Pattern Recognition and Anomaly Detection in High-Frequency Payments Data*

Ajit Desai[1] and Anneke Kosse[2]

May 16, 2022 - **Preliminary**

[1]Bank of Canada
[2]Bank for International Settlements

*The opinions here are those of the authors and do not necessarily reflect those of the Bank of Canada, the Bank for International Settlements or its Committee on Payments and Market Infrastructures. We thanks Leonard Sabetti for his contribution during earlier stages.

Motivation

⇒ High-value payments systems (HVPSs) are **core national infrastructures**. To ensure their safety it is crucial to understand participants' payment patterns

⇒ **Deviation from usual patterns** could hint at payments fraud, money laundering, a cyber-security event, terrorism financing, or an operational issue

- 2019 BIS Report: **highlights the importance** of using data & tools for wholesale payments fraud detection and prevention (at participant and system level)
- 2019 FED Report: **emphasizes the necessity** of and outlines steps to enhance payments system safety, lead to Fraud-Classifier model to sort unusual payments
- 2019 TARGET Report: **underlines the threat** posed by fraud in payments systems and recommends advance approaches to detect fraudulent transactions

Questions

$\Rightarrow$ What are the economic benefits of **real-time monitoring** in payments systems?

- What is the **system-wide impact** of potential anomalous payments?
- How to detect anomalous payments at the **system level**?
- (Our focus) what **type of model** to use to effectively predict participants' payment patterns and how to segregate anomalous payments in real-time?

Literature

$\Rightarrow$ Theoretical and empirical papers on payments timing:

- Bech and Garratt (2003), Martin and McAndrews (2008): Liquidity-delay trade-off
- McAndrews (2002), Bech and Garratt (2012): Impact of disruptions on timing
- Nellie Zhang (2015): Changes in payment timing in Canada's LVTS

$\Rightarrow$ Pattern recognition and anomaly detection in payments systems:

- Triepels et al. (2017): Anomaly detection in RTGS
- Sabetti and Heijmans (2020): Detecting anomalous flows in the Canadian ACSS
- Rubio et al. (2020): Classifying payment patterns with artificial neural networks
- Léon et al. (2020): Pattern recognition of FI's payment behavior
- Arévalo et al. (2022): Clustering of anomalous payments in salvadorian HVPS

Objective

$\Rightarrow$ Examine potential efficacy of advance **machine learning (ML)** tools for payment pattern recognition and anomaly detection in HVPS:

- Propose **layered approach** to segregate and study usual and unusual transactions
- Use the ML model to predict the **submission time** of payments in HVPS
- Use historical **transaction level** settlement data to learn usual patterns
- Evaluate the model on artificial anomalous transactions and on the **real data**
- **Interpret** model to study the impact of transaction features on payment patterns

Methodology

⇒ **Layered approach** for pattern recognition and anomaly detection:

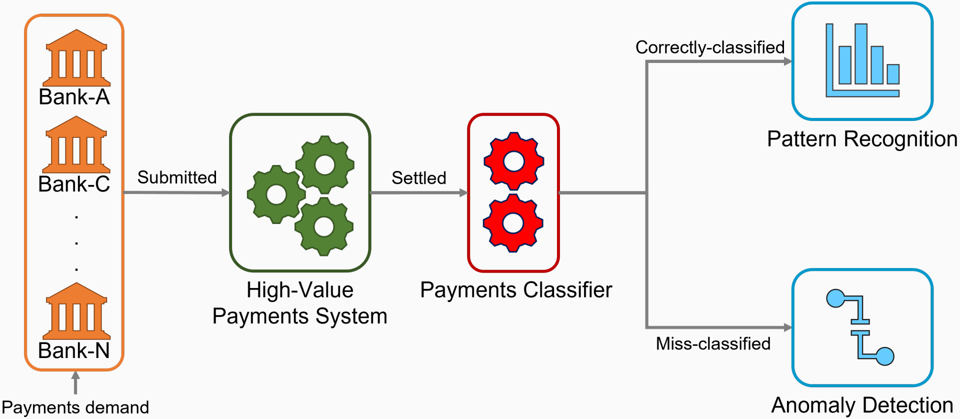

Results Preview

⇒ **This layered approach can help systematically focus on subsets of payments** for pattern recognition and anomaly detection, and could serve as a basis for real-time monitoring in HVPS.

- **Gradient boosting** model used as a **payments classifier** is able to predict payments timing with 96% accuracy (and outperform logistic regression by 35%)
- It can be used to study the impact of transaction features on payment patterns

- **Isolation forest** model used as an **anomaly detector** provides an effective way to identify and analyze anomalous transactions
- It can be interpreted to understand **potential causes** of anomalous payments

Outline

1. Data
2. Methodology
3. Preliminary Results
4. Conclusions & Next Steps

Data

Data: HVPS Transactions

⇒ **Transaction-level** settlement data from LVTS for three years period (2017-2019)

- Two settlement mechanisms (Tranche 1 (T1) and Tranche 2 (T2))
- 22 million transactions (settled between 6am and 6pm)

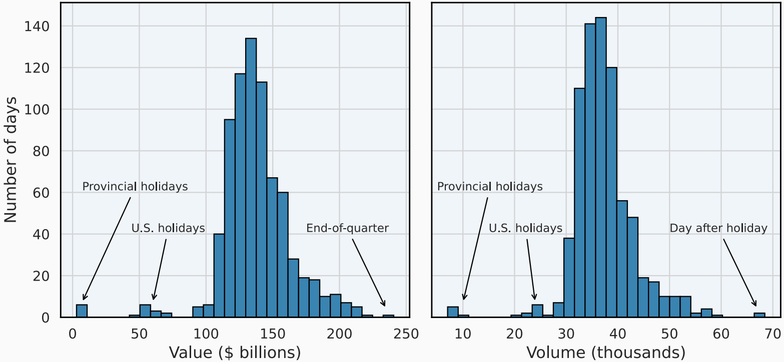

Data: Transaction Features

⇒ For each transaction settled in the system, we extract four sets of 20 features:

- **Transaction features**: sender, receiver, amount, payment-type, and tranche
- **Liquidity features**: total collateral pledged by the sender, sender's bilateral and multilateral credit limits, system-level liquidity, and overnight money market rate
- **Timestamp features**: year, month of the year, week of the year, day of the month, and day of the week
- **Intraday timing features**: The time elapsed since the last payment by the same sender (any type or same type as current payment) to any receiver (or to the same receiver), the time elapsed since the start of the current period

Methodology

Model: Payments Classifier using Supervised Learning

⇒ **Binary classifier:** Artificial neural network, logistic regression, gradient boosting

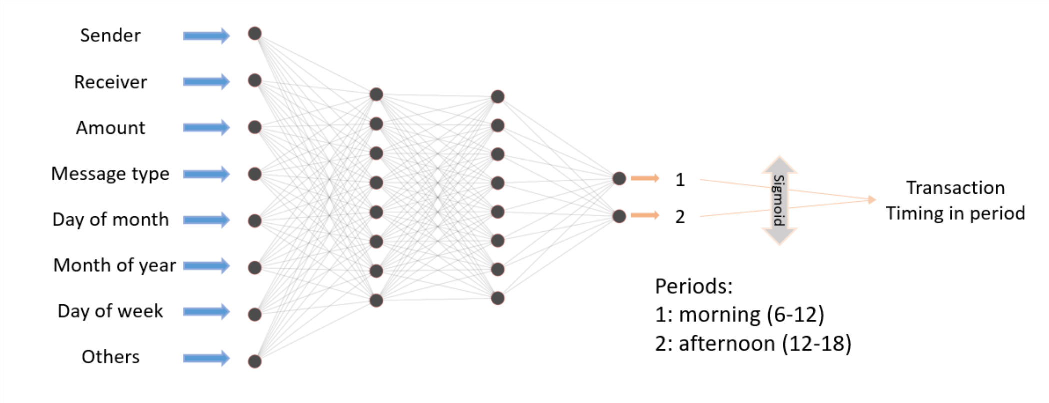

Model: Anomaly Detector using Unsupervised Learning

⇒ **Isolation forest:** Decision tree based anomaly (outlier) detection algorithm

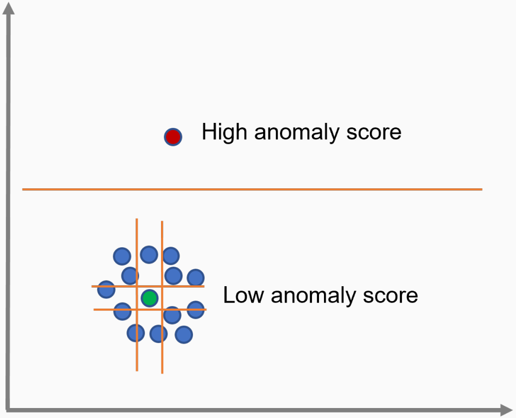

Model: Interpretation using Shapley values (SHAP)

⇒ **Example**: Consider classification is a "game", then the Shapley values can be used to fairly distribute the *payout* (= the prediction) among the *players* (= the predictors)

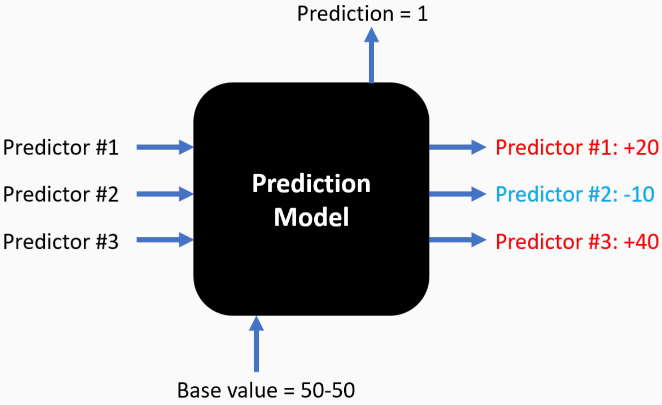

Preliminary Results

Payments Classifier on LVTS-T1 Dataset

⇒ **Model performance:** Normal days (clean) data for training and mix of clean and special days (holidays, post-closure and operational incident) data for testing

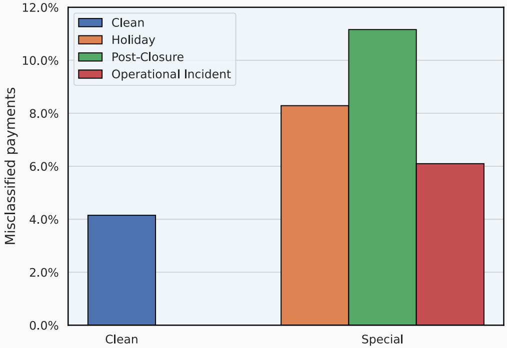

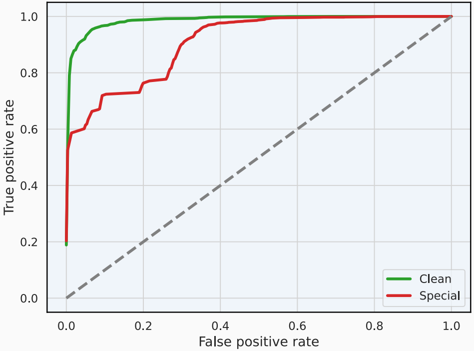

Payments Classifier on LVTS-T1 Dataset

⇒ **Average impact** of transaction features on the model output during training

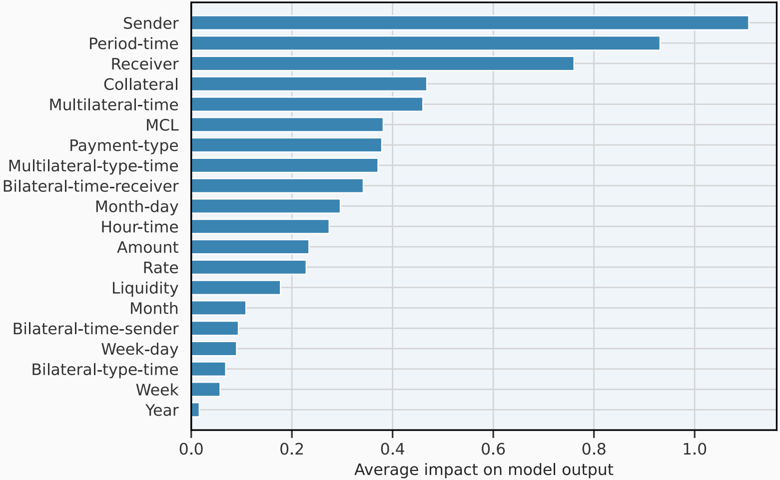

Payments Classifier on LVTS-T1 Dataset

⇒ **Interaction** between intraday timing and its impact on prediction

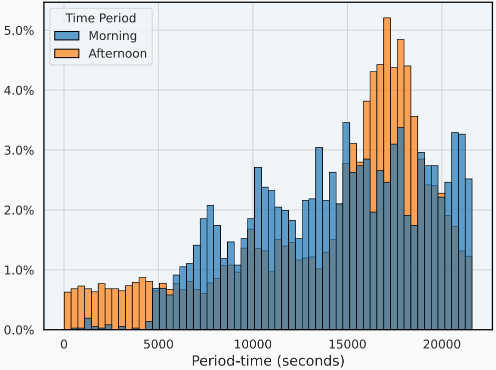

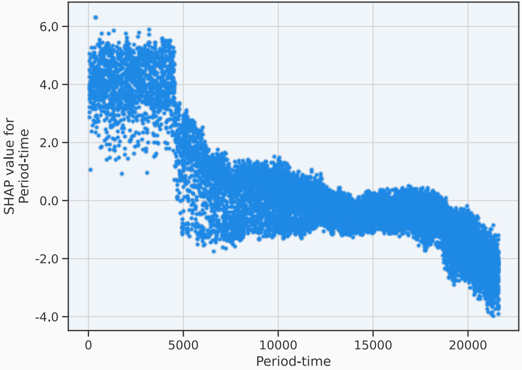

Payments Classifier on LVTS-T1 Dataset

⇒ **Interaction** between intraday timing (multilateral) and its impact on prediction

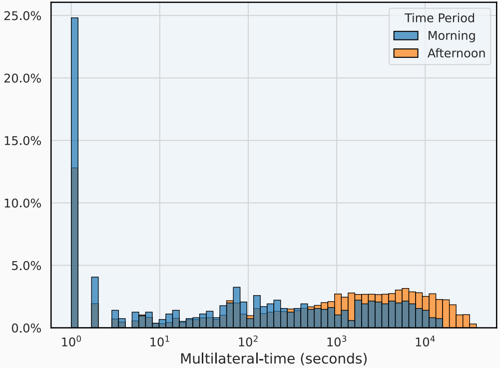

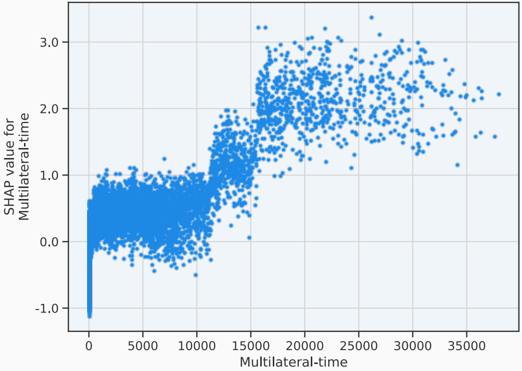

Payments Classifier Evaluation on Artificial Transaction

⇒ **Impact of features** on the model output for the indivisible payments. (Top) actual transaction, (bottom) with artificially manipulated period-time

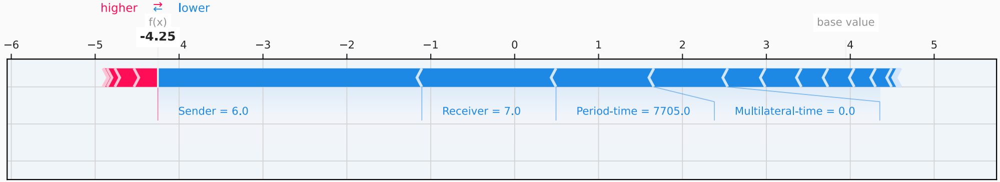

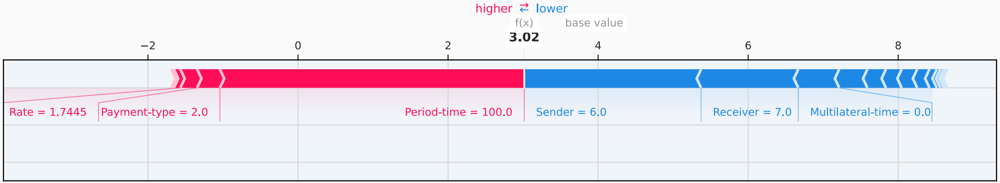

Low score (-4.25) → morning (more blue); and high score (3.02) → afternoon (more red)

Payments Classifier Evaluation on Special Days Transactions

⇒ **Impact of features**: (Top) predicted as morning but sent in afternoon (operational incident). (Bottom) predicted as afternoon but sent in morning (post closure)

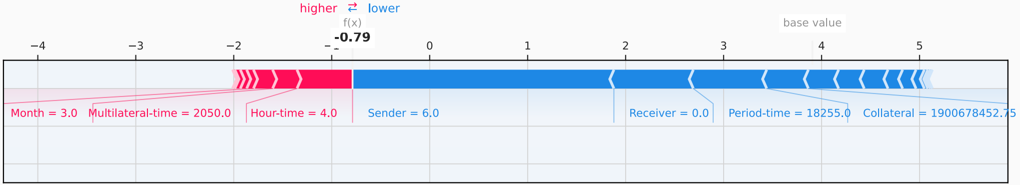

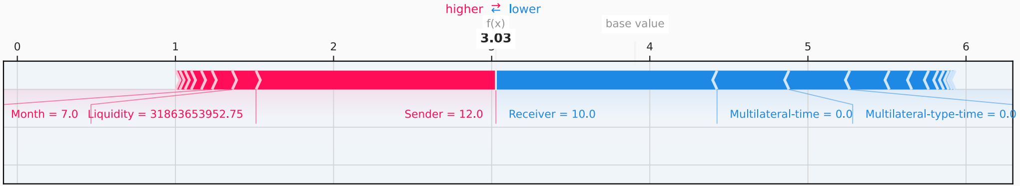

Anomaly Detector on LVTS-T1 Dataset

⇒ **Average impact** of transaction features on the model output during training

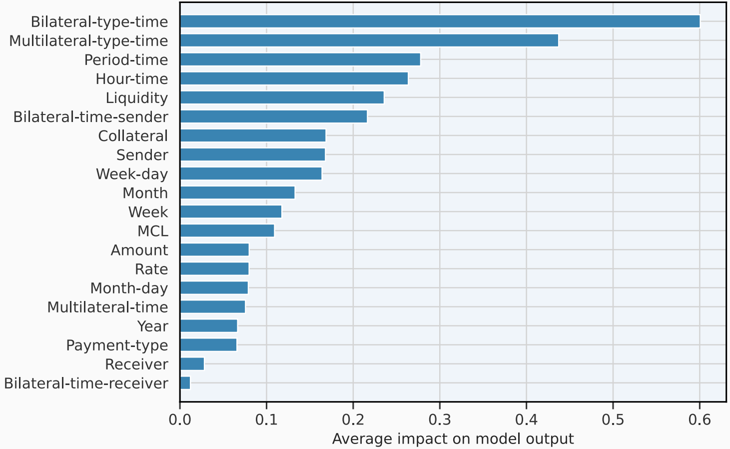

Anomaly Detector Evaluation on Artificial Transactions

⇒ **Anomaly scores:** Scores for the individual transactions for the subset of miss-classified transactions (original and artificially manipulated)

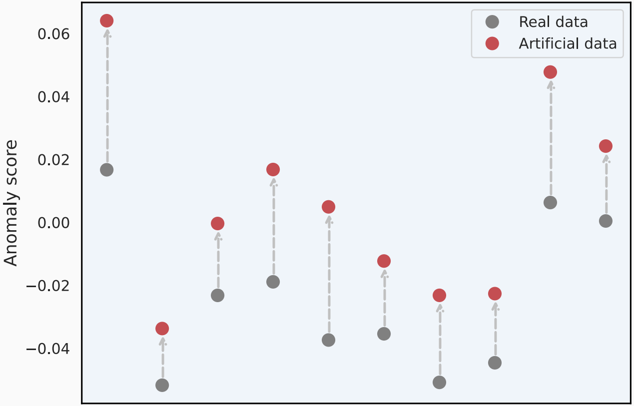

Anomaly Detector Evaluation on Real Transactions

⇒ **Anomaly scores:** Scores for the individual transactions for the subset of miss-classified transactions on regular (clean) and special days

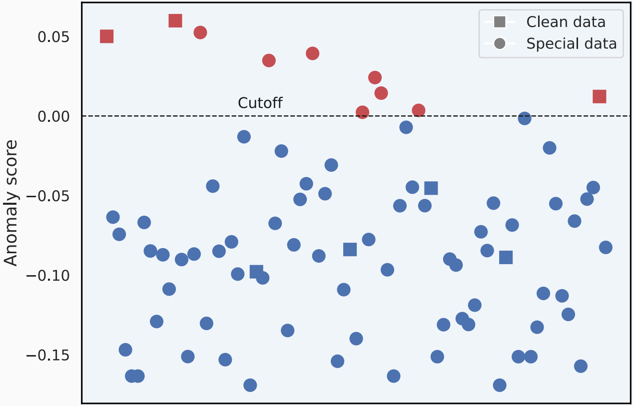

Conclusions & Next Steps

Conclusions

- **Layered approach simplifies the problem** and provides a systematic way to handle anomalies—especially in large datasets without (or few) known anomalies
- **Payments classifier** could help to classify the set of transactions into usual and unusual with $> 95\%$ accuracy
- **Anomaly detector** could be used to classify and rank anomalies
- **Model interpretation** (at local and global level) could be helpful for monitoring
- **Model is transferable**: trained on LVTS data, but could be used on Lynx data (Canada's HVPS since Aug 2021) to understand the potential change in patterns

Next Steps

⇒ More data (10 years), multi-class classification setup (3 periods model)

⇒ Test on LVTS-T2 transactions and Lynx settlement data

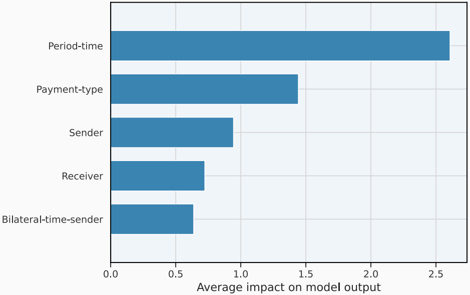

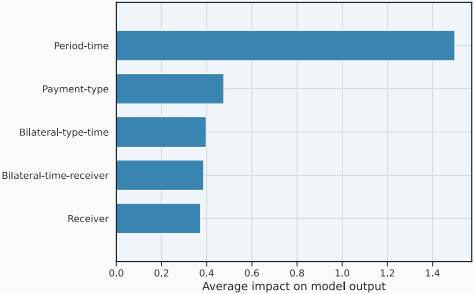

⇒ Use the model to study the impact of system characteristics—on top of transaction features—on payment patterns

→ Different settlement mechanisms (LVTS T1/T2, Lynx LSM/UPM)

→ Different payment systems (LVTS and Lynx)

Thank you!

References

- BIS-Report (2019), Reducing the risk of wholesale payments fraud related to endpoint security: a toolkit, Bank for International Settlements
- FED-Report (2019), Federal reserve fraud-classifier model, Technical report
- TARGET-Report (2019). Target Annual Report, European Central Bank
- Ke et al. (2017). LightGMB: A highly efficient gradient boosting decision tree, Advances in Neural Information Processing Systems
- Liu et al. (2008). Isolation forest, IEEE conference on data mining
- Lundberg et al. (2017). SHAP: A unified approach to interpreting model predictions

Gradient Boosting Machines

Gradient boosting is a decision tree (DT)-based non-parametric ensemble learning approach, here the sequence of weak learners, i.e., DTs are built on a repeatedly modified version of the training set.

For a given input features X and for each instance i using n DTs represented as T,

$$H_n(x_i) = \sum_{n=1}^{N} T_n(x_i),$$

Where the $H_n(x)$ is built as

$$H_n(x) = H_{n-1}(x) + \gamma h_n(x),$$

where γ is the learning rate used to regularize the contribution of new weak learner.

Isolation Forest

Isolation forest algorithm is designed with the idea that anomalies are few and distinct using decision trees. The procedure can be explained using the following steps:

- From given training data, a random sub-sample is selected and assigned to an DT
- From selected sub-sample, a random subset of features are chosen to build DT using a random threshold at each split
- The splitting of DT process is repeated until each data point is completely isolated or until max (predefined) depth is reached
- The above steps are repeated to construct many DTs by choosing different random subsets of features and sub-samples of data
- The anomaly score is then assigned to each data points in training sample based on the depth of the trees required to isolate each data point

Shapley Values

Shapley values: A method from coalitional game theory which provides a way to fairly distribute the *payout* among the *players* by computing average marginal contribution of each player across all possible coalitions.

Theorem: For player i in a coalition game (N, v):

$$\phi_i(N, v) = \underbrace{\frac{1}{N!} \sum_{S \subseteq N \setminus \{i\}}}_{\text{average over all } S} \underbrace{|S|! \left(|N| - |S| - 1\right)!}_{\text{possible coalitions}} \underbrace{\left[v(S \cup \{i\}) - v(S)\right]}_{\text{marginal value}}$$

where, N number of players, v is payoff (value) function, S are sets of coalitions