R42 Institute Fellowship

Stock Market Prediction

Our Team

Carlyle (Lyle) Morris

BS Computer Science @ Texas A&M University

Brian Lee

Master of Business Administration, Cornell University

Adi Lankipalle

Sophomore at the Wharton School, UPenn - Huntsman Program

Ektor Avlonitis

MSc in Computational Finance @ UCL

Iheb Hajji

MSc in Financial Engineering @ University of Montreal

Our Goal

Can we increase returns without taking on more risk?

We set out to do this by leveraging the predictive advantage of AlgoDynamix Flags to develop a **stock trading strategy to maximize returns while effectively managing risk.**

Risk Constraints:

- Long-only (no short selling)
- No leverage
- No derivatives or complex financial instruments

Table of Contents

R42

Part 1:
AlgoDynamix Background

What is AlgoDynamix?

- AI-based financial price forecasting analytics
- Functions as the key vehicle of the R42 Hedge Fund
- AlgoDynamix EXCLUSIVELY uses real-time data (no historical data)
- The technology focuses on identifying potential risks before they manifest, allowing for proactive decision-making

Detecting Price Movements Via Flags

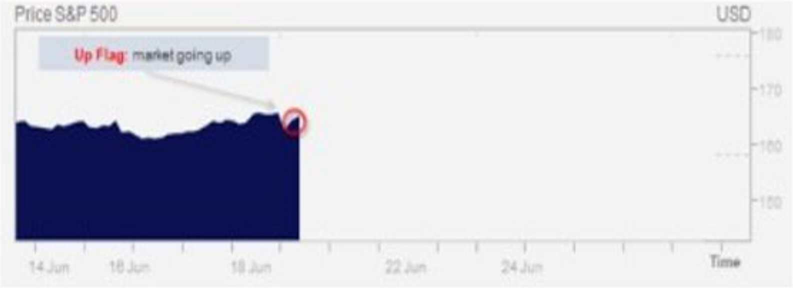

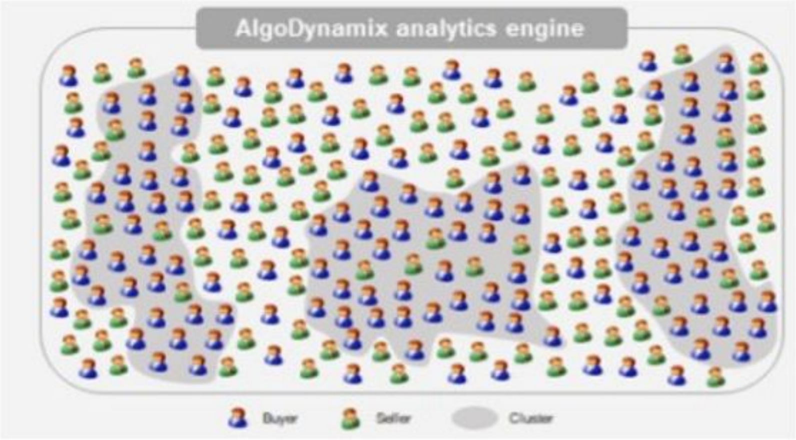

- Model generates Flags via unsupervised machine learning
- Utilizes order book (buy and sell history of an asset)
- Clusters created based on common feature sets → creates respective type of flag
- Issues an up or down flag when it detects activity likely to alter the price of an asset

AlgoDynamix Flags Illustration

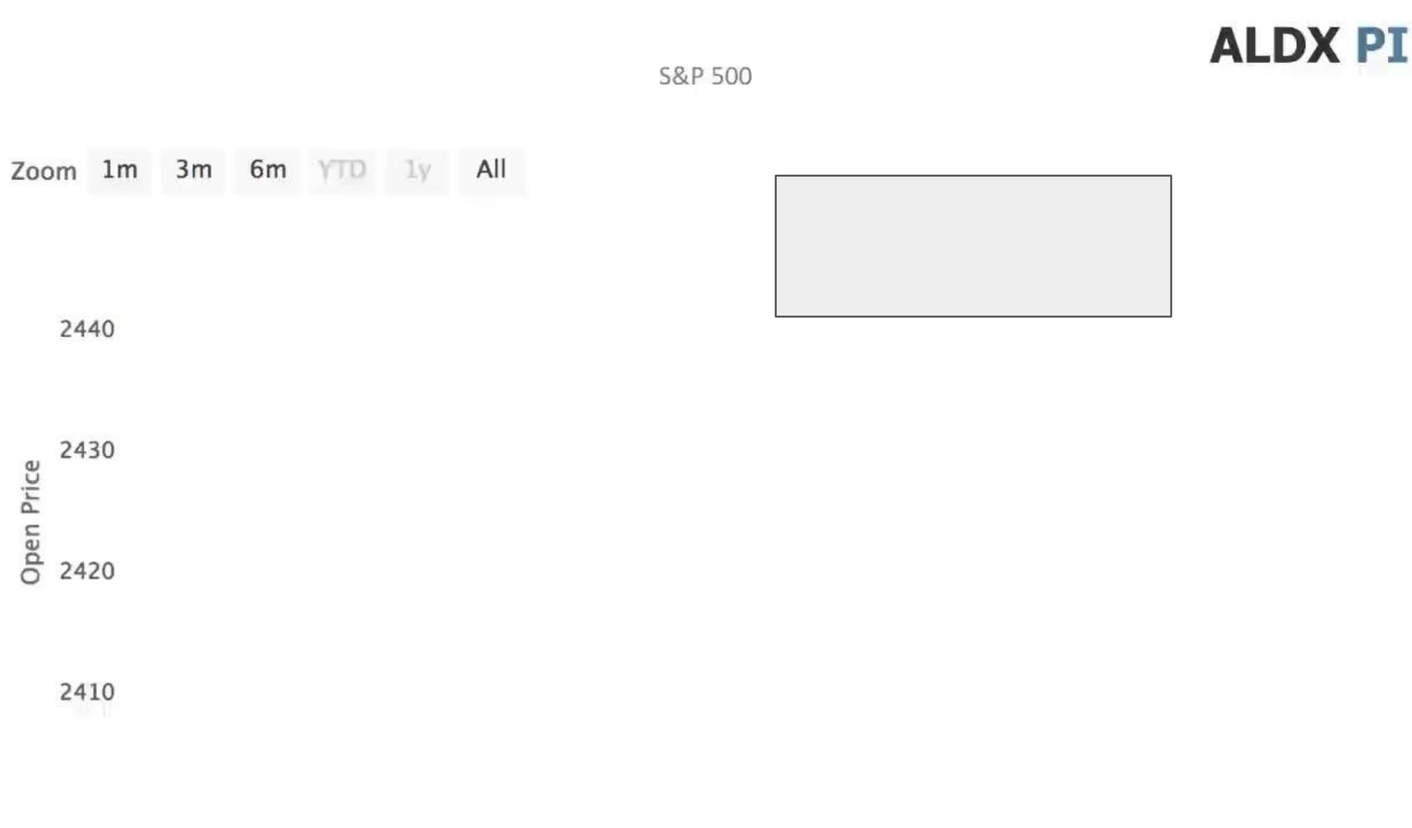

Key Points

- AlgoDynamix leverages machine learning algorithms to generate Flags

- Up and down Flags indicate buying and selling activity likely to influence the price of assets

- Flags are specific to a certain asset - each asset has a different flag history.

Part 2:
Existing R42 Equities Fund Strategy

Existing Portfolio Strategy

Select Assets

Choose the appropriate assets for two distinct portfolios.

Portfolios Composition

Use these assets to create two portfolios:

- Aggressive Portfolio
- Low Risk Portfolio

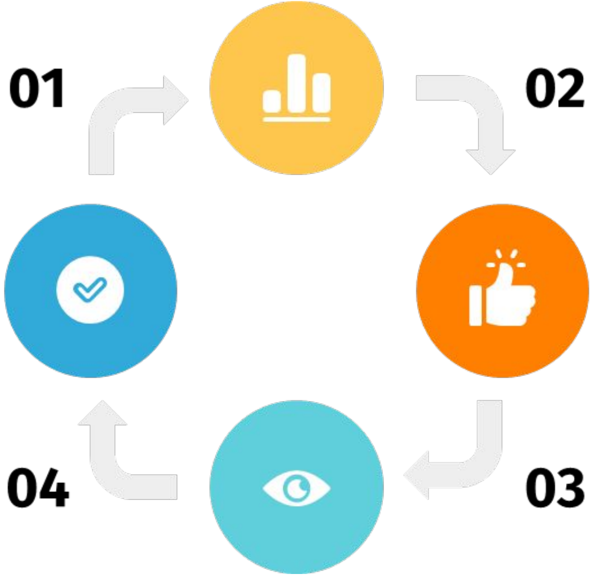

Optimize Investment

Use TPRs (calculated metric) to invest in the S&P 500 sectors with the highest potential returns.

Switching Portfolios

Utilize “Triple Up Flags” to switch between the aggressive and low risk portfolios based on market conditions.

Portfolio has protective and risky strategies

R42

Balanced Portfolio

- Mix of **Equities** and **Fixed Income**
- Sectors with positive *Tipping Point Ratios (TPRs)*.
- Suited to weather downturns

All Equities Portfolio

- **100% Equities**: sector ETFs from the S&P 500.
- Allocates based on the *Tipping Point Ratios (TPRs)* of each sector
- Higher risk higher reward

We pick S&P sector ETFs with the highest Tipping Point Ratios: Method to evaluate investments by comparing the potential good outcomes (returns above a target) versus the potential bad outcomes (returns below a target).

What is Tipping Point Ratio (TPR) ?

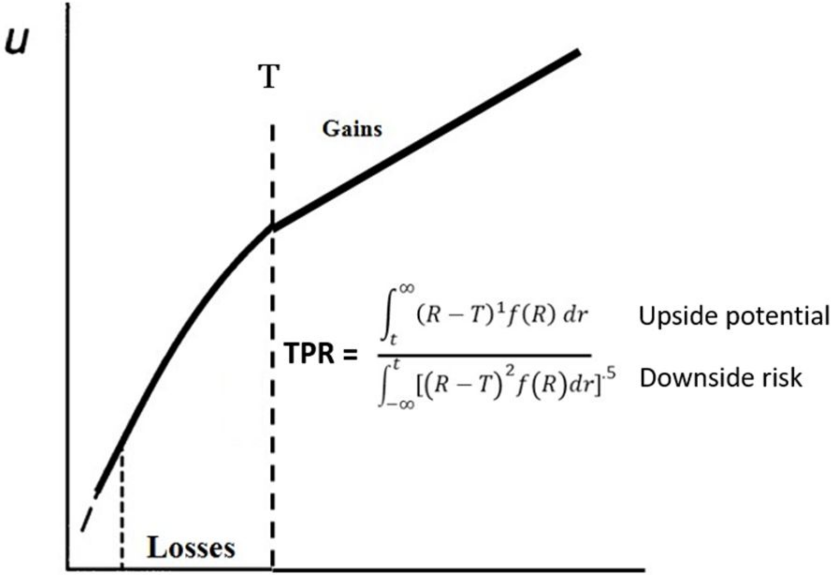

$$TPR = \frac{\int_t^{\infty} (R-T)^1 f(R)\, dr}{\left[\int_{-\infty}^{t} (R-T)^2 f(R) dr\right]^{.5}}$$

Higher TPR means more favorable outcomes compared to unfavorable ones.

Triple Up Flag

Triggered on:

- **SPY Flag**: tracks the performance of 500 largest U.S. firms.
- **QQQ Flag**: 100 of the largest non-financial companies listed on the NASDAQ stock exchange.
- **Decreasing VIX Volatility Index**: Measures how much stock prices are expected to fluctuate.

Key Points

Two Strategy Approach:

- **Balanced Portfolio**: Mixes equities and fixed income
- **All Equities Portfolio**: Focuses only on sectors ETFs from S&P 500.

Tipping Point Ratio (TPR):

- Select assets based on TPR to maximize upside potential while minimizing downside risk.
- Higher TPR = Better balance between gains and losses.

Triple Up Flag:

- Uses three key Flags to determine when to **switch** portfolios.
- Ensures we capitalize on market opportunities.

Part 3:
How we improved the strategy

How to improve R42 equities strategy?

Improving criteria to switch portfolios

"Triple Up Flags" → Bayesian Predictions

Our Improvements

- Predictive model instead of the “Triple Up Flags”.
- Should take into account as many Flags as possible - not just SPY and QQQ
- Several different trials and types of models to interpret flag data
- Eventually settled on Bayesian hierarchical model
  - Handles low frequency data
  - Represents group-level effects
  - Quantifies uncertainty about predictions

Flags don't grow on trees

- Flags are infrequent yet significant events
- Many model candidates had trouble fitting discrete rare events
- Bayesian model only updates parameters when relevant data is observed
- Particularly effective in modeling anomaly events.

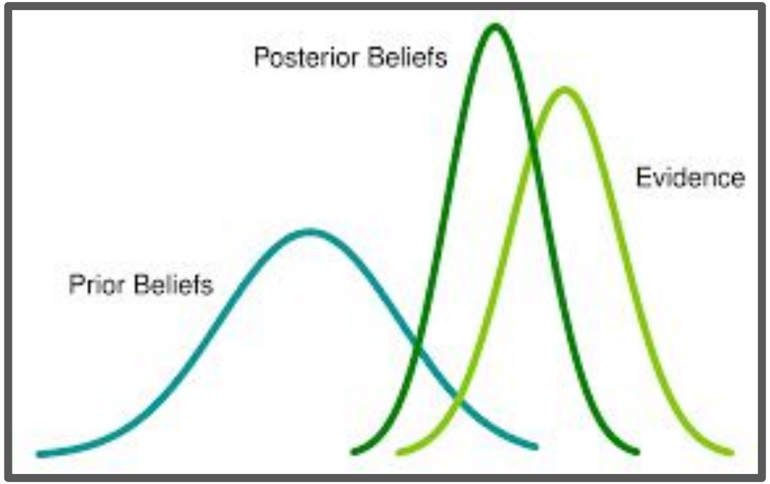

How long do Flags affect price?

- Engineered "Flag in the past n days" lookback features to learn how long a Flag is relevant for.
- Effect lengths for the same Flag generally have similar effects.
- Hierarchical model allows us to prevent overfitting

Information about Up Flags for SPY

Information about the different effect length features of SPY Up Flags

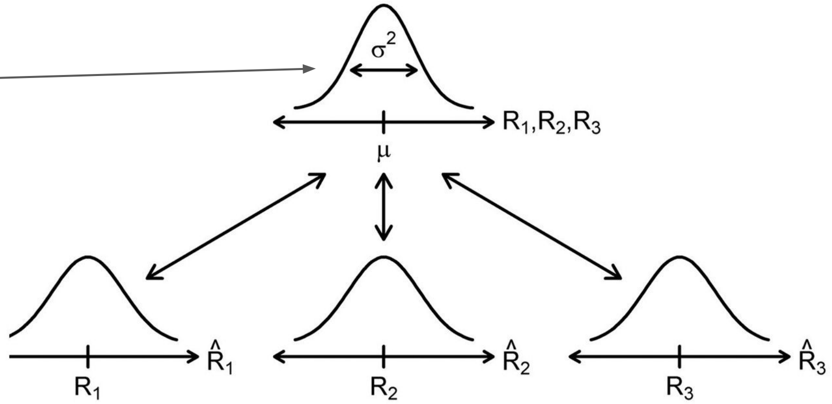

Aggregating info across assets

- We would like to consider flags on more than a single asset
- Predictive strength differs between Up Flags on different assets
- Up Flags have a different effect than Down Flags
- Hierarchy allows us to model the effects of different categories

  Flags fall into using a natural hierarchy:
    - 1) Up Flag or Down Flag
    - 2) Flag Asset (SPY QQQ ect)
    - 3) Effect duration (1-10 days)

Strengths of Bayesian Hierarchical Model

- Remembers rare events
- Shares relevant information across assets and effect duration
- **The model can "borrow strength" from Flags with more data to improve estimates for Flags with less data.**
- Bonus: Gives estimate of uncertainty about prediction

Promising predictive strength

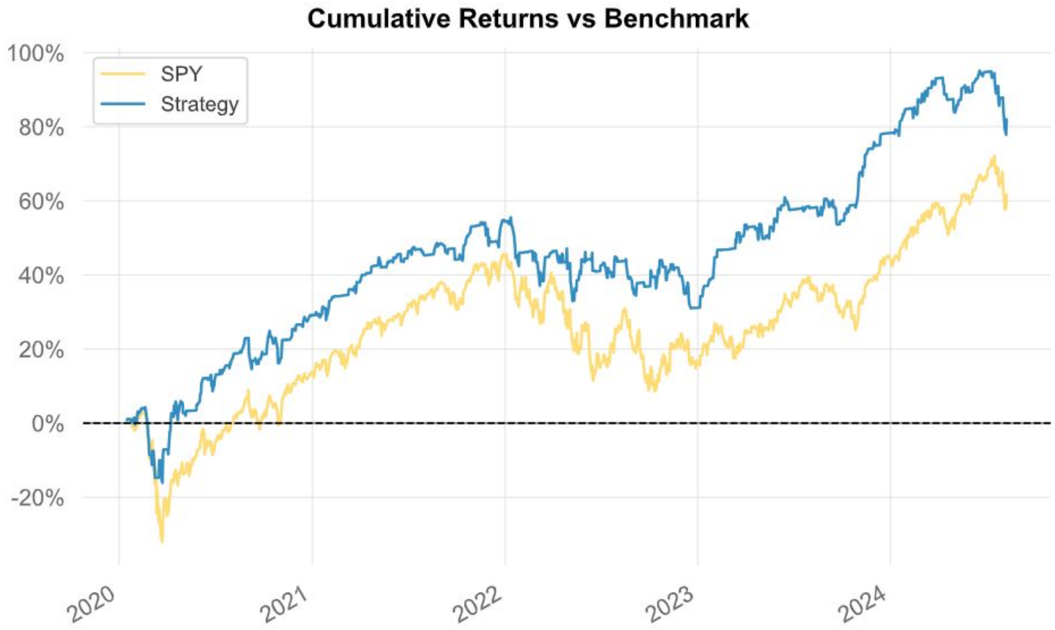

Key Performance Metrics

Metric	SPY	Strategy
Risk-Free Rate	0.0%	0.0%
Time in Market	100.0%	100.0%
Cumulative Return	61.69%	81.97%
CAGR﹪	7.54%	9.48%
Sharpe	0.59	0.97
Prob. Sharpe Ratio	89.52%	98.11%
Smart Sharpe	0.59	0.97
Sortino	0.83	1.43
Smart Sortino	0.82	1.43
Sortino/√2	0.59	1.01
Smart Sortino/√2	0.58	1.01
Omega	1.28	1.28

An example strategy selectively holding SPY or a fixed-income asset based on an ensemble of Bayesian hierarchical models.

R42 equities fund strategy with Bayesian Predictions

Improving criteria to Switch Portfolios

"Triple Up Flags"
Total Return: 34.8%
Sharpe Ratio: 0.8

→

Bayesian Predictions
Total Return: 38.16%
Sharpe Ratio: 0.86

*From 2021-07-06 to 2024-08-08

R42 Equities Fund Strategy with Bayesian Model

Existing R42 TPR strategy using Bayesian model instead of Triple Up Flag.

Tradeoffs of a Bayesian Model

- Issue: Trained model is different each time
  - Due to need to sample from probability distribution
- Our solution: use ensemble of models
- Issue: How do we trade optimally using this info
- Our solution: Research reinforcement learning

Key points

- Replace Triple Up Flag with predictive model
- A Bayesian hierarchical model has several advantages over other time series models for this application
- Can be thought of aggregating several discrete flags into one continuous higher-quality flag
- Increased returns without increasing risk

Part 4:
Future work with Reinforcement Learning

Future Work

How to interpret Bayesian predictions

Use prediction alone to switch portfolios
“static”

→

Use RL to incorporate more context into interpretation
“dynamic”

How to use continuous-Flags?

- Is switching mechanism a **learnable** mechanism?

- Reinforcement learning can serve as a powerful tool for developing a switching function that **dynamically** adjusts based on simulations of returns and the continuous-time Flags.

- The learning is achieved by trying to construct an approximation of a function
  that determines the optimal actions based on the **current state of the market**.

- The predictions use other indices such as NASDAQ that’s why we need a complex simulation model for capturing the relationships

RL-based trading: principle

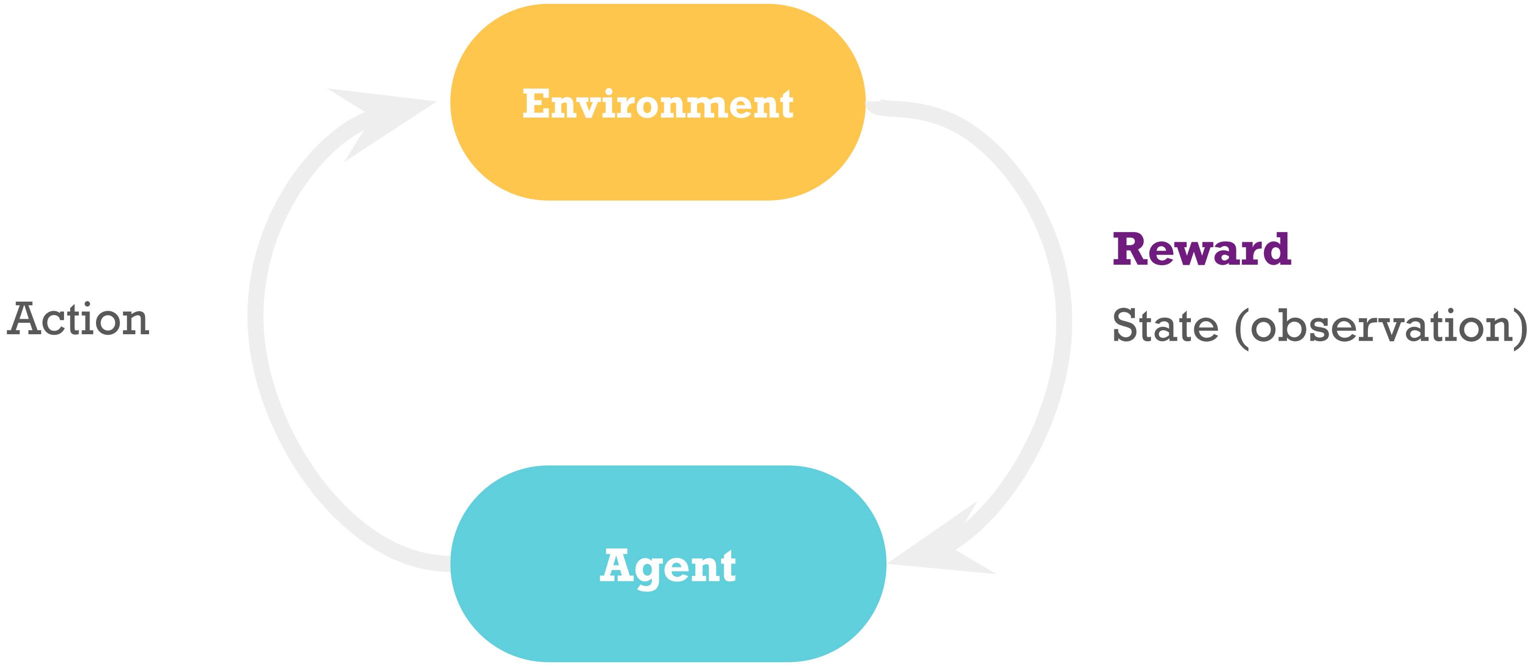

From prices to neural networks

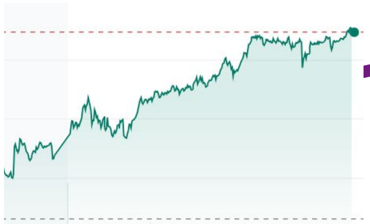

1 Fit statistical model for the returns

2 Simulate different scenarios for the returns

3 Generate Flags and bayesian predictions

4 Train the AI agent

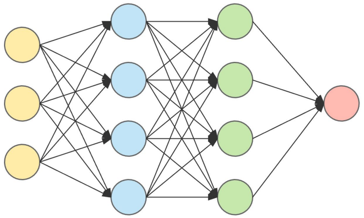

Simulation (steps 1–3) | **Training** (step 4)

Simulation of the returns

- Use **Beta distribution** for the returns of each index

  Offers flexibility for choosing the mean, variance, skewness and kurtosis

- Use **Gaussian Copula** to model the relationship between the indices

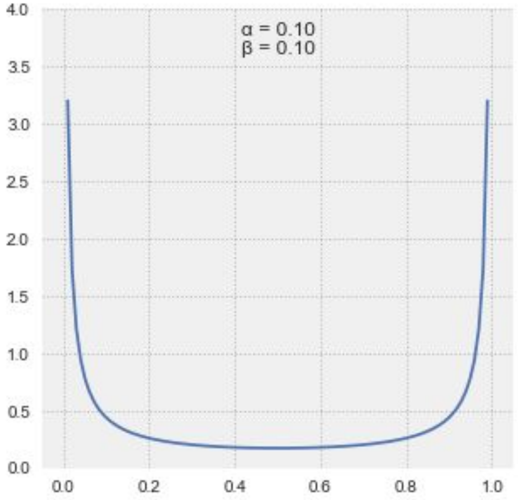

Early concept results

Resource intensive approach: Using our current and local resources, training an AI agent to trade on 1-year period takes approx **4 months.**

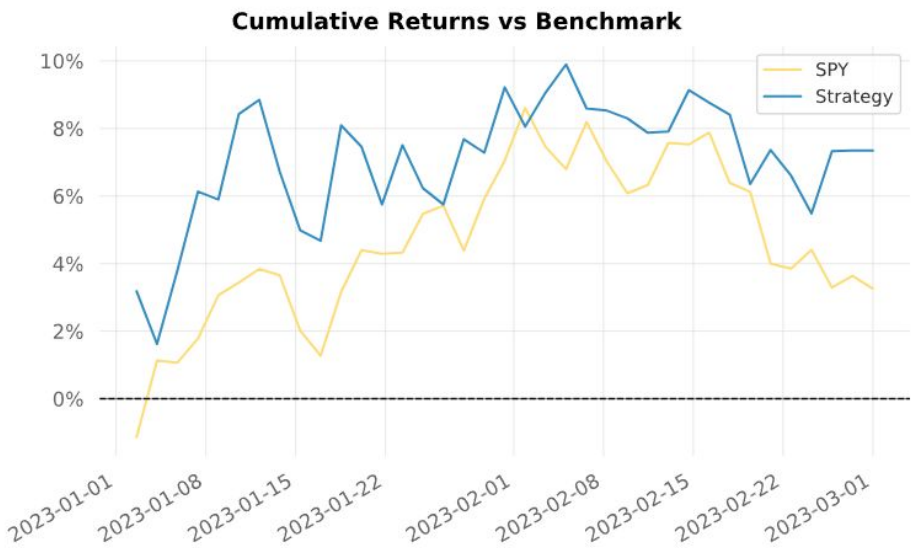

Key Performance Metrics

Metric	SPY	Strategy
Risk-Free Rate	0.0%	0.0%
Time in Market	100.0%	98.0%
Cumulative Return	3.26%	7.35%
CAGR﹪	15.22%	36.8%
Sharpe	1.4	2.25
Prob. Sharpe Ratio	70.21%	81.3%
Smart Sharpe	1.11	1.79
Sortino	2.17	4.18
Smart Sortino	1.73	3.33
Sortino/√2	1.54	2.95
Smart Sortino/√2	1.22	2.35
Omega	1.44	1.44

Conclusion

- We improved the existing R42 portfolio strategy by leveraging a Bayesian model to interpret Flags to determine when we deploy capital.

- Compared to the S&P500's (SPY) return of 27.88%, the AlgoDynamix portfolio with Bayesian prediction generated a 38.16% return between 2021-07-06 to 2024-08-08.

- Groundwork laid for reinforcement learning, AlgoDynamix has the ability to dynamically improve its accuracy in switching portfolios at the optimal moment

Acknowledgements

A special thank you to our mentors Dr. **Jeremy Sosabowski**, Dr. **Frank Sortino**, and Dr. **Ronjon Nag** for making this fellowship possible and for their unwavering guidance throughout the project.

We want to express our appreciation to all members of the **R42 Group** for their invaluable weekly feedback.

Dr. Ronjon Nag
Founder and President
R42 Group

Dr. Frank Sortino
Director
Pension Research Institute

Dr. Jeremy Sasobowski
Executive Director
AlgoDynamix

Contact Us

Carlyle (Lyle) Morris

Email: carlylemorris@tamu.edu

https://www.linkedin.com/in/carlylemorris/

Brian Lee

Email: bl684@cornell.edu

https://www.linkedin.com/in/brian-byungkwan-lee/

Adi Lankipalle

Email: adilanki@wharton.upenn.edu

https://www.linkedin.com/in/adi-lankipalle/

Ektor Avlonitis

Email: ektor.avlonitis.23@ucl.ac.uk

https://www.linkedin.com/in/ektor-avlonitis/

Iheb Hajji

Email: iheb.hajji@hec.ca

https://www.linkedin.com/in/ihebhajji/