Matrix profile enhanced Bayesian online change point detection for Bitcoin ETF trading

Masoud Neshastehriz
with Paul Laux

Institute for Financial Services Analytics
Alfred Lerner College of Business & Economics
University of Delaware

April, 2025

Overview

Dynamic trading strategies and portfolio allocation methodologies depend on a characterization of assets' return distributions

- Distributional characteristics are time varying, and changes are difficult to anticipate
- Time series dependencies in distributional characteristics are likely but also difficult to characterize, especially in real time

Machine learning approaches have been explored as a solution

- Lack of explainability and tendency to overfit with limited and low signal-to-noise market data
- **Wood et al. (2023)** develop a deep learning-based approach that demonstrates improved performance in momentum trading

This paper:

- We develop an explainable online workflow for detecting changes in distributional characteristics and detecting repeatable regimes in characteristics
- We apply it to a newly-available cryptocurrency exchange traded fund and to a major stock index exchange traded fund to understand the extent to which investible crypto may be similar or different from traditional assets (perhaps due to being more speculative or sentiment-driven), and whether the workflow can be a useful guide for trading

Key methods in our workflow

- Bayesian online change point detection (BOCPD) (**Adams and MacKay, 2007**), a method for detecting shifts in distributional parameters in real time
  - Bayesian change point detection is little-explored in the finance literature
  - **Thies and Molnár (2018)** used Bayesian models to detect regime changes in Bitcoin in an offline setting (required full dataset for analysis)
- Matrix profile analysis (MPA) (**Yeh et al., 2016**), a method to find repeated similarities in time series data in real time
  - The usefulness of MPA has not been explored at all in the finance literature

Focal market: IBIT, iShares Bitcoin Trust ETF

- Launched only in 2024, after SEC objections were struck down in court
- Institutional investor interest in spot bitcoin ETFs have been very strong since they remove compliance, custodial, and regulatory barriers
- An exchange traded fund that holds actual bitcoin, thus tracking crypto price patterns in a way that is both investible and not affected by derivative expiry cycles

Research Questions

1. Can a combined Bayesian Online Change Point Detection (BOCPD) and Matrix Profile Analysis (MPA) framework detect patterns in real time?
   - Traditional offline methods cluster return distributions, but we ask whether real-time algorithms can identify when regimes change and whether similar regimes have been seen before.
2. Do IBIT return patterns repeat over time?
   - That is, are there recognizable sequences in return data that recur, which could be exploited for market timing or risk mitigation?
3. How does IBIT compare to SPY, a traditional, highly liquid equity ETF in these regards?
   - SPY serves as our benchmark for pattern recurrence and model effectiveness.
4. Does a simple trading strategy that depends on repeatable patterns perform reasonably?
   - We evaluate whether exploiting recurring patterns leads to practical, out-of-sample trading performance.

Data Overview and Methodological Framework

Data

- Intraday trading data for IBIT and SPY
  - Sourced from Alpha Vantage using Nasdaq-licensed APIs
  - Covers the full 2024 calendar year for IBIT
  - Includes all trading sessions:
    - Pre-market (4:00 AM ET)
    - Regular hours (9:30 AM–4:00 PM ET)
    - Post-market (until 8:00 PM ET)
  - Returns are calculated over 5, 15, 30, and 60-minute intervals

Methodological Framework

- Bayesian Online Change Point Detection (BOCPD)
  - Detect when the behavior of asset returns changes in real time
- Matrix Profile Analysis (MPA)
  - Detect repeating patterns

Bayesian Online Change Point Detection (BOCPD)

Adams, MacKay (2007) identifies change points as they occur by estimating the probability that a new regime has begun.

- Detects shifts in distribution parameters like mean and variance without needing future data.
- Continuously updates with each new data point.
- Change point declares if P(change at time t | data up to t) $>$ 0.5.
- **Library used:** `Y-Bar/BOCPD`

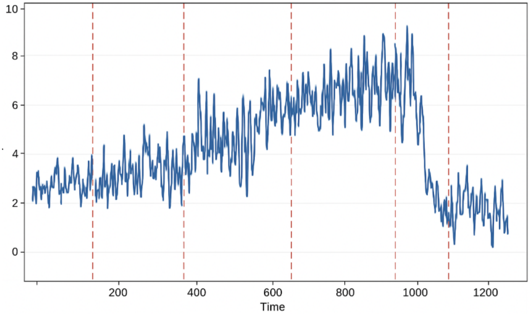

Image source: ChatGPT

Matrix Profile Analysis (MPA)

Yeh et al. (2016) detect recurring patterns, not change points

- MPA identifies motifs (similar sequences)
- Compares windows using similarity measures, typically Euclidean distance
- Causal implementation for real-time analysis
  - We apply "causal" MPA—only past windows are used for comparison
  - Ensures results are usable in a live trading context

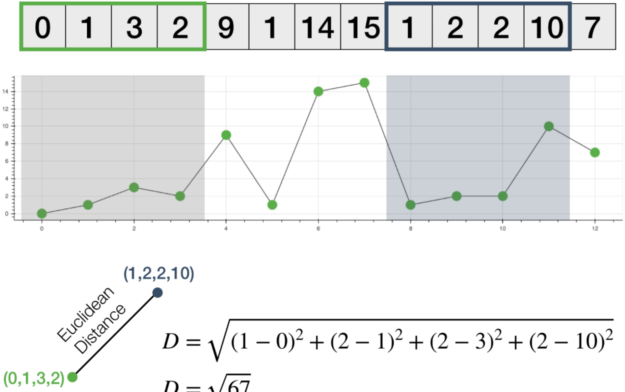

$$D = \sqrt{(1-0)^2 + (2-1)^2 + (2-3)^2 + (2-10)^2}$$

$$D = \sqrt{67}$$

Image source: `https://stumpy.readthedocs.io/en/latest/Tutorial_The_Matrix_Profile.html`

Workflow Summary

Our analysis follows a 5-step methodology to detect, segment, compare, and evaluate time-series regimes across the two ETFs:

1. **Apply BOCPD** to intraday returns. This partitions the return series into contiguous segments ("runs") with stable statistical properties (mean and variance).
2. **Create windows and Run Matrix Profile Analysis (MPA)** on each window to identify nearest past neighbors based on Euclidean distance for different windowing strategies like half-day
3. **Windows are mapped to BOCPD-detected "runs"**, aligned with different windowing strategies like half-day
4. **Find the matched runs** using the nearest neighbors in the second step if they pass the coverage threshold.
5. **Group similar runs into "regimes"** and calculate metrics such as match rate, regime size, and coverage to evaluate pattern recurrence.

Example: Matching Runs and Clustering

- Consider a run with windows at time stamps: **10, 11, 12, 13, 14**.
- Its nearest neighbors have time stamps: **3, 7, 6, 5, 9**.
- Another run has windows at time stamps: **2, 3, 4, 5, 6**.

Matching Criterion (50%, 80%, 90%)

Using a **50% cover threshold**, we compare the two runs. They share **3 out of 5** windows (**3, 5, 6**), so they are considered a **match**.

Clustering into Regimes

Matched runs that exhibit direct connections (strong similarity) are grouped into clusters called **"regimes"**.

BOCPD Results: Change Point Visualization

Detected change points reveal structural contrasts

Panel A: IBIT 60 minute returns and change points

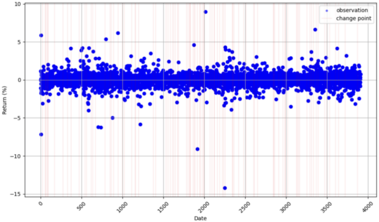

Panel B: SPY 60 minute returns and change points

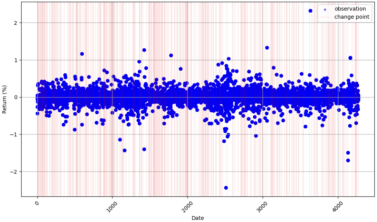

- SPY shows more frequent change points, reflecting a more dynamically evolving return process
  - SPY tracks a broad equity index, sensitive to macroeconomic shifts, earnings news, and global events

Analyzing distributional parameter shifts for two consecutive runs

IBIT

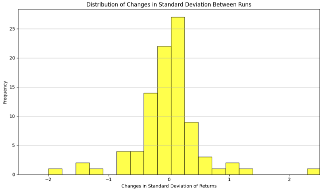

SPY

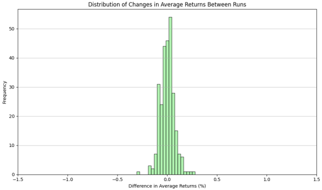

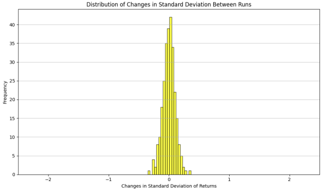

- IBIT shows larger parameter shifts than SPY, but they are somewhat less frequent (as shown on previous slide)
- SPY undergoes more frequent parameter shifts (as shown on previous slide), though with smaller magnitude

Identifying regimes: IBIT big picture

Visualizing the structure of regimes found in IBIT

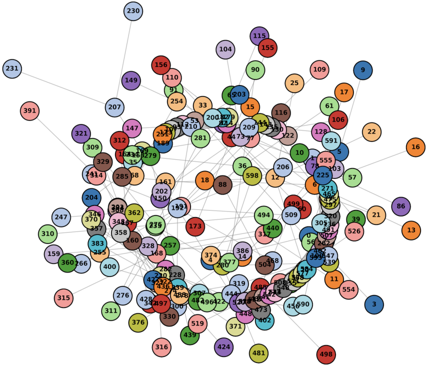

- Each node represents a detected run, and edges connect runs with sufficient matching windows (e.g., here, 50% overlap)
- Too many runs but not too many matches

MPA Results

Contrasting pattern structures in IBIT vs. SPY

Dataset	Number of Runs	Number of Matches	Number of Regimes
IBIT	3016	782	219
SPY	5205	1485	340

- IBIT has somewhat fewer changepoints (although recall they are more substantial)
- In the MPA, IBIT also shows fewer matches and fewer regimes than SPY— may seem to suggest that pattern-based trading tactics may not work

IBIT and SPY are similar when it comes to repeatability

Regime Coverage: Are we often in a regime?

$$\text{Regime Coverage} = \frac{\text{Total number of runs in all regimes}}{\text{Total number of runs}}$$

Ticker	Regime Coverage (Threshold 90%)
IBIT	33%
SPY	35%

Components of Regime Coverage

$$\text{Regime Coverage} = \text{Size of Regimes} \times \text{Importance of Regimes, where:}$$

$$\text{Size of Regimes} = \frac{\text{Total number of runs in all regimes}}{\text{Number of regimes}}$$

$$\text{Importance of Regimes} = \frac{\text{Number of regimes}}{\text{Total number of runs}}$$

$$\text{IBIT:} \quad 0.33 = \frac{219}{3016} \times \frac{995}{219} = 4.5443 \times 0.0726$$

$$\text{SPY:} \quad 0.35 = \frac{340}{5205} \times \frac{1821}{340} = 5.3559 \times 0.0653$$

- Regime coverage, regime importance, and regime size are roughly similar for IBIT and SPY

Is there enough repeatability to be useful?

Regime-Based Trading Strategy

- Data Splitting
  - In-sample (training) – used for regime identification.
  - Out-of-sample (testing) – used for strategy implementation.
- Regime Characterization
  - Mean return (μ_r)
  - Standard deviation (σ_r)
  - Assign portfolio weights using the mean-variance approach.
- Trading Logic
  - At each time t, identify the current regime and apply the corresponding weight w_r.
  - Allocate capital accordingly in period $t + 1$; compute the realized return.
- Performance Evaluation
  - Benchmark against a passive buy-and-hold strategy.

Period	Strategy Return	Buy-and-Hold Return
Q3	41%	0.05%
Q4	49%	47%

Conclusion

- **Hybrid, real-time detection framework**
  - Combined BOCPD and MPA to analyze return regimes and pattern recurrence
  - Applied to IBIT (crypto ETF) and SPY (traditional equity ETF)
- **Key findings and implications**
  - SPY shows more frequent parameter shifts, but they tend to be smaller
  - IBIT and SPY are nonetheless similar when it comes to regime coverage
  - A simple trading strategy using IBIT regimes achieves an apparent improvement over buy-and-hold, suggesting that the repeatability of patterns is enough to be useful—and that the workflow has detected actual economic patterns

Thank You!

Questions? Comments?

mneshast@udel.edu
laux@udel.edu