From Outliers to Insight: Using Anomaly Detection to Strengthen NHS Oversight of Credit and Fuel Card Spend

INTRODUCTION

NHS organisations process thousands of credit and fuel card transactions each year, yet manual audits struggle to spot unusual or high-risk activity early. Using multi-year spend data ranging from £1.63 to £18,295.20 across diverse cardholders and categories, machine-learning was applied to proactively detect anomalous spending patterns.

AIM

- Identify transaction and behavioural anomalies
- Highlight high-risk spend patterns
- Support audit teams with clear insights

METHODS

HOW ANOMALIES WERE DETECTED

The model analysed transaction value, spending frequency and timing, category patterns, cardholder behaviour, refund activity, and credit limit rules. These features were normalised to ensure fairness between high and low spenders.

K-Means clustering was used to distinguish typical spending profiles and identify outliers, supported by distance-based anomaly scoring to rank transactions by risk level. Diagnostic checks using STL decomposition confirmed that anomalies were not seasonal patterns.

WORKFLOW

End-to-end anomaly detection pipeline.

Clusters Highlighting Transaction Outliers

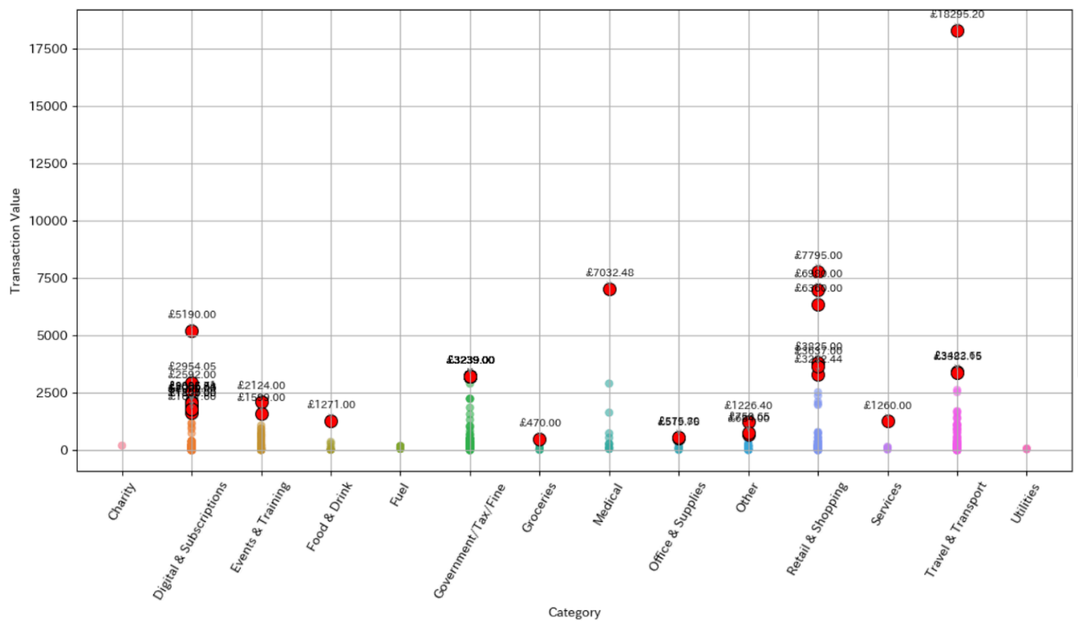

Top 5 Cardholders by Anomaly Rank

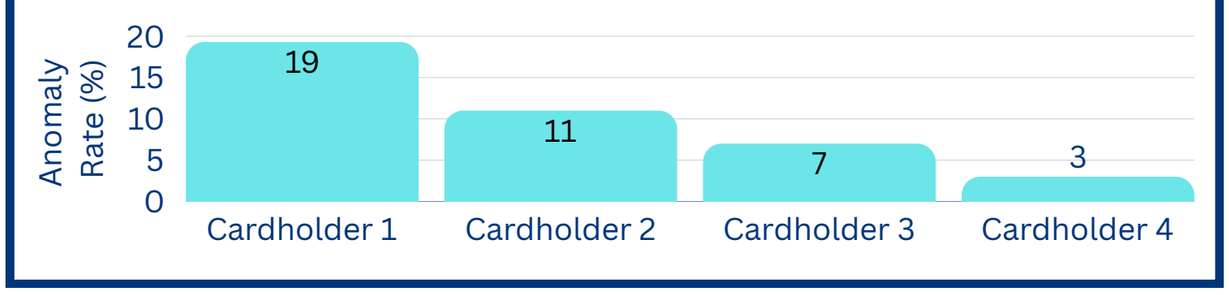

POLICY & BEHAVIOURAL ANOMALIES

There were 145 instances where monthly spending exceeded the £10k policy limit, including a highest breach of £26,302 and total above-limit values of £105,595 in 2023 and £57,717 in 2024. Behaviour-based anomalies included sudden spending spikes, high-frequency transactions, repeated similar-value payments, and irregular refunds without operational justification, indicating potential governance or approval risks.

IMPACT & LESSON LEARNED

IMPACT

- Focused audit attention on the top 5-10% highest-risk transactions, revealing spending patterns that manual checks would not identify.
- Strengthened oversight and fraud-prevention assurance, with an approach that can be scaled across other NHS card types and organisations.

LESSONS LEARNED

- Better data quality and normalisation improve anomaly accuracy and enable fair comparison between cardholders, but outliers still require known explanations rather than assumptions of misuse
- Anomalies are reviewed with the client and the model is refined so that future dashboards show updated, validated outliers — creating an iterative approach that improves accuracy each year and supports continuous reuse