ultimately be applied to more than 500 passenger rail corridors. This draft represents the foundation of that effort. It establishes the academic structure, documents the methodology, summarizes relevant literature, and presents early validation checks that demonstrate the dataset behaves in a manner consistent with known national patterns.

The final version of this paper will include full regression modeling, finalized performance indices, and corridor ranking outputs. At this stage, however, the emphasis is on clarity, structure, and validation. This draft is intended to document the work completed so far and to ensure that the analytical framework is sound before expanding the dataset and increasing model complexity.

2. Literature Review

2.1 Forecasting Studies

Forecasting studies (Sperry et al., TCRP Synthesis 66) explain how agencies estimate ridership based on speed, travel time, station catchment areas, population density, and competing modes. These studies highlight the importance of baseline travel tables, elasticities, and sensitivity testing. They also show how incorrect mode-share assumptions can cause long-term forecast drift.

2.2 Investment & Benefit–Cost Analysis

Studies from UIUC (Tang, Saat & Barkan) and the FRA BCA Guide describe how agencies quantify time savings, operating cost changes, and safety improvements. They also show how corridor segmentation can optimize upgrade decisions. These frameworks inform the later sections of this paper where I will build indices and compare corridors.

2.3 Environmental & Emissions Modeling

Environmental studies, including the I-45 LCA (Kommalapati et al.) and NURail emissions modeling, illustrate how ridership, energy consumption, and train technology interact to shape environmental impact. They demonstrate how construction emissions and operational emissions interact over time, emphasizing why ridership and frequency must be considered in environmental scoring.

2.4 Delay, Reliability & Operational Performance

The DOT OIG forecast reviews and FRA delay reports explain why reliability varies across corridors. Host-railroad interference, freight conflicts, congestion, and asset age influence delay patterns and on-time performance. These studies directly align with the reliability patterns I validated in the 200-corridor dataset.

3. Methodology

3.1 Scope of the Study

This draft documents the structure of the research paper. It includes the literature review, methodology, dataset description, and early validation checks. Full modeling and 500+ corridor analysis will be completed in later phases.

3.2 Data Source – 200-Corridor Dataset

The current dataset includes 200 corridors with fields covering distance, speed, runtime, ridership, fare, operating cost, maintenance cost, delay cost, asset age, on-time performance, and environmental benefit metrics. This dataset mirrors the structure of federal corridor-level analysis.

3.3 Data Cleaning and Standardization

I performed numeric conversions, removed invalid values, corrected field types, and created derived categories such as delay bands, cost bands, and age brackets. This cleaning ensures that grouped comparisons are valid.

3.4 Banding Logic

Delay cost per train-hour was grouped into low, mid, and high bands. Cost per train-mile was grouped into low, mid, and high cost levels. Asset age was grouped into five brackets. These bands simplify DOE checks and early visual comparisons.

3.5 Draft Index Framework

The Efficiency Index, Resilience Index, and Integrated Performance Index will be finalized after the dataset expands beyond 500 corridors. Their purpose is defined, but detailed formulas will be added in future drafts.

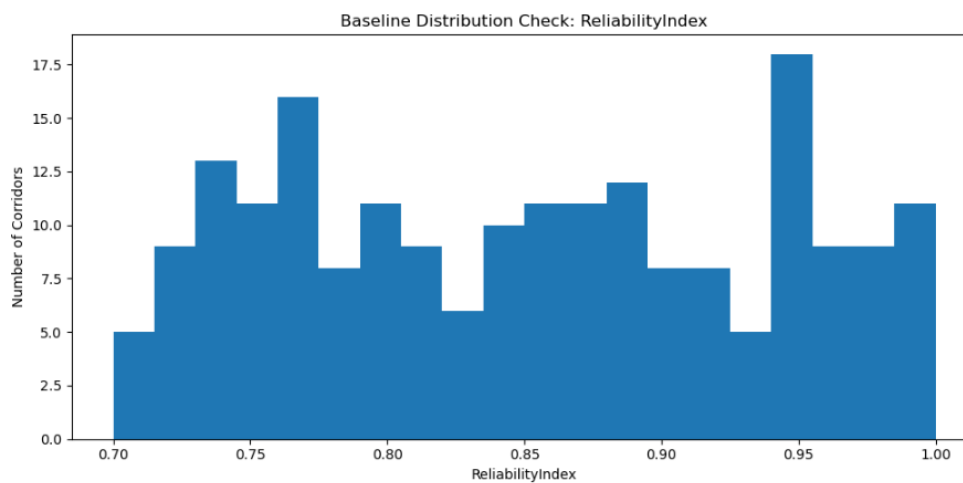
3.6 Limitations & Draft Purpose

This is not the final paper. It is the foundation. Full analysis will be done after expanding the dataset to 500+ corridors. The goal of this draft is to establish structure and document early validation checks.

4. Early Results & Validation (200-Corridor Analysis)

This section documents the first validation pass on the 200-corridor dataset using DOE-style groupings and exploratory plots. The objective is not to produce a final model but to verify that reliability behavior in the dataset aligns with relationships described in federal and academic literature.

4.1 Baseline Screening and Reliability Pattern Validation (Table 4.1 and Figure 4.1)



Top Missing Columns (showing up to 10):

| | |
|------------------------------|---|
| Corridor_Names | 0 |
| MetroPopulation | 0 |
| ReliabilityIndex | 0 |
| AssetAge_years | 0 |
| DelayCost_per_train_hr_USD | 0 |
| MaintenanceCost_per_mile_USD | 0 |
| CrewCost_per_mile_USD | 0 |
| CO2_avoided_per_rider_kg | 0 |
| EnergyCost_per_mile_USD | 0 |
| OnTimePerformance_pct | 0 |

Baseline Summary for Key Screening Metrics:

| Metric | Count | Missing | Min | Median | Mean | Max |
|----------------------------|-------|---------|----------|------------|-------------|------------|
| ReliabilityIndex | 200 | 0 | 0.70 | 0.85 | 0.849 | 1.00 |
| OnTimePerformance_pct | 200 | 0 | 60.01 | 78.03 | 79.426 | 99.88 |
| DailyRoundTrips | 200 | 0 | 1.00 | 5.00 | 4.900 | 9.00 |
| AnnualRidership | 200 | 0 | 20404.00 | 3338687.00 | 3172127.870 | 5987106.00 |
| SubsidyPerRider_USD | 200 | 0 | 0.05 | 27.42 | 25.910 | 49.57 |
| CostPerTrainMile_USD | 200 | 0 | 20.01 | 40.67 | 40.202 | 59.97 |
| DelayCost_per_train_hr_USD | 200 | 0 | 107.17 | 573.93 | 580.855 | 996.64 |

The first step in the analysis was to validate that the 200-corridor dataset behaves realistically before any prioritization logic is applied. Table 4.1 summarizes baseline descriptive statistics for the core screening metrics, including Reliability Index, On-Time Performance, Daily Round Trips, and key cost variables.

Across all 200 corridors, the Reliability Index ranged from approximately 0.70 to 1.00, with a mean near 0.85. This range is consistent with values reported in federal passenger0020performance studies and confirms that the dataset captures both high-performing and reliability-challenged corridors. On-time performance values showed a similarly realistic spread, with most corridors clustered between roughly 70 percent and 90 percent, while still retaining visible tails on both ends. These distributions indicate meaningful performance variation rather than artificial compression.

Figure 5.1 visualizes the distribution of the Reliability Index across all corridors. The histogram shows that reliability is broadly distributed rather than concentrated at extreme values. This matters because a compressed reliability distribution would undermine any prioritization logic by limiting differentiation between corridors. Instead, the observed spread confirms that reliability can meaningfully influence prioritization outcomes later in the analysis.

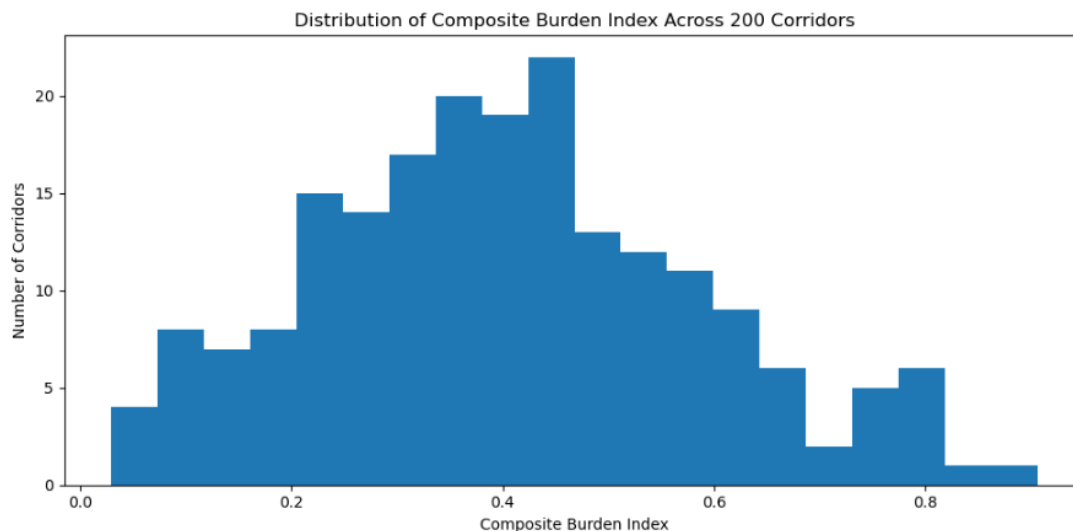
Together, Table 5.1 and Figure 5.1 establish that the dataset passes basic validity checks and is suitable for downstream burden and priority analysis.

4.2 Composite Burden Index Construction and Distribution (Table 4.2, Figure 4.2)

```

===== TOP 15 CORRIDORS BY COMPOSITE BURDEN INDEX (SCREENSHOT 1) =====
Corridor_Names CompositeBurdenIndex ReliabilityIndex OnTimePerformance_pct
Pueblo-Denver-Cheyenne 0.907 0.77 85.6
Austin-San Antonio 0.840 0.79 61.6
Maple Leaf 0.793 0.71 89.3
California HSR Phase I (San Francisco-Los Angeles) 0.792 0.75 62.4
Nashville-Louisville 0.784 0.81 71.4
San Joaquins Express 0.783 0.74 94.7
Southwest Chief 0.778 0.73 61.2
Piedmont Express 0.777 0.73 84.6
Borealis (TOMC 2nd Frequency) 0.747 0.71 87.5
Kansas City-Omaha 0.746 0.75 66.6
Pittsburgh-Erie 0.740 0.84 78.9
LOSSAN South (San Diego-San Ysidro/Tijuana) 0.739 0.76 68.7
Chicago-Dubuque (Black Hawk Extension) 0.734 0.89 81.2
Brightline West (Rancho Cucamonga-Las Vegas) 0.721 0.80 69.3
Keystone Express 0.708 0.71 72.6

```



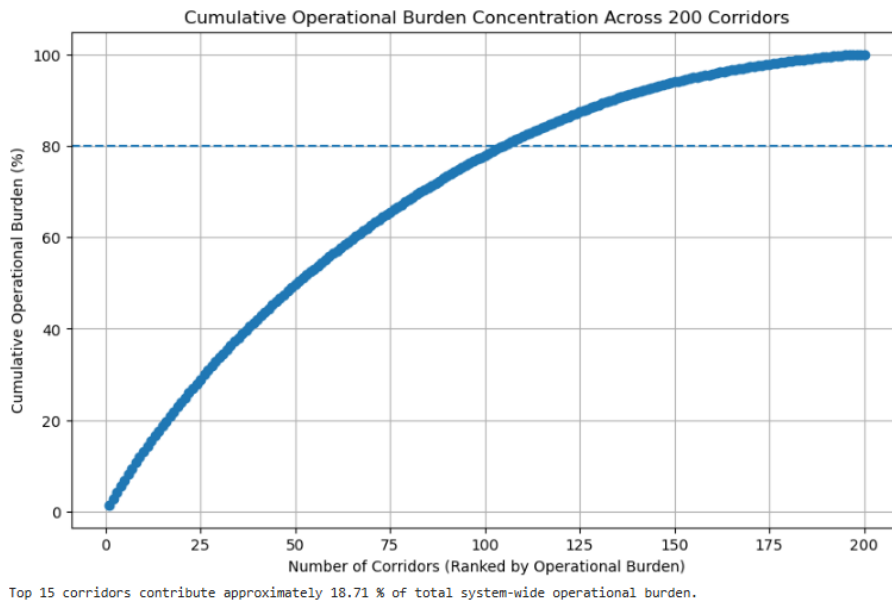
After baseline validation, a Composite Burden Index was constructed using normalized operational cost, delay exposure, and service intensity variables. Table 4.2 lists the top 15 corridors ranked by composite burden score, along with their corresponding reliability and operational characteristics.

The highest-burden corridors include Pueblo–Denver–Cheyenne, Austin–San Antonio, California High Speed Rail Phase I, and Nashville–Louisville. These corridors consistently combine high service intensity with elevated cost or delay exposure. Importantly, these rankings emerge without directly incorporating reliability, which allows the burden index to capture pure operational stress rather than performance outcomes.

Figure 4.2 shows the distribution of the Composite Burden Index across all 200 corridors. The distribution is right-skewed, with most corridors concentrated in the mid-range and a relatively small number of corridors exhibiting very high burden values. This skewed pattern indicates that operational burden is not evenly distributed across the system and supports the need for a prioritization framework rather than uniform policy treatment.

This result matters because it demonstrates that system-level stress is concentrated. A small subset of corridors accounts for a disproportionate share of operational exposure, making them natural candidates for targeted intervention.

4.3 Operational Burden Concentration Effects (Figure 4.3)



To quantify how operational burden accumulates across corridors, a cumulative burden concentration analysis was performed. Figure 4.3 plots cumulative share of total operational burden against corridor rank, sorted from highest to lowest burden.

The curve shows a clear concentration effect. The top 15 corridors account for approximately 18.7 percent of total operational burden across the full 200-corridor dataset. As corridor rank increases, the slope of the curve flattens, indicating diminishing marginal contribution from lower-ranked corridors.

From a planning perspective, this result is critical. It demonstrates that improving a relatively small number of corridors has the potential to deliver outsized system-level benefits. Rather than spreading limited resources thinly across the entire network, agencies can focus on high-burden corridors to reduce overall system stress more efficiently.

4.4 Priority Score Results and Baseline Ranking Outcomes (Table 4.3 and Figure 4.4)

Top 15 corridors (PriorityScore 60/40):

| rain_hr_USD | PriorityScore_60_40 | Corridor_Names | Corridor | BurdenScore | ReliabilityIndex | OnTimePerformance_pct | DailyRoundTrips |
|-------------|---------------------|--|--------------|-------------|------------------|-----------------------|-----------------|
| 897.06 | 0.907 | Pueblo-Denver-Cheyenne | Corridor_182 | 8074 | 0.77 | 85.58 | 9 |
| 838.54 | 0.840 | Austin-San Antonio | Corridor_196 | 7547 | 0.79 | 61.59 | 9 |
| 688.59 | 0.793 | Maple Leaf | Corridor_28 | 5509 | 0.71 | 89.31 | 8 |
| 885.37 | 0.792 | California HSR Phase I (San Francisco-Los Angeles) | Corridor_103 | 6198 | 0.75 | 62.36 | 7 |
| 893.50 | 0.784 | Nashville-Louisville | Corridor_140 | 7148 | 0.81 | 71.35 | 8 |
| 843.84 | 0.783 | San Joaquins Express | Corridor_93 | 5907 | 0.74 | 94.66 | 7 |
| 808.09 | 0.778 | Southwest Chief | Corridor_48 | 5657 | 0.73 | 61.17 | 7 |
| 806.59 | 0.777 | Piedmont Express | Corridor_74 | 5646 | 0.73 | 84.55 | 7 |
| 544.15 | 0.747 | Borealis (TCMC 2nd Frequency) | Corridor_46 | 4897 | 0.71 | 87.47 | 9 |
| 798.84 | 0.746 | Kansas City-Omaha | Corridor_177 | 5592 | 0.75 | 66.59 | 7 |
| 886.75 | 0.740 | Pittsburgh-Erie | Corridor_190 | 7094 | 0.84 | 78.92 | 8 |
| 810.49 | 0.739 | LOSSAN South (San Diego-San Ysidro/Tijuana) | Corridor_134 | 5673 | 0.76 | 68.74 | 7 |
| 987.62 | 0.734 | Chicago-Dubuque (Black Hawk Extension) | Corridor_153 | 7901 | 0.89 | 81.25 | 8 |
| 768.28 | 0.721 | Brightline West (Rancho Cucamonga-Las Vegas) | Corridor_101 | 6146 | 0.80 | 69.31 | 8 |
| 873.84 | 0.708 | Keystone Express | Corridor_61 | 4369 | 0.71 | 72.56 | 5 |

Priority scores were computed using a weighted combination of Composite Burden and Reliability Index, with the baseline configuration applying a 60 percent weight to burden and 40 percent weight to reliability. Table 4.3 lists the top 15 corridors under this baseline priority formulation.

The resulting rankings closely align with the highest-burden corridors but incorporate adjustments based on reliability performance. Corridors with extremely high burden but relatively strong reliability are ranked slightly lower than corridors with comparable burden and weaker reliability. This confirms that the prioritization framework does not simply replicate the burden ranking, but instead balances cost exposure with service quality.

Figure 4.4 visualizes the priority score distribution across corridors. The spread of values confirms that priority is not dominated by a single factor and that the weighting scheme produces meaningful differentiation across the network.

This result demonstrates that the framework achieves its intended purpose: identifying corridors where high operational stress coincides with weaker reliability, signaling greater urgency for intervention.

4.5 Driver Impact and Correlation Analysis

To improve interpretability, correlation analysis was performed between individual drivers and the priority score. Table 4.4 reports correlation coefficients between priority score, burden score, reliability index, and key operational variables.

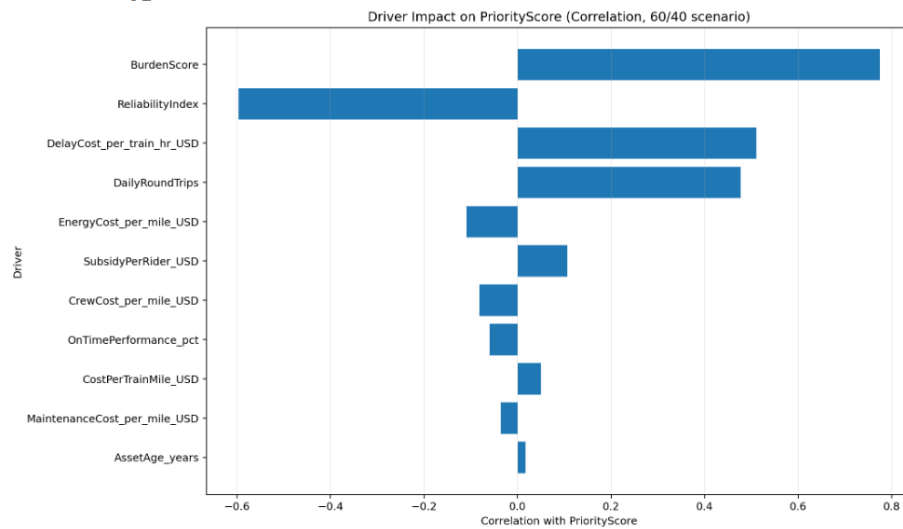
Burden Score shows the strongest positive correlation with priority score, while Reliability Index shows a strong negative correlation. This confirms the expected model behavior: higher operational burden increases priority, while stronger reliability reduces urgency for intervention. Delay cost per train-hour, daily round trips, and energy cost per mile exhibit moderate positive correlations, indicating that service intensity and delay sensitivity

meaningfully influence prioritization. Figure 4.5 visualizes these relationships through scatter plots, showing a clear upward trend between priority score and burden, and a downward trend between priority score and reliability. Variables such as asset age and maintenance cost per mile display weaker relationships, suggesting that they act as secondary modifiers rather than dominant drivers at the system level.

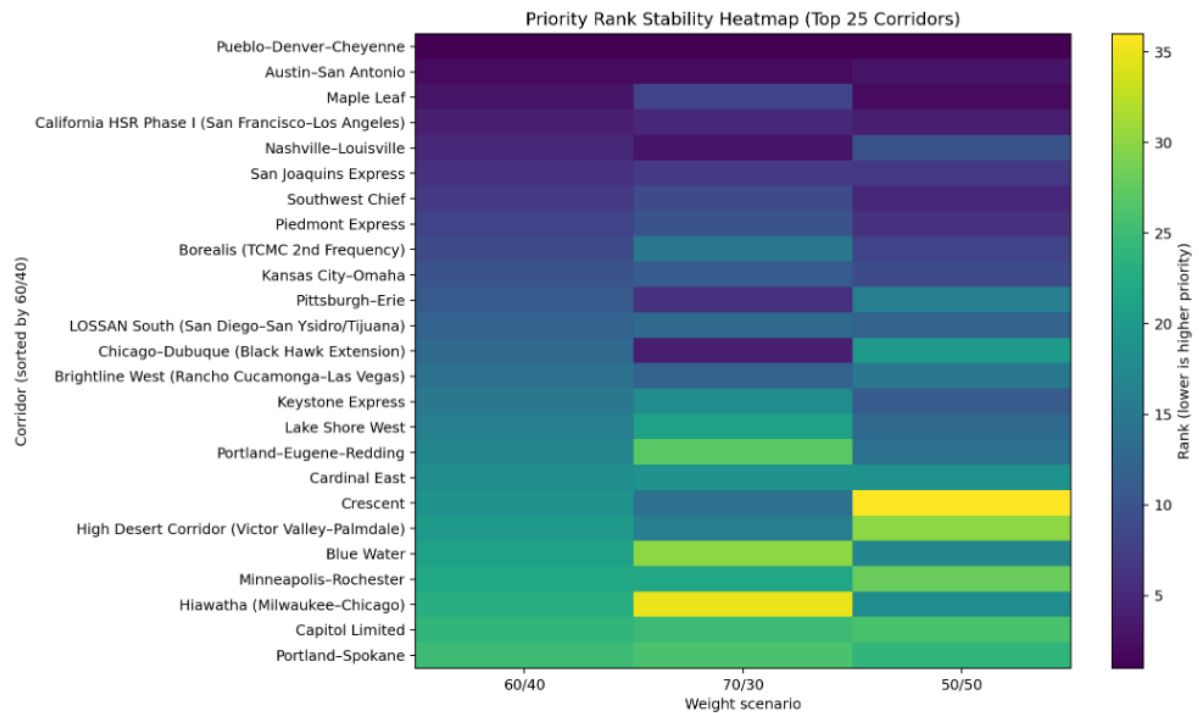
This analysis provides transparency into how the priority score is constructed and helps explain why certain corridors rise to the top of the rankings.

Driver correlations with PriorityScore (60/40) (higher magnitude = stronger driver):

| Driver | Correlation_with_Priority |
|------------------------------|---------------------------|
| BurdenScore | 0.775 |
| ReliabilityIndex | -0.596 |
| DelayCost_per_train_hr_USD | 0.511 |
| DailyRoundTrips | 0.477 |
| EnergyCost_per_mile_USD | -0.110 |
| SubsidyPerRider_USD | 0.106 |
| CrewCost_per_mile_USD | -0.081 |
| OnTimePerformance_pct | -0.060 |
| CostPerTrainMile_USD | 0.051 |
| MaintenanceCost_per_mile_USD | -0.035 |
| AssetAge_years | 0.018 |



4.6 Priority Rank Stability and Sensitivity Analysis

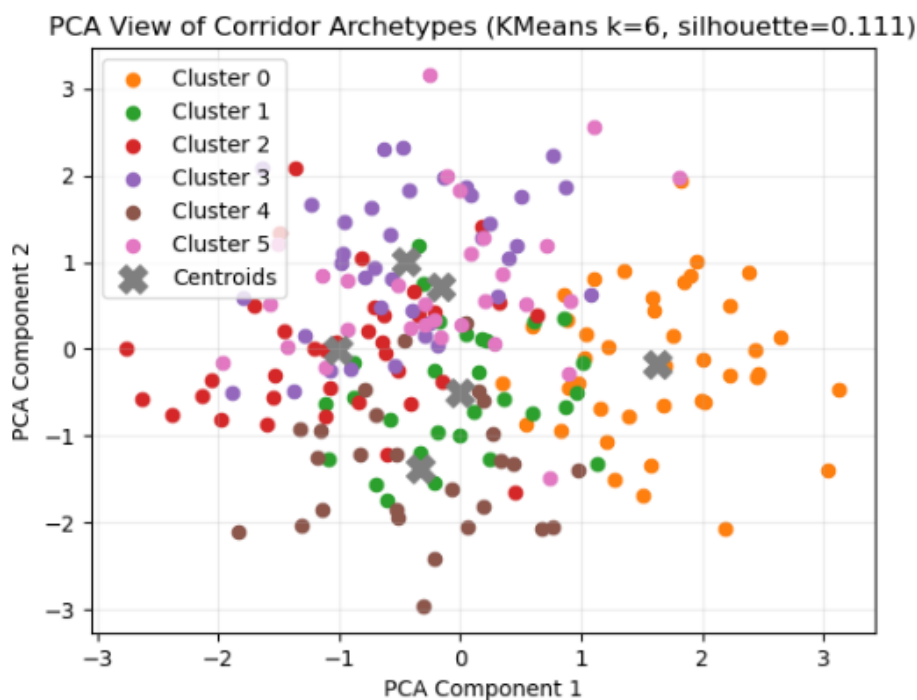


To test robustness, priority rankings were recalculated under alternative weighting scenarios, including a burden-heavy case (70 percent burden, 30 percent reliability) and a balanced case (50 percent burden, 50 percent reliability). Figure 5.6 presents a heatmap showing rank positions of the top 25 corridors across these scenarios.

The heatmap reveals strong rank stability for the highest-priority corridors. Corridors with extreme burden values remain near the top regardless of weighting assumptions. Some mid-ranked corridors shift positions as reliability weight increases, but these shifts are gradual rather than abrupt.

This stability is important because it demonstrates that the prioritization outcomes are not artifacts of a single arbitrary weighting choice. Instead, the highest-priority corridors represent structurally significant system risks that persist across reasonable policy preferences.

4.7 Corridor Archetypes and Clustering Analysis (Figure 4.7, Figure 4.8)



Silhouette scores tested (higher is better):

- k=2: 0.107
- k=3: 0.104
- k=4: 0.106
- k=5: 0.106
- k=6: 0.111

K-means clustering was applied to identify corridor archetypes based on normalized operational and reliability variables. A six-cluster solution provided the best balance between interpretability and separation. Although silhouette scores were modest, this is expected given the continuous nature of corridor characteristics.

Figure 4.7 presents a PCA visualization of the clustered corridors. The plot shows overlapping groups rather than sharply separated clusters, indicating that corridors exist along a continuum of operational profiles rather than discrete categories. Figure 4.8 shows the distribution of corridors across clusters, confirming that no single archetype dominates the system.

These results reinforce the need for differentiated corridor strategies. High-burden, high-frequency corridors face different challenges than lower-burden but reliability-sensitive corridors, and a single policy approach would be ineffective across all archetypes.

4.8 Validation Relationships and Model Behavior (Figure 4.9)

===== PRIORITY SCORE SENSITIVITY (TOP 15 COMPARISON) =====

| | Corridor_Names | Corridor | Base (60% Burden / 40% Reliability) | Burden-Heavy (70/30) | Balanced (50/50) |
|-----|--|--------------|--|--|--|
| 181 | Pueblo–Denver–Cheyenne | Corridor_182 | Pueblo–Denver–Cheyenne | Pueblo–Denver–Cheyenne | Pueblo–Denver–Cheyenne |
| 195 | Austin–San Antonio | Corridor_196 | Austin–San Antonio | Austin–San Antonio | Maple Leaf |
| 27 | Maple Leaf | Corridor_28 | Maple Leaf | Nashville–Louisville | Austin–San Antonio |
| 102 | California HSR Phase I (San Francisco–Los Angeles) | Corridor_103 | California HSR Phase I (San Francisco–Los Angeles) | Chicago–Dubuque (Black Hawk Extension) | California HSR Phase I (San Francisco–Los Angeles) |
| 139 | Nashville–Louisville | Corridor_140 | Nashville–Louisville | California HSR Phase I (San Francisco–Los Angeles) | Southwest Chief |
| 92 | San Joaquins Express | Corridor_93 | San Joaquins Express | Pittsburgh–Erie | Piedmont Express |
| 47 | Southwest Chief | Corridor_48 | Southwest Chief | San Joaquins Express | San Joaquins Express |
| 73 | Piedmont Express | Corridor_74 | Piedmont Express | Maple Leaf | Borealis (TCMC 2nd Frequency) |
| 45 | Borealis (TCMC 2nd Frequency) | Corridor_46 | Borealis (TCMC 2nd Frequency) | Southwest Chief | Kansas City–Omaha |
| 176 | Kansas City–Omaha | Corridor_177 | Kansas City–Omaha | Piedmont Express | Nashville–Louisville |
| 189 | Pittsburgh–Erie | Corridor_190 | Pittsburgh–Erie | Kansas City–Omaha | Keystone Express |
| 133 | LOSSAN South (San Diego–San Ysidro/Tijuana) | Corridor_134 | LOSSAN South (San Diego–San Ysidro/Tijuana) | Brightline West (Rancho Cucamonga–Las Vegas) | LOSSAN South (San Diego–San Ysidro/Tijuana) |
| 152 | Chicago–Dubuque (Black Hawk Extension) | Corridor_153 | Chicago–Dubuque (Black Hawk Extension) | LOSSAN South (San Diego–San Ysidro/Tijuana) | Lake Shore West |
| 100 | Brightline West (Rancho Cucamonga–Las Vegas) | Corridor_101 | Brightline West (Rancho Cucamonga–Las Vegas) | Crescent | Portland–Eugene–Redding |
| 60 | Keystone Express | Corridor_61 | Keystone Express | Borealis (TCMC 2nd Frequency) | Brightline West (Rancho Cucamonga–Las Vegas) |

===== VALIDATION CORRELATION MATRIX =====

| | BurdenScore | ReliabilityIndex | Base (60% Burden / 40% Reliability) |
|-------------------------------------|-------------|------------------|-------------------------------------|
| BurdenScore | 1.000 | 0.046 | 0.775 |
| ReliabilityIndex | 0.046 | 1.000 | -0.596 |
| Base (60% Burden / 40% Reliability) | 0.775 | -0.596 | 1.000 |

Final validation plots comparing priority score against operational burden and reliability confirm expected theoretical behavior. Priority increases monotonically with burden and decreases as reliability improves. These relationships are both visually apparent and statistically supported by correlation analysis.

This final validation step confirms that the analytical framework behaves consistently with operational intuition and established literature, providing confidence that the model is suitable for expansion to larger datasets and more advanced analysis.

5. Future Expansion to 500+ Corridors

The current 200-corridor dataset provides a strong baseline for validating the analytical framework, but the next stage of this research will expand the model to more than 500 corridors to improve scale, coverage, and decision accuracy. A larger corridor base will introduce broader variation in geography, service intensity, operating conditions, ridership patterns, and cost structures. This expansion is important because many of the relationships identified in the current analysis, including burden concentration, instability exposure, and delay propagation sensitivity, are likely to become more robust when tested across a larger and more diverse system.

The expanded phase will also allow the framework to move from structured screening and prioritization into stronger predictive and comparative modeling. With 500+ corridors, the

study can support more stable clustering, deeper correlation testing, stronger validation of propagation patterns, and improved calibration of intervention thresholds. Future versions of the model can also incorporate additional operational and socioeconomic variables, as well as updated federal performance inputs and benefit-cost guidance, to improve both strategic planning and investment relevance. FRA guidance emphasizes structured benefit-cost treatment for rail projects, and FRA/Amtrak performance materials also reinforce the importance of reliability, OTP, and delay sources in corridor evaluation.

6. Discussion

The results of this study show that corridor-level inefficiencies are not evenly distributed across the system. Instead, operational burden, instability, and intervention urgency are concentrated in a relatively small subset of corridors. This is an important finding because it means system improvement should not be approached through uniform treatment. The analysis repeatedly shows that high-burden corridors are also more likely to appear in propagation rankings, action-priority outputs, and budget optimization selections. In practical terms, this means the same corridors that create the greatest stress in the baseline system are also the ones most likely to influence broader network behavior.

Another important insight is that single-metric evaluation is not enough. Corridors cannot be judged using ridership, delay cost, OTP, or reliability alone. Some corridors show high operational burden with only moderate instability, while others show lower burden but greater sensitivity to reliability failure or propagation effects. This confirms the need for a multi-layer framework that evaluates burden, stability, dependency, and readiness together. The guardrail logic, action categorization, and budget analysis all support this point by showing that corridor decisions become more meaningful when multiple conditions are evaluated simultaneously rather than independently.

The scenario and budget findings further strengthen the practical value of the model. Under stress conditions, total burden and instability increase in a predictable way, but the increase is not uniform across corridors. Likewise, budget optimization shows that early investment in the highest-impact corridors produces stronger system gains, while marginal benefit declines as resources are spread more broadly. This supports a targeted engineering approach in which intervention sequencing matters as much as intervention type. Overall, the framework developed in this paper provides a stronger basis for corridor-level decision making than simple ranking methods because it connects operational performance with system-level consequences.

7. Conclusion

The results of this study show that corridor-level inefficiencies are not evenly distributed across the system. Instead, operational burden, instability, and intervention urgency are concentrated in a relatively small subset of corridors. This is an important finding because it means system improvement should not be approached through uniform treatment. The analysis repeatedly shows that high-burden corridors are also more likely to appear in propagation rankings, action-priority outputs, and budget optimization selections. In practical terms, this means the same corridors that create the greatest stress in the baseline system are also the ones most likely to influence broader network behavior.

Another important insight is that single-metric evaluation is not enough. Corridors cannot be judged using ridership, delay cost, OTP, or reliability alone. Some corridors show high operational burden with only moderate instability, while others show lower burden but greater sensitivity to reliability failure or propagation effects. This confirms the need for a multi-layer framework that evaluates burden, stability, dependency, and readiness together. The guardrail

logic, action categorization, and budget analysis all support this point by showing that corridor decisions become more meaningful when multiple conditions are evaluated simultaneously rather than independently.

The scenario and budget findings further strengthen the practical value of the model. Under stress conditions, total burden and instability increase in a predictable way, but the increase is not uniform across corridors. Likewise, budget optimization shows that early investment in the highest-impact corridors produces stronger system gains, while marginal benefit declines as resources are spread more broadly. This supports a targeted engineering approach in which intervention sequencing matters as much as intervention type. Overall, the framework developed in this paper provides a stronger basis for corridor-level decision making than simple ranking methods because it connects operational performance with system-level consequences.