LocaPerk

LocaPerk Labs Working Paper Series

Working Paper No. 002

LocaPerk Labs Research Note

MBD-mini Behavioural Stress Modelling and Temporal Probability Stabilization

---

Author: Damilola Olayemi

Founder, LocaPerk

Lead Researcher, LocaPerk Labs

Date: 2026

Public Research Note

locaperk.com/labs

dami@locaperk.com

Executive Summary

This research note documents how a streamed transaction dataset can be converted into behavioural signals, future stress proxies, and concern-ranking outputs suitable for early-warning triage. The work was designed as a methodological milestone for LocaPerk Labs rather than a final institutional benchmark. Its purpose is to show how behavioural transaction data can be engineered into interpretable signals, how those signals can support future stress ranking, and why probability stabilisation and percentile-based tiering are necessary when behaviour shifts across time. The results support a cautious operational conclusion: behavioural signals can help concentrate future concern into smaller review groups, and percentile-based concern tiers are more defensible than hard score thresholds when behavioural regimes shift.

Keywords

Behavioural financial intelligence • Behavioural signals • Liquidity resilience • Spending pressure • Cashflow volatility • Concern tiers • Early-warning systems

Table of Contents

1. The Financial Intelligence Gap

1.1 Why a Transaction-Led Research Note Matters

Transaction data offers one of the clearest windows into how financial life is evolving over time, yet it is still underused in operational financial intelligence. Traditional records can describe formal history, but they often say less about how resilience is changing month to month. Behavioural transaction analysis therefore provides a valuable research setting for testing whether earlier signs of stress can be detected before more formal warning indicators fully adjust.

This note focuses on a compact dataset for speed of iteration and engineering validation. The question is not whether this single dataset is sufficient for final production claims. It is whether the methodological path is credible: can streamed transaction records be transformed into behavioural signals, can those signals support future concern ranking, and can the resulting outputs be stabilised enough for percentile-based triage?

2. Behavioural Signal Engineering

2.1 From Raw Transactions to Interpretable Features

The first stage of the pipeline converts raw transaction activity into monthly behavioural snapshots. These snapshots summarise recurring financial behaviour through interpretable concepts rather than raw ledger history. Liquidity resilience approximates how much buffer appears to exist relative to typical outflows. Income stability describes the rhythm and consistency of inflows. Spending pressure captures whether spending appears to be intensifying relative to historical norms or incoming resources.

Additional features describe obligation coverage, cashflow volatility, payment consistency, and irregularity in recurring commitments. These are not presented as universal or final variables. They are pragmatic proxies designed to help institutions read behavioural stress in operational language. Their value lies in legibility as much as in predictive utility.

2.2 Why Feature Legibility Matters

- Liquidity resilience is easier to explain operationally than a latent hidden factor.
- Spending pressure and cashflow volatility align with how relationship and support teams already think about strain.
- Payment consistency and obligation coverage create a bridge between behavioural interpretation and risk oversight.

3. System Architecture

3.1 Behavioural Modelling Pipeline

The research workflow follows the same conceptual pipeline as the broader LocaPerk Labs programme. Transaction data is aggregated into behavioural features, behavioural features are used for temporal modelling, probability outputs are stabilised where possible, percentile concern ranking converts outputs into more usable orderings, and an operational interface interprets the result as concern tiers rather than hard decisions.

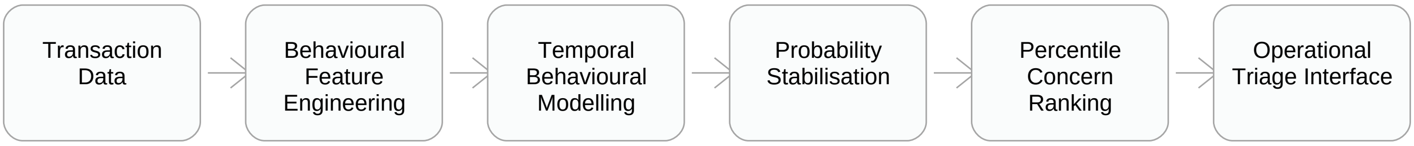

Figure 1: Behavioural Intelligence Pipeline

3.2 Why Probability Stabilisation Is a Separate Layer

A central lesson from transaction-based behavioural modelling is that strong ranking does not automatically imply stable probabilities. Behavioural regimes can shift. Cohorts can drift. Label prevalence can change. As a result, the system needs an explicit probability-stabilisation layer rather than assuming that a single model output can be used safely across all periods. This is why temporal validation and percentile-based interpretation are integral to the research design.

4. Behavioural Stress as a Future Outcome

4.1 Proxy Label Construction

The note treats behavioural stress as an emerging condition rather than a simple static label. Proxy future stress flags are built from patterns such as declining inflow stability, worsening spending-to-income relationships, falling liquidity resilience, and increasing irregularity in recurring obligations. This makes the target closer to an early-warning concept than a classical default label.

That choice is deliberate. LocaPerk is not being developed as a hard approve or decline engine. It is being developed as an early-warning and decision-support layer. The research target should therefore reflect later concern and deterioration rather than a narrow downstream event alone.

5. Ranking and Review Prioritisation

5.1 What Operational Success Looks Like

For this kind of system, the most important practical question is whether it can concentrate later concern into relatively small review groups. If an institution can review only a limited share of customers, does the behavioural ranking help those reviews capture a disproportionate share of future concern? This ranking-focused framing is more relevant than raw accuracy because the product is intended for queue prioritisation and triage.

The evidence from the MBD-mini work supports that framing. Behavioural transaction signals can improve the focus of small review buckets, which in turn makes them useful for watchlisting and proactive review. This does not imply full stability of absolute probabilities, but it does support the value of relative ordering.

6. Temporal Probability Stabilization

6.1 Why Time Changes the Interpretation

Temporal robustness is a central challenge for behavioural systems. When cohorts change, the same absolute output can mean something different from one period to another. The note therefore compares baseline outputs, calibration strategies, rolling retraining, recency weighting, and recent-cohort approaches to understand which strategy produces the most trustworthy operational probabilities.

6.2 Operational Conclusion on Probability Use

The most responsible conclusion is that probabilities are most useful when translated into percentile concern bands rather than treated as rigid universal cutoffs. Relative tiering is better aligned with real-world drift. It also reduces the temptation to over-claim precision in settings where the evidence is still developing.

Normal Monitoring
Below 80th percentile

Watchlist
80th to 95th percentile

Proactive Review
95th to 99th percentile

Urgent Support Flag
Top 1%

Figure 2: Example Concern Tier Structure

7. Product Interpretation

7.1 How This Research Fits the LocaPerk Platform

The MBD-mini work supports the same product interpretation as the wider LocaPerk Labs programme. The most defensible role for the platform is as an early-warning and decision-support layer. It is intended to support watchlisting, review prioritisation, proactive engagement, and vulnerability triage. It is not designed for automated credit approvals or declines. That boundary is both methodologically and operationally important.

8. Research and Validation Roadmap

8.1 What Comes Next

- Phase 1: Behavioural signal engineering from transaction data.
- Phase 2: Temporal modelling and probability stabilisation under cohort shift.
- Phase 3: Institutional pilot validation in shadow monitoring workflows.
- Phase 4: Explainability and governance frameworks for operational use.
- Phase 5: Longitudinal validation on richer institutional datasets.

The current note sits between phases one and two. It demonstrates that the engineering and modelling path is viable. It does not claim that the evidence base is already complete enough for broad production certainty. That next step will require richer time coverage and live institutional data.

9. Conclusion

9.1 A Credible Early Research Milestone

The MBD-mini research note shows that transaction-level behaviour can be translated into a practical behavioural intelligence layer. It supports the claim that liquidity resilience, spending pressure, cashflow volatility, and related behavioural signals can help identify emerging concern earlier than static records alone. It also reinforces a more cautious operational lesson: relative ranking and concern tiers are currently more robust than claims about universally stable probabilities.

That still represents meaningful value. Behavioural intelligence can complement traditional financial assessment, help institutions prioritise human review more effectively, and enable earlier support. The strongest current use case is triage and prioritisation. Future research should deepen this evidence with richer datasets, real institutional partnerships, and longer-term validation.

10. References

• Bazarbash, M., and Beaton, K. (2020). Filling the gap: Digital credit and alternative data. IMF Working Paper.

• Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. Journal of Finance, 77(1), 5–47.

• Hand, D. J., and Henley, W. E. (1997). Statistical classification methods in consumer credit scoring: A review. Journal of the Royal Statistical Society: Series A, 160(3), 523–541.

• Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2022). Forecasting and temporal instability in applied predictive systems. International Journal of Forecasting.

• Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1, 206–215.