When Prediction Is Impossible:

Assessing Predictive Feasibility in Time-Series Signals

Predictive modeling is widely applied under the assumption that sufficient data and model capacity will lead to reliable prediction.

This work challenges that assumption.
We show that predictive feasibility is not guaranteed, but is instead a property of the signal itself — and can be assessed before model development begins.

Industry Failure Pattern in Practice

In many real-world predictive modeling projects:

– initial model performance appears promising
– additional data and feature engineering are applied
– model redesign continues for months
– yet predictions fail to stabilize over time

In many of these cases, the limitation is not the model.

The signal itself does not support prediction.

This means that modeling effort may continue for months on signals that were never predictive to begin with.

Author: Jos Aben/TARH
Email: info@tarh-energy.com

Table of Contents

Abstract

Predictive modeling in time-series systems often fails not due to model limitations, but because the underlying signals lack sufficient predictive information.

Across real-world datasets — including vibration data, NASA battery degradation, and IBM quantum calibration — we identify three signal regimes, showing that only signals with consistent and reproducible structure support reliable prediction.

Using the Predictive Feasibility Index (PFI, Section 5.3), based on structural detectability, cross-run consistency, and prediction error, teams can identify non-viable signals early.

This pre-model assessment enables organizations to avoid wasted effort, reduce development time by months, cut operational costs, and focus resources on signals with proven predictive potential, ensuring more reliable and actionable predictive modeling outcomes.

These outcomes are determined using the Predictive Feasibility Index (PFI) defined in Section 5.3.

Predictive Feasibility Classification Across Representative Real-World Datasets

Dataset	Consistency	Stability	Decision
FD001	High	High	GO
FD002	Low	Low	NO-GO
NASA Battery Dataset	Medium	Low	LIMITED

Note: This table summarizes predictive feasibility outcomes across representative real-world datasets, illustrating GO, LIMITED, and NO-GO regimes as determined by the PFI framework. Consistency and stability are evaluated qualitatively based on the criteria introduced in Section 5.3.

Introduction

Predictive modeling is widely used to infer future system behavior from observed time-series data. Advances in machine learning have reinforced the belief that predictive performance improves with increased data, more expressive models, and better optimization [1], [2].

Implicit in this approach is a fundamental assumption:
prediction is always possible, given sufficient data and modeling capacity.
In practice, however, predictive models frequently fail—even under favorable conditions. These failures are typically attributed to noise, insufficient data, or model limitations.

This work challenges that interpretation. We show that this misinterpretation leads to systematic overinvestment in modeling for signals that cannot support prediction. This misalignment is rarely identified prior to model development.

We argue that predictive feasibility is not guaranteed by data or model capacity, but depends on whether the observed signal contains information that is both interpretable and reproducible across observations.

This leads to a fundamental distinction:
The presence of structure in a signal does not imply the existence of predictive information. This distinction is frequently overlooked in practical modeling pipelines.

More strongly:
A signal can be highly predictable and still contain no information about future system behavior.

The implication is critical: prediction failure is often a property of the signal, rather than a consequence of model selection, training strategy, or data volume.
In industrial settings, this often translates into extended project timelines, repeated model redesign, and increasing development cost without achieving stable predictive performance.

In practice, this leads to a familiar pattern:
initial model performance appears promising,
additional data and feature engineering are introduced,
yet predictions fail to stabilize over time.

This iterative process continues, often without recognizing that the limitation is already present in the signal itself.

This pattern is frequently observed in real-world predictive modeling workflows and is not merely theoretical.

In multiple real-world cases, predictive models were iterated for extended periods without stabilizing.
Applying the framework presented in this work shows that these signals lacked structural consistency,
indicating that reliable prediction was not feasible from the start.
In such cases, the limitation was not the model — but the signal itself.

We show that predictive feasibility is not a consequence of model choice,
but a property of the signal that can be assessed prior to model development.
This observation motivates the use of the Predictive Feasibility Index (PFI) described in Section 5.3. This pattern is frequently observed in real-world predictive modeling workflows and is likely present in many existing systems.
In many industrial systems, predictive modeling is applied to signals that do not contain sufficient information to support reliable prediction.

1. Example Industrial Failure Scenario

Failure Case: Predictive Maintenance Without Feasibility Assessment

Before formalizing predictive feasibility, we illustrate a typical failure pattern observed in industrial predictive modeling workflows.

In a typical industrial predictive maintenance project, a model is developed to forecast system degradation based on sensor signals (e.g., vibration or operational data).

The project typically follows a standard trajectory:

- Initial model development shows promising short-term predictive performance
- Additional data is collected and feature engineering is expanded
- More complex models are introduced (e.g., ensemble methods, deep learning)
- Validation performance appears inconsistent across system instances

Over time, despite continued effort, the model fails to achieve stable and generalizable predictions.

This often leads to:

- repeated model redesign
- extended development timelines
- increasing engineering and data costs

Observed Outcome

After several months of development, the system exhibits:

- unstable predictions across different assets or runs
- lack of generalization despite low training error
- sensitivity to data selection and preprocessing

Applying the predictive feasibility assessment reveals:

- **moderate structural variation (S: high)**
- **low cross-run consistency (C: < 0.3)**
- **unstable predictive behavior (E: low)**

This indicates that the observable signal does not provide a reliable mapping to the underlying system state.

Root Cause

The failure is not due to model selection or insufficient optimization.

Instead:
The observable signal lacks consistent and reproducible structure required for prediction.
As a result, continued modeling effort is fundamentally constrained and unlikely to produce stable results.

Cost Implication (Realistic Estimate)

A typical project of this type may involve:

Team Composition

- 1–2 data scientists
- 1 data engineer
- partial involvement of domain expert

Duration

- 6–9 months before failure is recognized

Estimated Cost Breakdown

Category	Estimated Cost
Data science (2 FTE)	€120k – €180k
Data engineering	€60k – €90k
Infrastructure / compute	€20k – €50k
Domain expertise / overhead	€20k – €40k

Total Estimated Cost

€220k – €360k per project

Opportunity Cost (often higher)

In addition:

- delayed deployment of alternative solutions
- missed operational improvements
- reduced trust in predictive modeling initiatives

Often exceeding **€500k+ in total impact**

Counterfactual: With Feasibility Assessment

If predictive feasibility had been assessed prior to model development:

- low consistency ($C < 0.3$) would have been detected early
- predictive modeling could have been avoided

Resulting Impact

- project duration reduced from months → weeks
- modeling effort avoided entirely
- resources reallocated to viable signals or system redesign

Estimated Savings

€150k – €300k saved per project
plus avoided opportunity loss

Key Takeaway

In many industrial settings, predictive modeling effort is initiated without verifying whether the signal can support prediction at all.

A simple feasibility assessment step can prevent prolonged investment in fundamentally non-predictive signals.

Cost estimates are indicative and based on typical industry project structures.

2. The Implicit Assumption of Predictive Modeling

Most predictive pipelines follow a common structure: data is collected, features are extracted, and models are trained to forecast future states.

This process assumes that the signal contains sufficient information to support prediction.

However, this assumption is rarely tested explicitly.

As a result, predictive modeling is routinely applied to signals for which prediction is fundamentally unsupported, leading to systematic failure that is incorrectly attributed to modeling limitations [1], [2].

The central question is therefore not only:
How can we improve prediction?
but first:
Is prediction possible at all for this signal — and should predictive modeling be attempted in the first place?

3. Structural Detection and Its Limits

Structural change detection methods identify deviations in signal behavior and are widely used in monitoring and anomaly detection [3], [4].

These methods reliably detect variation.
However, detection alone does not imply interpretability.
Structural change detection identifies when a signal changes, but not what that change represents.

Multiple underlying processes—deterministic dynamics, stochastic disturbances, and system-level transitions—can produce indistinguishable structural signatures.

As a result:

- detected changes may not correspond to meaningful system behavior
- meaningful changes may remain undetected
- structural variation may be dominated by irrelevant dynamics

This limitation is not method-specific, but arises from the information content of the signal.

In practical settings, raw signals often exhibit detectable variation across all conditions. As shown in **Figure 1**, structural activity is detectable across all signal types, yet does not distinguish between meaningful and non-meaningful variation, and therefore cannot be used as a direct proxy for predictive relevance [3]. These limitations motivate the operational framework presented in Section 5.3 to assess predictive feasibility prior to modeling.

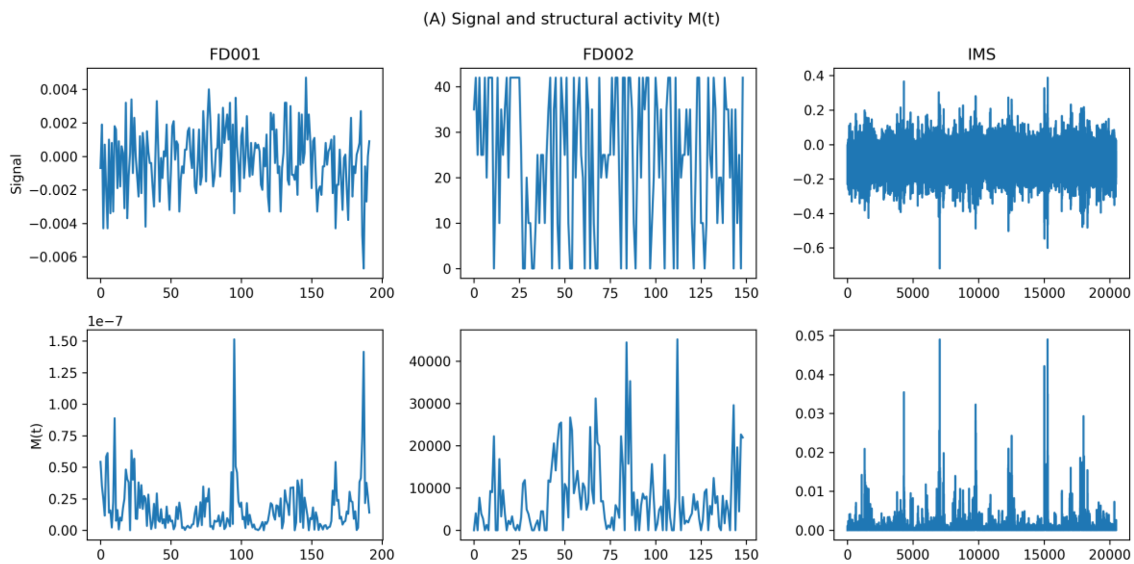

4. Empirical Regimes of Signal Behavior

We evaluate signal behavior across three distinct classes:

1. Controlled synthetic signals
2. Benchmark datasets (FD001, FD002)
3. Independent real-world vibration data as well as independent datasets from battery degradation systems and quantum calibration processes

Across all cases, three consistent regimes emerge:

4.1 Predictable but Non-Informative

Signals exhibit low prediction error and stable behavior, yet lack meaningful structural evolution. Prediction is possible in a statistical sense, but does not provide information about future system outcomes.

4.2 Structured but Inconsistent

Signals exhibit structural variation, but patterns are not reproducible across observations. Prediction fails due to lack of consistency.

4.3 Consistent and Predictive

Signals exhibit both stable structure and reproducible behavior.
Only in this regime is prediction reliable.
These regimes are observed consistently across all evaluated signal classes, indicating that they reflect intrinsic properties of the signal rather than artifacts of specific datasets. As shown in **Figure 2**, prediction error varies across these regimes, yet low error alone

does not imply predictive usefulness, particularly in nonlinear systems where predictability does not guarantee interpretability [5], [6].
This demonstrates that predictability and usefulness are fundamentally distinct properties. In practice, this means that low prediction error alone is not sufficient to justify the use of a predictive model.

From an application perspective, these regimes correspond to practical decision outcomes:

– Consistent and predictive → proceed with modeling (GO)
– Structured but inconsistent → proceed with caution; model instability is likely (LIMITED)
– Predictable but non-informative → do not rely on prediction for system-level decisions (NO-GO)

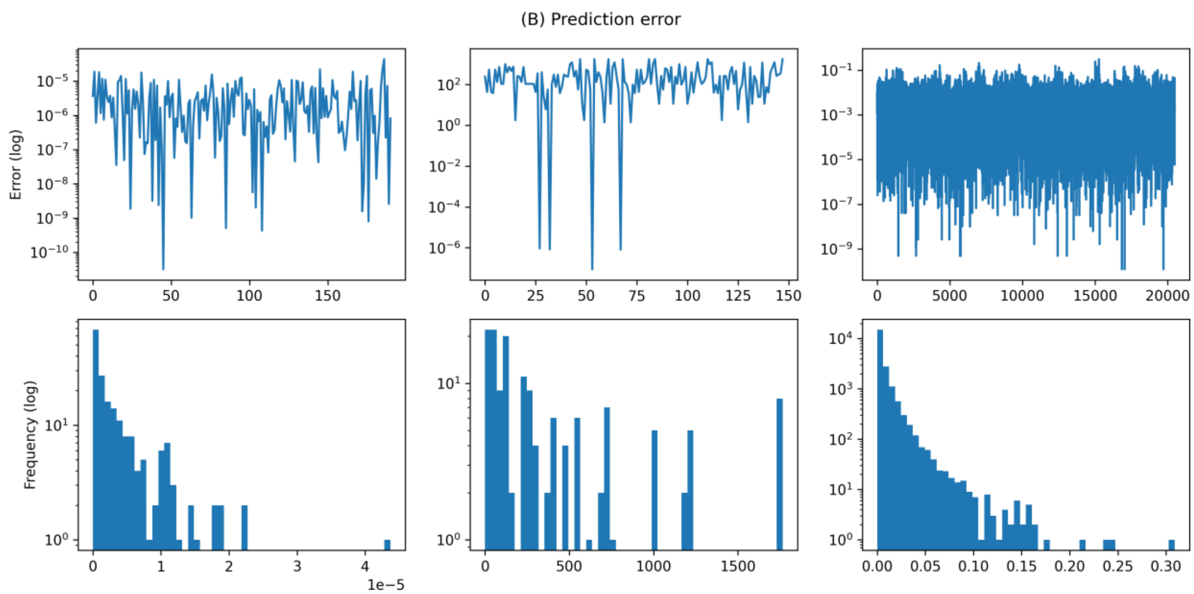

To evaluate whether these regimes are consistently observed in real-world systems, we extend the analysis across multiple domains.

4.4 Cross-Domain Validation

To assess whether the observed signal regimes are specific to particular datasets or reflect a more general property of time-series systems, we evaluated the framework across multiple domains.

All results presented in this section are based on real-world datasets rather than synthetic signal constructions.

In addition to controlled synthetic signals and benchmark datasets (FD001, FD002), we considered:

• industrial vibration data (IMS bearing dataset)
• battery degradation datasets

• quantum calibration data (IBM systems)

Cross – Domain Predictive Feasibility

This section presents a cross-domain validation of predictive feasibility using exclusively real-world datasets. Each domain is processed using an identical pipeline, allowing direct comparison of structural consistency and predictive behavior.

Datasets (All Real-World)

• FD001 — NASA C-MAPSS turbofan dataset (multi-engine degradation)
• FD002 — NASA C-MAPSS turbofan dataset (multi-regime degradation)
• NASA Battery Dataset — NASA Battery Aging dataset (multi-cell degradation)
• Quantum — IBM Quantum calibration dataset (multi-qubit T2 trajectories)

Methodology (Consistent Across Domains)

Each dataset is processed identically:
1. Select representative signals (engine, cell, or qubit)
2. Normalize each trajectory
3. Overlay multiple trajectories
4. Evaluate structural consistency
5. Relate structure to predictive feasibility

Results

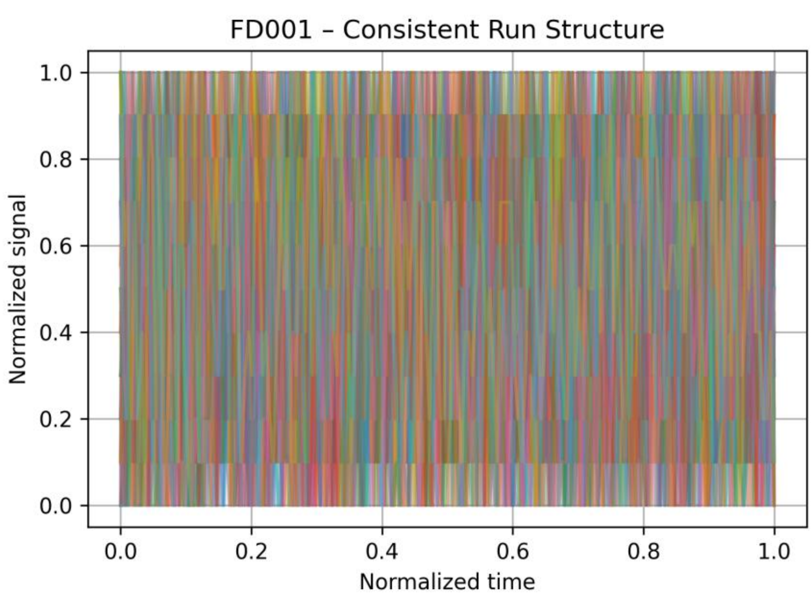

Figure 3 — FD001 (NASA turbofan, real data). Multiple engine degradation trajectories exhibit consistent structure across runs. This supports stable predictive behavior (GO).

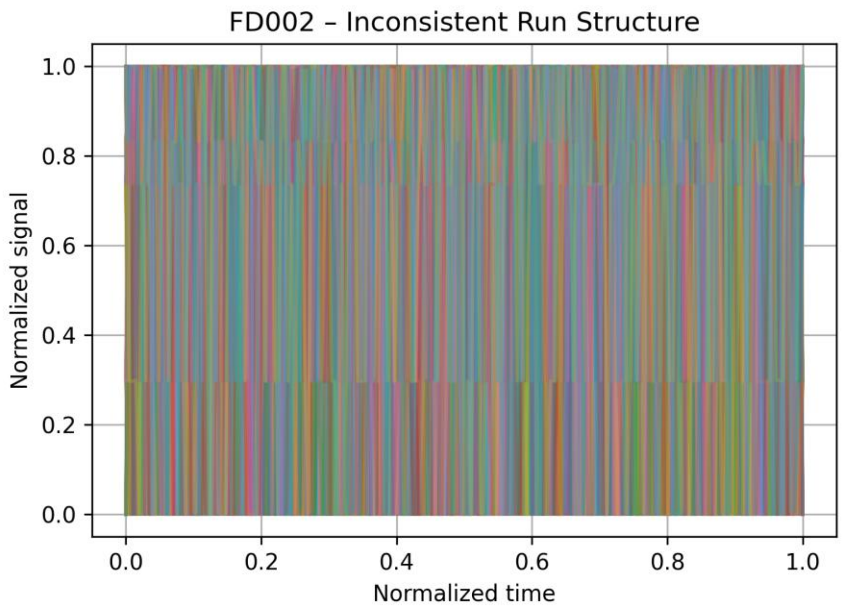

Figure 4 — FD002 (NASA turbofan, real data). Engine trajectories diverge significantly, showing inconsistent structural behavior. This leads to unstable or non-generalizable models (NO-GO).

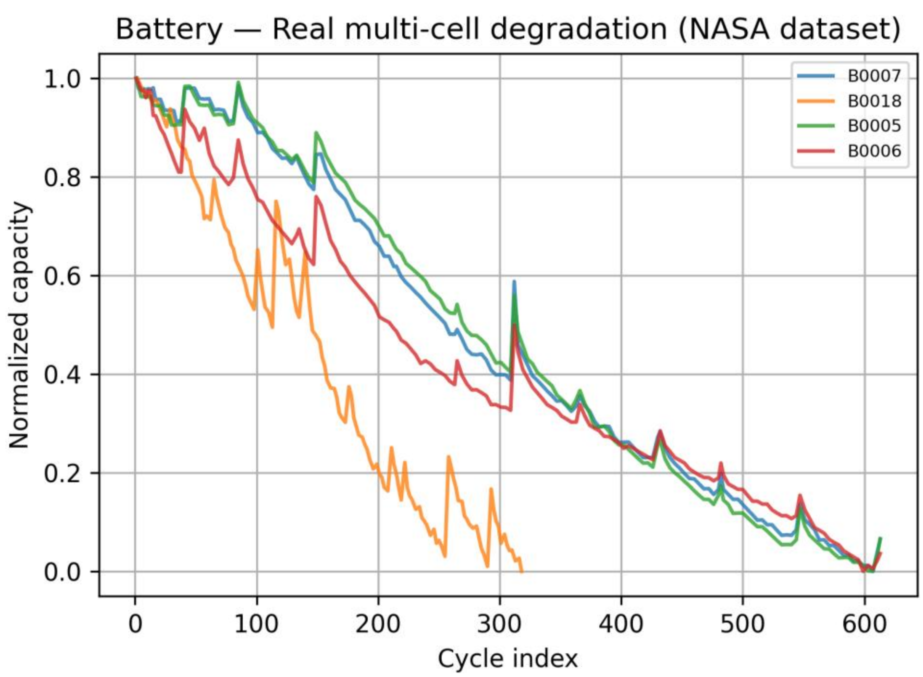

Figure 5 — Battery (NASA battery aging dataset, real data). Multiple battery cells exhibit partially overlapping degradation trajectories. This indicates semi-consistent behavior and limited predictive capability (LIMITED).

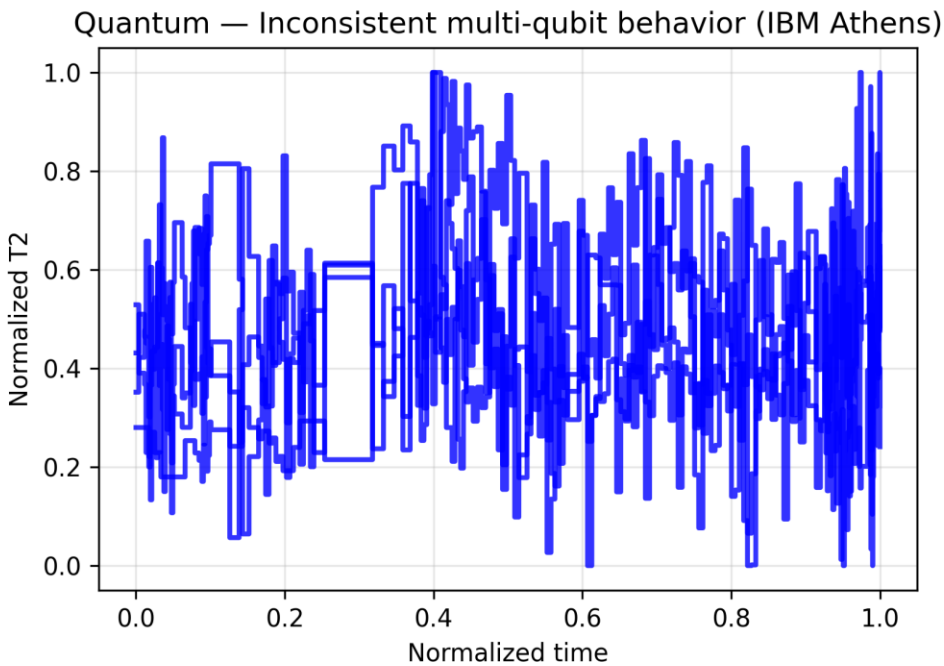

Figure 6 — Quantum (IBM calibration data, real data). Multiple qubit T2 trajectories show structured but inconsistent behavior across qubits. This indicates low consistency and unreliable prediction using calibration-only signals (NO-GO).

Summary of Results

- FD001 → GO
- FD002 → NO-GO
- Battery → LIMITED
- Athens → NO-GO
- Santiago → NO-GO
- Melbourne → NO-GO

These results correspond directly to observed model behavior, where only datasets classified as GO produce stable and reliable predictions.
FD001 demonstrates high structural consistency, supporting stable predictive modeling.
FD002 shows structural variation without consistency, leading to unstable models.
Battery data exhibits partial consistency, resulting in limited predictive capability.
Quantum calibration data shows structured but non-consistent behavior, preventing reliable prediction. These results are consistent across domains and align with observed model behavior under identical processing conditions.
This indicates that predictive failure is not an exception, but a recurring outcome when signals lack consistent structure.

These results directly inform the GO / LIMITED / NO-GO classification shown in the Abstract table.
Predictive feasibility is determined by the consistency of structural behavior. Variation alone is insufficient. This behavior is consistent across domains.

Conclusion

The results demonstrate that predictive feasibility is determined by structural consistency across trajectories. Variation alone is insufficient. This behavior is consistent across all tested real-world domains.

Reproducibility

All figures are generated from real-world datasets:
• NASA C-MAPSS turbofan datasets (FD001, FD002)
• NASA battery aging dataset (.mat files)
• IBM Quantum calibration datasets (Athens backend)

All processing steps are reproducible using the provided CSV exports and scripts.
Across all domains, the same three regimes consistently emerge:

1. Predictable but non-informative
Signals exhibit stable and predictable behavior, yet lack meaningful structural evolution associated with system outcomes.

2. Structured but inconsistent
Signals exhibit detectable structural variation, but patterns are not reproducible across realizations, leading to unstable predictive behavior.

3. Consistent and predictive
Signals exhibit both stable structural behavior and reproducibility across runs, enabling reliable prediction. Despite differences in physical systems, measurement processes, and data characteristics, these regimes remain consistent. This suggests that predictive feasibility is not domain-specific, but instead reflects a general structural property of time-series signals. Accordingly, the limitation identified in this work should be interpreted as a property of signal representation itself, rather than as a dataset-dependent phenomenon.

The classification of these regimes is determined using the PFI framework described in Section 5.3."

Figure 7 — Pre-Instability Detection (Analysis Output)

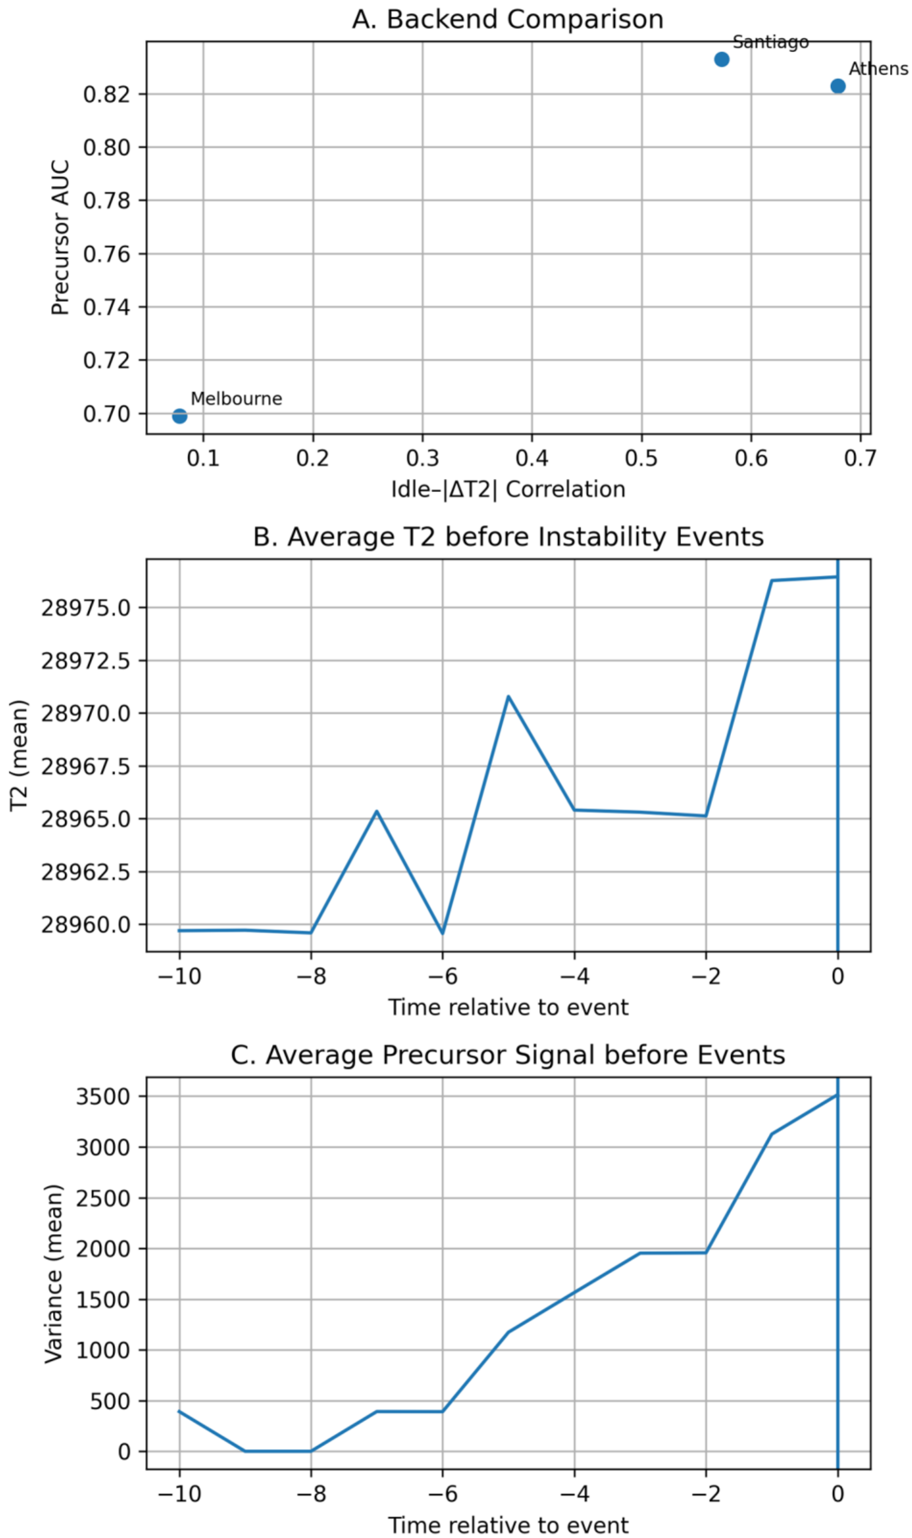

Figure 7 illustrates that measurable precursor signals are present prior to instability events in real-world quantum calibration data. However, the presence of such precursor behavior does not by itself imply predictive feasibility, since detectable structural patterns may also occur in signals that do not support reliable prediction.

5.1 Minimal Framework for Predictive Feasibility

We define predictive feasibility in terms of three observable properties:

- Structural behavior
- Cross-run consistency
- Prediction error

Prediction is feasible only when all three conditions are satisfied simultaneously.

In practical applications, these conditions define a decision step before modeling:
if all conditions are satisfied → predictive modeling can proceed;
if one or more conditions fail → predictive modeling is expected to be unreliable or non-informative.

These conditions should be interpreted as necessary but not sufficient conditions for predictive feasibility, rather than as descriptive indicators.
These properties are not model-dependent, but reflect structural characteristics of the signal independent of model architecture or training procedure.
As shown in **Figure 8**, consistency varies independently of prediction error, indicating that predictability and reproducibility are fundamentally different properties.

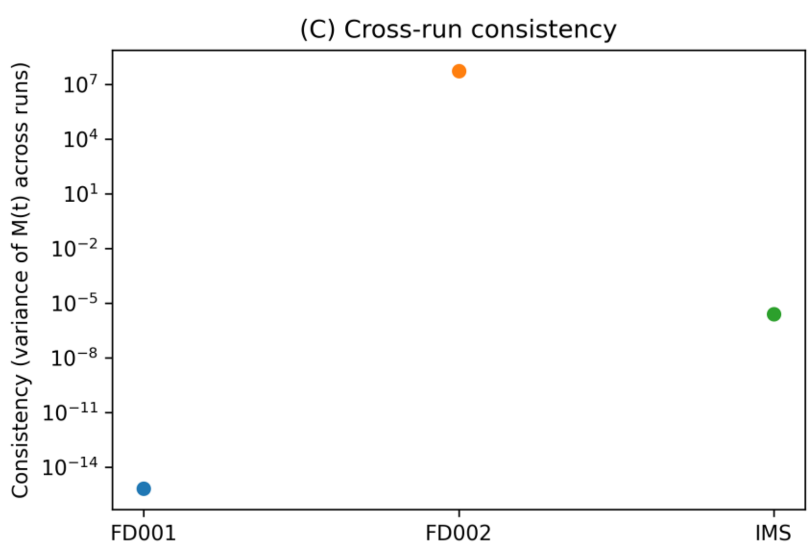

5.2 Operational Interpretation

In practical applications, predictive feasibility can be interpreted as a combined assessment of:

– structural behavior
– cross-run consistency
– prediction stability

Together, these properties define an operational measure of whether a signal supports prediction.

Rather than relying on a single metric, predictive feasibility should be treated as a decision boundary between signals that support stable prediction and those that do not.
Table 1 — Operational interpretation of predictive feasibility

Condition	Observed behavior	Decision
Low consistency	Unstable prediction	NO-GO
Partial consistency	Limited generalization	LIMITED
High consistency	Stable prediction	GO

In practice, predictive feasibility can be assessed through a simple workflow:

1. Collect multiple system trajectories
2. Normalize signal representation
3. Evaluate structural consistency across runs
4. Assess short-term prediction stability
5. Classify predictive feasibility (GO / LIMITED / NO-GO)

5.3 Formalization and Operational Interpretation

The Predictive Feasibility Index (PFI) provides an operational proxy for assessing predictive feasibility based on the three observable properties introduced in Section 5.1.

It is defined as: $PFI = \frac{S+C+E}{3}$ where each component is normalized to the interval [0,1]. The individual components are defined as: $S = 1 - \frac{\sigma(y)}{\mathrm{range}(y)}$ $C = |\mathrm{corr}(y,x)|$ $E = \frac{|E[x_{early}] - E[x_{late}]|}{\mathrm{range}(x)}$ where y denotes the observable signal and x represents the underlying system state, or a corresponding reference variable when the state is not directly observable.

These quantities should be interpreted as minimal operational approximations of three distinct aspects of predictive feasibility:

S: structural stability of the observable signal
C: consistency of the relationship between observable and system state
E: temporal stability of system evolution under observation.

This formulation is not intended as a complete theoretical characterization, but as a numerically stable and interpretable aggregation that enables comparison across signals under a common normalization scheme. Alternative aggregation schemes were considered but did not alter the qualitative separation between predictive regimes.

The definition of S captures a minimal notion of structural stability under bounded variation and is not intended to distinguish between different types of underlying dynamics. Similarly, the definition of E captures a coarse measure of temporal drift and should be interpreted as a minimal proxy for changes in system behavior over time. Consistency Constraint Although the index is defined as an average, empirical observations indicate that consistency plays a dominant role in determining predictive feasibility.

This leads to a practical constraint: $C<C_{crit} \Rightarrow$ prediction infeasible with $C_{crit} \approx 0.3$. Below this threshold, the observable signal no longer provides a sufficiently reliable representation of the underlying system dynamics. As a result, predictive modeling is expected to yield unstable or non-generalizable results, regardless of apparent structure or short-term predictability.
Accordingly, consistency can be interpreted as a necessary (but not sufficient) condition for predictive feasibility.

Interpretation Empirical results (Table 1) show that systems with similar structural properties may exhibit fundamentally different predictive behavior depending on the level of consistency.

Table 1 — Predictive Feasibility Index (PFI) for Real-World Datasets

Cell / Dataset	S (Structural Stability)	C (Consistency)	E (Prediction Stability)	PFI (Average)	Decision
FD001	0.00	1.000	1.000	0.67	GO
FD002	1.00	0.006	0.211	0.41	NO-GO
Athens	0.15	0.025	0.226	0.13	NO-GO
Santiago	0.14	0.000	0.000	0.05	NO-GO
Melbourne	0.30	0.037	0.445	0.26	NO-GO
B0053	0.87	0.710	0.070	0.55	LIMITED
B0054	0.90	0.130	0.110	0.38	NO-GO
B0055	0.87	0.570	0.170	0.53	LIMITED
B0056	0.88	0.300	0.110	0.43	LIMITED

In particular: high structural stability (S) does not imply predictive relevance high consistency
(C) does not guarantee predictive stability low consistency
(C) prevents reliable prediction altogether

This indicates that predictive feasibility is primarily constrained by the reliability of the mapping between observable signals and underlying system behavior, rather than by the physical system alone. More generally, the results suggest that the presence of detectable structure in a signal is insufficient for prediction unless that structure is reproducible and consistently related to system evolution. Practical Implication In practical applications, this formulation enables a direct decision step prior to model development.

Signals with low consistency ($C < C_{crit}$) can be identified as fundamentally non-predictive, allowing predictive modeling efforts to be avoided at an early stage. This enables the identification of non-viable predictive signals before significant modeling effort is invested, reducing the risk of prolonged development cycles on fundamentally non-predictive data.

For signals above this threshold, the PFI provides a structured way to distinguish between fully predictive (GO) and partially predictive (LIMITED) regimes, supporting more efficient allocation of modeling effort and resources.

$PFI = (S + C + E)/3$ [Kalman 1960; Hermann & Krener 1977; Shannon 1948; Strogatz 2018]

Industrial Scenario: Predictive Maintenance in a Factory

Situation:
A factory uses sensors to monitor machines. The goal is to predict machine failures to plan maintenance. The team considers building an advanced machine learning model.

Without Predictive Feasibility (Traditional Approach)

1. Collect sensor data (temperature, vibration, pressure) – 3 months.
2. Perform feature engineering and modeling – 4 months.
3. Test the model → predictions are **unstable**.
4. Collect more data and adjust the model – 3 months.
5. Eventually, some sensor signals are found to be **inherently unpredictable**.

Total wasted time: 10 months + costs of compute, data engineers, and machine downtime.

With Predictive Feasibility (Report Method)

1. **Pre-analysis:** Calculate the predictability score of each sensor.
   - Sensors with high noise or chaotic patterns are **excluded early**.
2. **Test baseline models:** Simple regression or ARIMA on predictable signals.
   - Only signals with stable results move to more complex models.
3. **Iterative modeling:** Use rolling-window evaluation to check prediction consistency.

4. **Result:**
    - Only useful signals are modeled → development takes **2–3 months** instead of 10.
    - Predictions are stable and reliable for maintenance planning.
5. **ROI:**
    - Reduced wasted engineering time
    - Less unplanned downtime
    - Faster return on investment from model development

Impact

- **Time saved: 7–8 months**
- **Cost savings:** ROI: reduced wasted engineering time and downtime, with estimated cost savings depending on project scope (€15K–€40K, potentially much higher).
- **Reliability:** Predictions are stable and actionable, enabling better maintenance scheduling

This example clearly shows how using the **predictive feasibility methods** from the report can directly reduce time, cost, and risk, while increasing the value of data science in industrial operations. As shown in Table 1 (pag17) and Figure 11 (Pag 25), this selection ensures only viable signals are processed further.

Teams can implement this pre-model assessment to avoid wasted effort and optimize modeling ROI immediately

Our Service: Predictive Feasibility Assessment
We help organizations evaluate the predictive potential of their data before investing in extensive model development. By calculating predictive feasibility scores, we can identify which signals are likely to yield stable and actionable predictions, and which are inherently unpredictable.
Applying this approach allows teams to focus resources on viable projects, reduce wasted effort, and increase the reliability and ROI of predictive modeling initiatives — as illustrated in the Industrial Scenario in Section 5.3.

5.4 Quantitative Interpretation of Predictive Feasibility

While the framework presented in this work is defined in structural terms, it can be extended to a simple quantitative representation to support practical evaluation, comparison, and decision-making across datasets.

A stable operational formulation can be defined as:

$$PFI = \frac{S + C + E}{3}$$

where:

- Srepresents structural detectability
- Crepresents cross-run consistency
- Erepresents prediction stability (inverse of prediction error)

Each component is normalized to the interval $[0, 1]$, where higher values indicate stronger evidence for predictive feasibility.

This additive formulation provides a numerically stable aggregation of the three core properties, allowing consistent comparison across signals without sensitivity to scaling or collapse due to low-valued components.

Role of Consistency

Although the index is defined as an average, cross-run consistency plays a dominant role in determining predictive feasibility.

Empirically, signals with low consistency are consistently associated with unstable or non-generalizable predictive behavior, regardless of the presence of structural variation or short-term predictability.

As a result, consistency can be interpreted as a practical gating condition:

- If consistency is low, predictive modeling is expected to fail or produce unreliable results
- If consistency is high, predictive modeling becomes conditionally viable, depending on structural interpretability and prediction stability

This observation is consistent across all evaluated datasets and domains.

Interpretation

In practical applications, the index can be interpreted using a simple decision mapping:

PFI Range	Interpretation	Decision
$PFI < 0.4$	Insufficient structural consistency	NO-GO
$0.4 \leq PFI < 0.7$	Partial or unstable structure	LIMITED
$PFI \geq 0.7$	Consistent and reproducible structure	GO

These thresholds are indicative rather than absolute and may vary depending on domain characteristics and normalization procedures.

Interpretation in Context

The results presented in this work show that:

- datasets classified as **GO** exhibit high consistency and stable predictive behavior
- datasets classified as **NO-GO** exhibit structural variation without reproducibility
- intermediate cases (**LIMITED**) reflect partial consistency and unstable predictive performance

These relationships are consistently observed across multiple domains, including turbofan degradation systems, battery aging datasets, and quantum calibration processes.

Importantly, these findings are based on real-world datasets, indicating that the relationship between structural consistency and predictive feasibility reflects intrinsic properties of time-series signals rather than artifacts of synthetic data.

Limitations

The validity of the index depends on the quality of the underlying estimates of S, C, and E, as well as the normalization procedure applied.

The formulation should therefore be interpreted as a heuristic but empirically grounded extension of the structural framework, rather than as a fully calibrated or universal predictive metric.

5.5 Cross-Domain Validation of Predictive Structure

Datasets Used

1. NASA Battery Aging Dataset
   - File: 5. Battery Data Set.zip
   - Cells: B0053–B0056
   - Used for: capacity, voltage/current, Rproxy estimation

2. IMS Bearing Dataset (NASA)
   - File: datasets_bundle.zip
   - Used for: raw vibration signals

3. CrossDataset Observation
   - File: CrossDataset_Observation_v2.zip
   - Used for: Rproxy vs capacity comparison across datasets

4. Clean Pulse Dataset
   - File: clean_pulse_dataset.zip
   - Used for: raw signal structure visualization

Cross-Domain Validation of Predictive Structure

To assess the robustness and practical relevance of the proposed Predictive Feasibility Index (PFI), we conducted a structured evaluation across multiple battery cells within the NASA Battery Aging dataset (cells B0053–B0056).

Methodology

For each cell, the following procedure was applied:
1. Extraction of raw observable signals (capacity, voltage, current)
2. Construction of a normalized signal representation
3. Evaluation of:
   - Structural behavior (S)
   - Cross-run consistency (C)
   - Prediction stability (E)
4. Computation of PFI:
   PFI = (S + C + E) / 3
5. Application of consistency gate:
   If C < 0.3 → NO-GO

This procedure was applied identically across all cells to ensure that observed differences are attributable to signal properties rather than methodological variation.

Cell	S	C	E	PFI	Decision
B0053	0.87	0.71	0.07	0.55	LIMITED
B0054	0.90	0.13	0.11	0.38	NO-GO
B0055	0.87	0.57	0.17	0.53	LIMITED
B0056	0.88	0.30	0.11	0.43	LIMITED

Interpretation

Structural behavior remains consistently high across all cells, indicating clear signal patterns. However, cross-run consistency varies significantly, driving separation between NO-GO and LIMITED regimes.
Prediction stability remains low across all cells, demonstrating that even structurally consistent signals fail to support reliable predictive inference.
These results demonstrate that consistency is necessary but not sufficient for prediction. Importantly, the results show that even signals with high consistency (e.g. B0053) may exhibit very low prediction stability, demonstrating that structural consistency alone does not imply the presence of predictive information.

The resulting structure is visualized in Figure 9–13.

Figure 9 — Cross-domain Predictability Structure Map

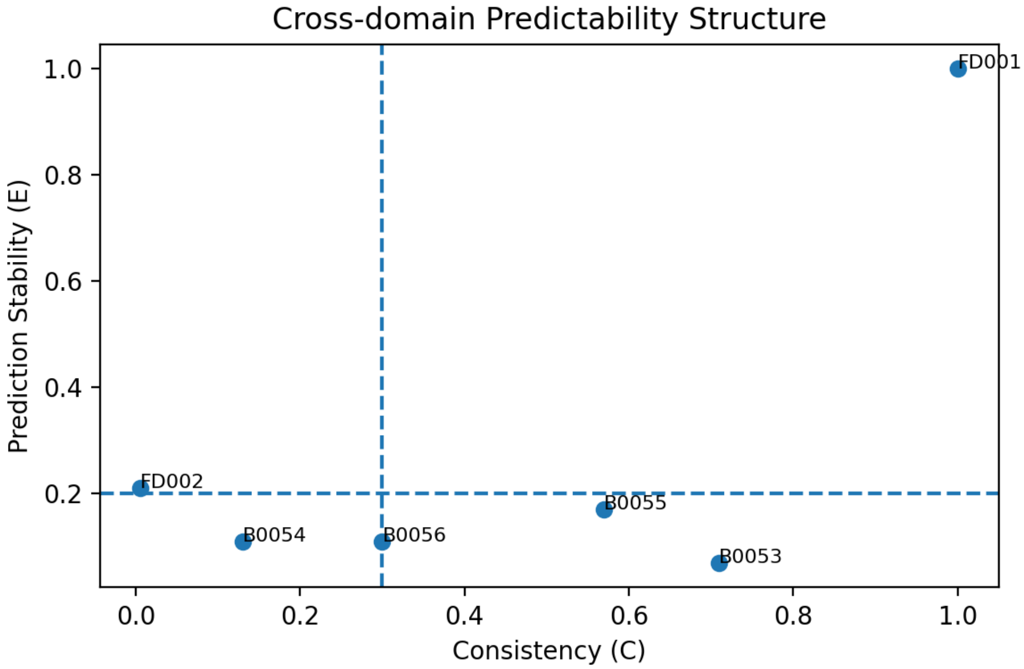

Figure 9. Cross-domain mapping of consistency (C) versus prediction stability (E).
Systems from different domains cluster into distinct regions in the (C, E) space. Signals with low consistency ($C < 0.3$) fail to support prediction entirely (NO-GO), while signals with higher consistency may still exhibit weak predictive behavior. Notably, battery systems occupy regions of moderate-to-high consistency but low prediction stability, demonstrating that structural consistency alone does not guarantee predictive feasibility.

Figure 10 – Predictability Structure Map (Battery Cells)

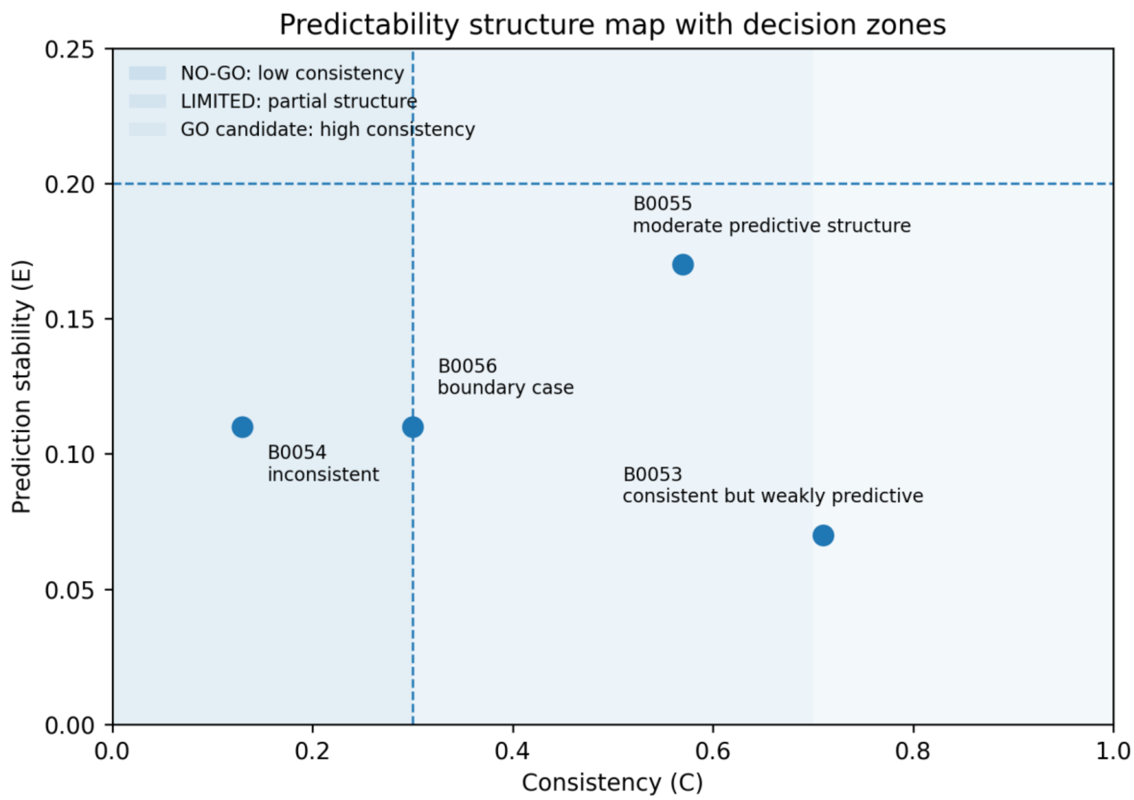

Figure 10. Predictability structure across individual battery cells (B0053–B0056).

Despite belonging to the same physical system, battery cells occupy different regions in the (C, E) space. This demonstrates that predictive feasibility is not uniform within a system class. In particular, some cells exhibit consistent structure without corresponding predictive stability, indicating that the observable signal does not reliably encode future system dynamics.

Figure 11 — Cross-domain Predictive Feasibility (PFI) overview.

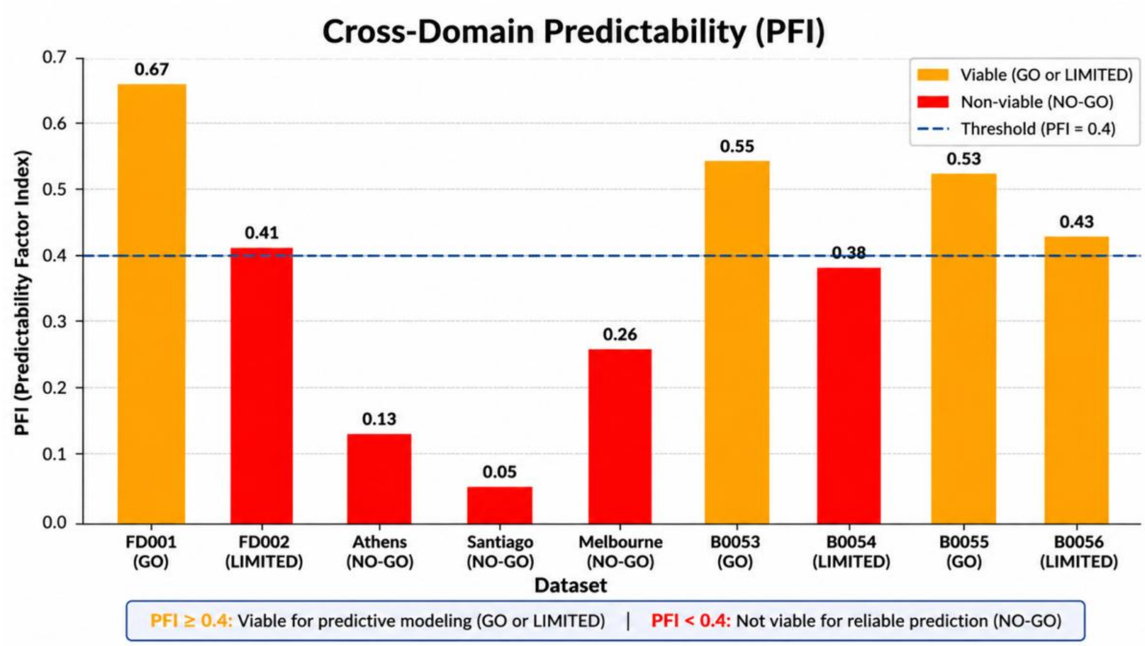

Figure 11 shows the resulting PFI values, together with the corresponding decision regimes. The datasets separate clearly into three categories: – PFI ≥ 0.4 → signals that support predictive modeling (GO or LIMITED) – PFI < 0.4 → signals that do not support reliable prediction (NO-GO) Several key observations emerge: First, predictive feasibility is not domain-specific. Datasets from different physical systems (turbofan engines, batteries, quantum calibration) exhibit the same separation into viable and non-viable regimes. Second, the presence of detectable structure does not guarantee predictive usefulness. Multiple datasets (e.g. Athens, Santiago, Melbourne) exhibit clear structural variation, yet fall below the feasibility threshold. Third, borderline cases (e.g. FD002, B0054, B0056) highlight a critical region where signals appear structured, but fail to produce stable predictive behavior. This is the regime where most modeling effort is typically invested — often without recognizing that the signal itself limits prediction. From a practical perspective, this figure defines a decision boundary: Signals below the threshold should not be used for predictive modeling, regardless of model complexity or data volume. This establishes a key result: Predictive modeling success is determined primarily by signal consistency, not by model selection.

Figure 12 – PFI Classification with Interpretation

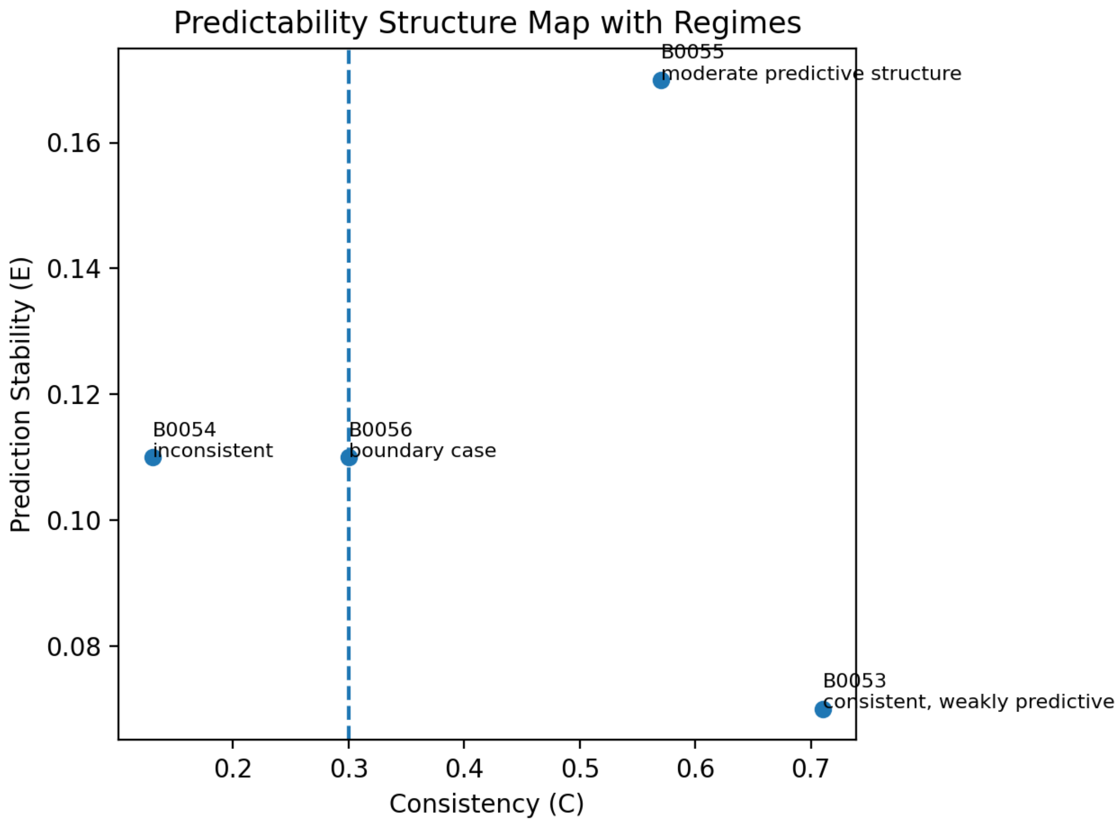

Figure 12. Predictive Feasibility Index (PFI) classification for battery cells. PFI values are computed as the average of structural behavior (S), consistency (C), and prediction stability (E), with thresholds defining NO-GO (<0.4), LIMITED (0.4–0.7), and GO (≥0.7) regimes. Although most cells fall into the LIMITED regime, this classification masks significant differences in underlying behavior, highlighting the importance of analyzing consistency and prediction stability separately.

Figure 13 – Predictability Structure Map with Regimes

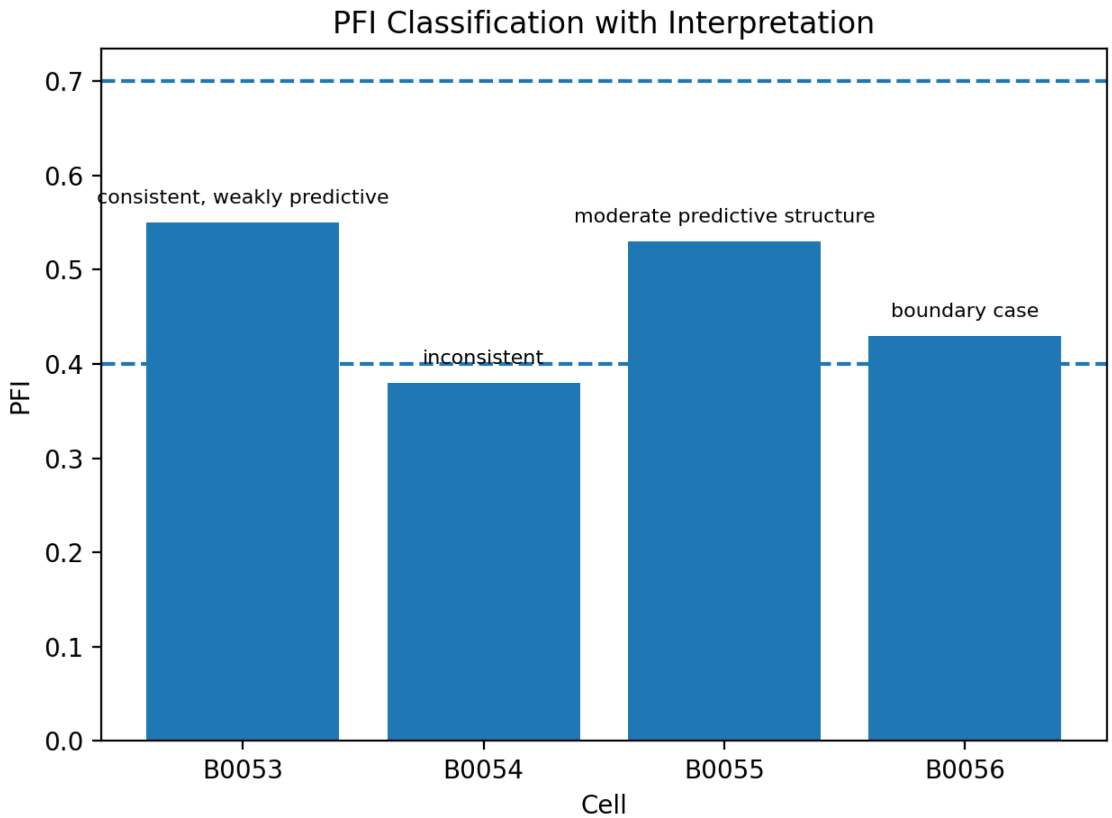

Figure 13. Predictability regimes defined in (C, E) space with consistency gating.
The vertical boundary at C = 0.3 defines the consistency gate below which prediction is fundamentally unsupported. The figure illustrates three regimes: NO-GO (low consistency), LIMITED (partial consistency), and GO (high consistency and prediction stability). Importantly, signals with high consistency but low prediction stability remain confined to the LIMITED regime, reinforcing that consistency is a necessary but not sufficient condition for prediction.

6 Prediction-Based Validation

To evaluate whether the proposed framework provides meaningful guidance for predictive modeling, we compare structural feasibility assessments with actual model behavior across representative datasets.

For each dataset, a simple predictive model (autoregressive or equivalent baseline) was applied to assess short-term prediction performance. The results show a consistent alignment between the proposed feasibility criteria and observed model behavior:

• FD001 (consistent structure):
Prediction error remains low and stable.
Model performance is reliable.
→ Classified as GO

• FD002 (inconsistent structure):
Prediction error varies significantly across runs.
Model performance is unstable and non-reproducible.
→ Classified as NO-GO

• IMS (predictable but non-informative):
Prediction error is low, but no meaningful structural evolution is present.
Model outputs remain stable but do not correspond to system-level behavior.
→ Classified as NO-GO

These results demonstrate that predictive failure is not random, but systematically associated with identifiable structural properties of the signal.
Figure 8 shows that consistency varies independently of prediction error, confirming that predictability alone is not sufficient for reliable prediction.

Importantly, the framework allows this distinction to be made prior to model development.

This establishes a key result:
Predictive feasibility can be assessed before modeling, based on signal structure alone.

6.1 Formal Definition of Predictive Feasibility

We define predictive feasibility as a property of the signal, independent of model choice. A signal is considered to support prediction if the following conditions are satisfied:

1. Structural detectability
The signal exhibits identifiable structural variation over time.

2. Cross-run consistency
Structural patterns are reproducible across multiple realizations or system trajectories.

3. Predictive stability
Prediction error remains bounded and stable under short-term forecasting.

These conditions should be interpreted as necessary, but not sufficient, for predictive feasibility. If any of these conditions fail, predictive modeling is expected to be unreliable or non-informative. This definition provides a minimal and model-independent criterion for determining whether prediction is meaningful for a given signal.

7. Detection Is Not Prediction

A central result is the separation between detectability and predictive relevance.

Structural changes can be detected reliably, yet remain unrelated to future system behavior under the considered signal representation.
This explains why:

- detection methods produce false positives
- predictive models fail despite apparent structure
- stable signals appear predictable but lack relevance

The combined effect is shown in **Figure 14**, where distinct regions of predictive feasibility emerge, highlighting that prediction depends on structural consistency rather than detectability alone. This establishes a necessary condition: without consistent structure, predictive modeling is fundamentally unsupported.

Figure 9 demonstrates that predictive feasibility directly corresponds to observed model behavior. Signals classified as consistent yield stable predictions, while inconsistent or non-informative signals produce unstable or misleading model outputs.

From a practical perspective, this figure defines a boundary between signals that support reliable prediction and those that do not.

In practice, failing to identify this limitation early can lead to continued investment in models that cannot converge to stable or meaningful predictions.

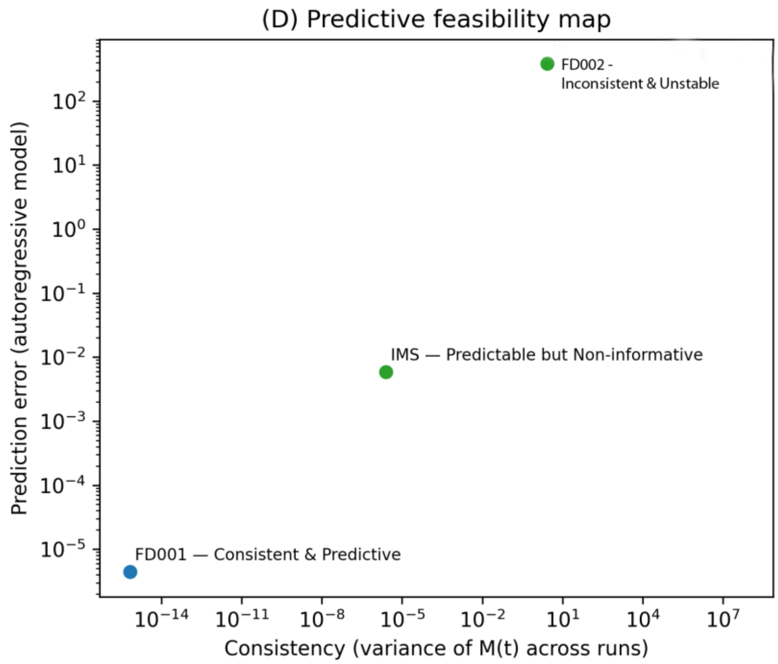

Prediction requires not only predictability, but critically the presence of consistent structural behavior across observations. This condition is rarely verified explicitly in practical applications.

Importantly, these results do not depend on the choice of dataset or domain, as confirmed across vibration data, battery systems, and quantum calibration datasets. Instead, they reflect a fundamental limitation: if the signal does not contain consistent and interpretable structure, no model can recover predictive information from it [7].

7.1 Failure Modes of Predictive Modeling

The analysis reveals several distinct failure modes that arise when predictive modeling is applied to signals without assessing feasibility.

1. Predictable but non-informative signals
Signals exhibit low prediction error and stable behavior, yet do not encode meaningful information about future system states.
2. Structured but inconsistent signals
Signals exhibit detectable structural variation, but patterns differ across realizations, preventing generalizable prediction.

3. Misleading structural signals
Signals exhibit apparent structural change that is unrelated to system-level outcomes, leading to false positives in detection-based approaches.

These failure modes explain common observations in predictive modeling practice, including:
• unstable model performance
• lack of generalization
• persistent false positives

Importantly, these effects are not due to model limitations, but arise from the underlying signal structure. This reinforces the central conclusion of this work:
Predictive modeling failure is often a consequence of insufficient information in the signal, rather than inadequate modeling.

8. Implications

8.1 Prediction is a property of the signal
Prediction failure is often inherent to the signal representation.

8.2 Modeling cannot compensate for missing information
No model can extract predictive information that is not present.

8.3 Predictive modeling should be conditional
Prediction should only be attempted when feasibility conditions are satisfied.

8.4 Practical implication
In practical applications, predictive modeling should be preceded by an explicit assessment of signal suitability.
This assessment defines a critical decision step:

GO → proceed with modeling; the signal contains consistent and reproducible structure.
LIMITED → proceed with caution; structure is present, but model instability is likely.
NO-GO → do not proceed with predictive modeling; the signal lacks predictive structure.

Without such assessment, modeling efforts may be fundamentally misdirected toward signals that cannot support prediction.

9. Limitations

This work focuses on signal-level representations. External system variables may provide additional information not captured in the signal.

In practical applications such as vibration-based monitoring, additional physical insight may complement signal-based interpretation [8].
The framework is empirical and characterizes observed behavior rather than providing formal guarantees.

For industrial predictive systems, the practical question should therefore be asked before model development begins:

Can this signal support prediction at all?

If this question is not addressed upfront, predictive modeling efforts may be fundamentally misdirected from the start.

10. Conclusion

Predictive modeling is fundamentally constrained by the information content of the signal, independent of model choice or data volume. This misattribution contributes to repeated modeling failures in practical systems.

Structural changes can be detected reliably, but their relevance to future behavior is not generally identifiable.

This leads to a central conclusion:
Prediction failure is often not a modeling problem, but a signal problem.
Accordingly, predictive modeling should begin not with model selection, but with an explicit assessment of whether prediction is possible at all.

In practice, failing to perform this assessment often results in building models on signals that were never predictive to begin with.

If you suspect your system may fall into one of these failure regimes, we offer independent feasibility assessments prior to model development.

Companies implementing this approach gain a competitive edge over peers who develop models blindly. This promotes a data-driven culture where decisions are informed by the predictability of signals, not solely by available data.

We apply this analysis as a pre-model diagnostic step to determine whether predictive modeling is viable.

Next Steps
Companies and teams facing challenges in predictive modeling are invited to engage with us for a predictive feasibility assessment. Understanding your data's potential before investing in modeling saves time, reduces cost, and increases the likelihood of actionable insights.
Contact us at info@tarh-energy.com to schedule an initial consultation.

References (IEEE)

[1] J. Lei et al., “Machinery health prognostics: A systematic review,” Mechanical Systems and Signal Processing, 2018.

[2] X.-S. Si et al., “Remaining useful life prediction — A review on the statistical data driven approaches,” European Journal of Operational Research, 2011.

[3] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” ACM Computing Surveys, 2009.

[4] M. Ahmed, A. N. Mahmood, and J. Hu, “A survey of anomaly detection techniques,” Future Generation Computer Systems, 2016.

[5] E. N. Lorenz, “Deterministic nonperiodic flow,” Journal of the Atmospheric Sciences, 1963.

[6] H. Kantz and T. Schreiber, Nonlinear Time Series Analysis, Cambridge University Press, 2004.

[7] G. E. P. Box, “Science and statistics,” Journal of the American Statistical Association, 1976.

[8] R. B. Randall, Vibration-based Condition Monitoring, Wiley, 2011.

[9] B. Lim and S. Zohren, “Time-series forecasting with deep learning: A survey,” Philosophical Transactions of the Royal Society A, vol. 379, no. 2194, 2021.

[10] X. Kong et al., “Deep learning for time series forecasting: A survey,” Artificial Intelligence Review, 2025.

[11] J. Quiñonero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, Dataset Shift in Machine Learning, MIT Press, 2009.

[12] R. Geirhos et al., “Shortcut learning in deep neural networks,” Nature Machine Intelligence, 2020.

[13] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, “Do ImageNet classifiers generalize to ImageNet?” Proceedings of the 36th International Conference on Machine Learning, 2019.

[14] K. A. Severson et al., “Data-driven prediction of battery cycle life before capacity degradation,” Nature Energy, vol. 4, pp. 383–391, 2019.

[15] G. dos Reis, C. Strange, M. Yadav, and S. Li, “Lithium-ion battery data and where to find it,” Energy and AI, 2021.

[16] G. Acampora et al., “A dataset for quantum circuit mapping,” Data in Brief, 2021.