Predicting NBA Mispricings: CS1090b MS4

Ethan Aidam, Aaron Mcdonald, & Sam Burgess

Background

- Data comes from Kalshi API and ESPN API
- What is Kalshi?
    - Prediction Market
    - Allows people to “bet” on real world events (sports, politics, etc.)
    - Bet by “buying” or “selling” outcomes
        - Price should be reflective of market sentiment
- Basketball!
    - Sport
    - Tradable on Kalshi
    - Lebron James
- ESPN
    - Major statistics data provider for NBA
    - Offers metrics like win probabilities for each game
    - Win probabilities often fluctuate with plays that have great influence on the outcome of the game

Motivation

Human Emotion & Big Plays

Big plays generate excitement, impacting Kalshi market activity.

Market Irrationality

Markets may not always react rationally to high-intensity game moments.

Identifying Mispricings

Determining where markets react appropriately vs. where events expose Kalshi mispricings.

Problem Statement

How do prediction market prices adjust to NBA events with large impacts on win probability?

To what extent do these adjustments exhibit systematic overreactions that lead to persistent mispricing?

Data & key EDA findings

Dataset Description

Alignment & Cleaning

- Aligned Kalshi market candles with ESPN game data
- Built 20 minute candle sequences around win-prob shocks (5% change in win-prob)
- Filtered for sufficient coverage and quality

Modeling Dataset

Primary: lstm_dataset_df

Analysis: sequence_analysis_df

12,422 Rows

Engineered Features:

- Target return & volatility
- Coverage & direction
- Alignment diagnostics

Dataset Benefits

- Provides full event-sequences for modeling
- Preserves temporal behavior and analytics around shocks

Distribution of Mispricing Events

We considered restricting the events in our data set to only the end of the quarter or Q4 where win probabilities could potentially swing the most but EDA led us to reject this hypothesis.

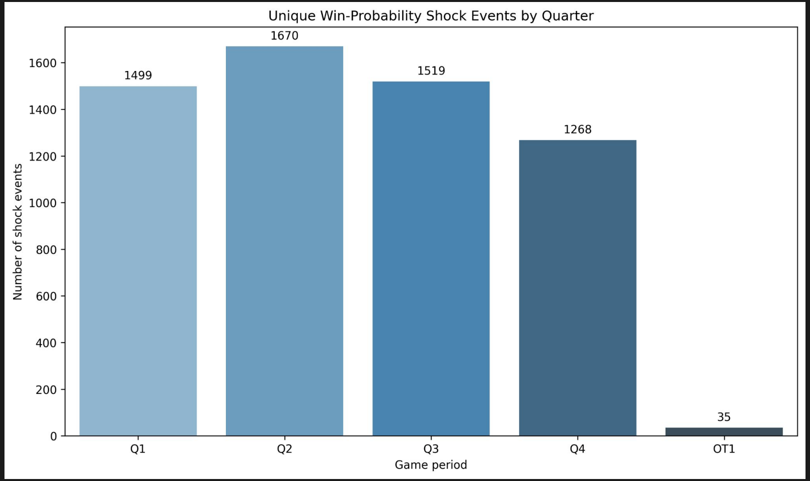

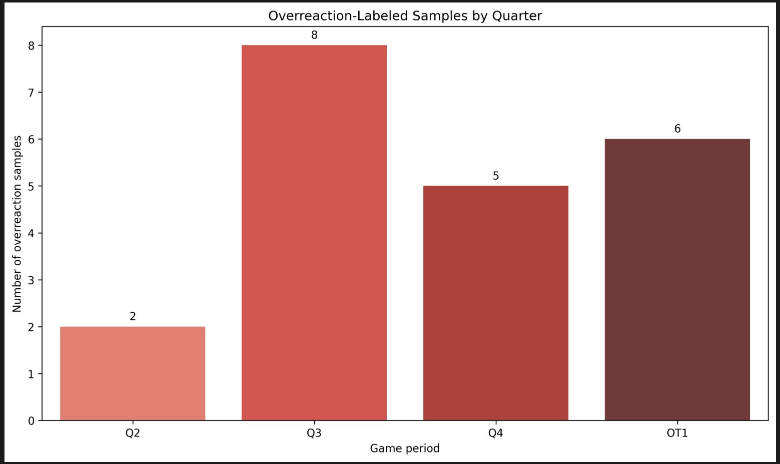

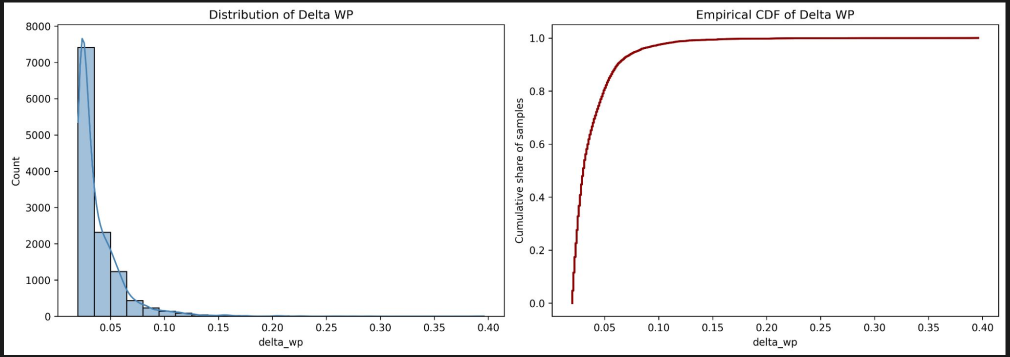

Example Mispricing Event: LAL @ HOU | 2026-04-24

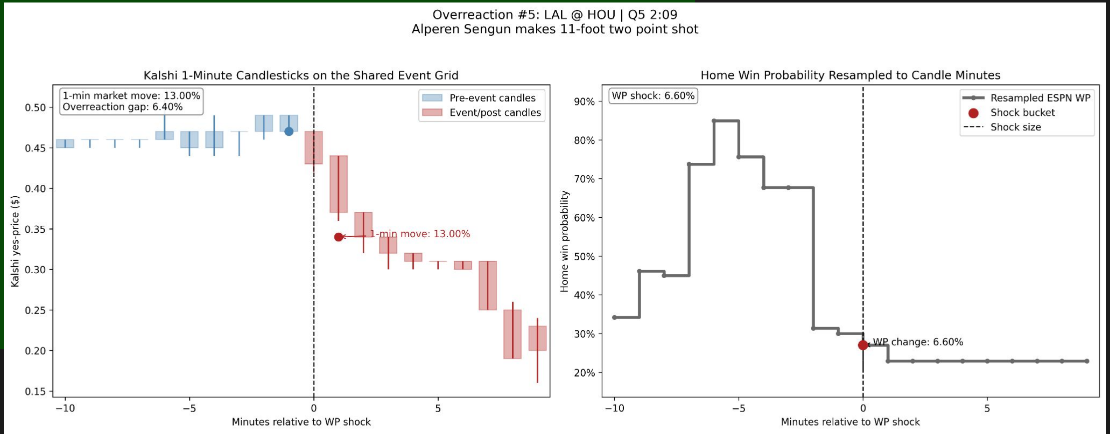

EVENT DESCRIPTION

Time: 2:09 (Q5)

Play: Alperen Sengun makes 11-foot shot

Final Outcome: LAL Win

MARKET REACTION

WP Shock: **6.60%**

Overreaction Gap: **6.40%**

13.00%

1-Min Market Move

Modeling approach and Training Details

Our project goal is to evaluate whether NBA win-probability shocks and recent Kalshi market behavior can predict short-horizon market reactions. The models are designed to answer this in increasing levels of complexity using one of two modeling frameworks: an interpretable **multiclass direction classification model**, and a **short-term price prediction regression model**. This results in two potential sequence model architectures, LSTM and a Time Series Transformer, to predict both the direction and price 1-minute after a win-prob shock event.

Task	Model	Train BS	Eval BS	LR	WD	Max Ep	Ep Run	Patience	Time (s)
1-min direction	LSTM	32	64	1e-3	0	60	50	10	53.59
1-min direction	Transformer	32	64	5e-4	1e-4	60	60	12	80.71
1-min price	LSTM	32	64	1e-3	0	60	60	10	65.77
1-min price	Transformer	32	64	5e-4	1e-4	60	31	12	41.36

Results and Interpretation: Classification

Model	Results
Logistic Regression	- Accuracy: 0.4440 - Balanced accuracy: 0.4272
HistGradientBoosting Classifier	- Accuracy: 0.5073 - Balanced accuracy: 0.4996
LSTM Classifier	- Accuracy: 0.8209 - Balanced accuracy: 0.8217
Transformer Classifier	- Accuracy: 0.7975 - Balanced accuracy: 0.7966

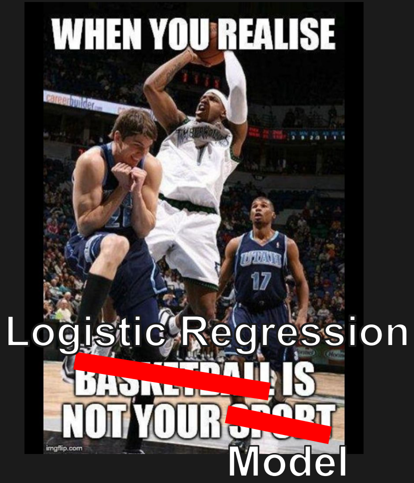

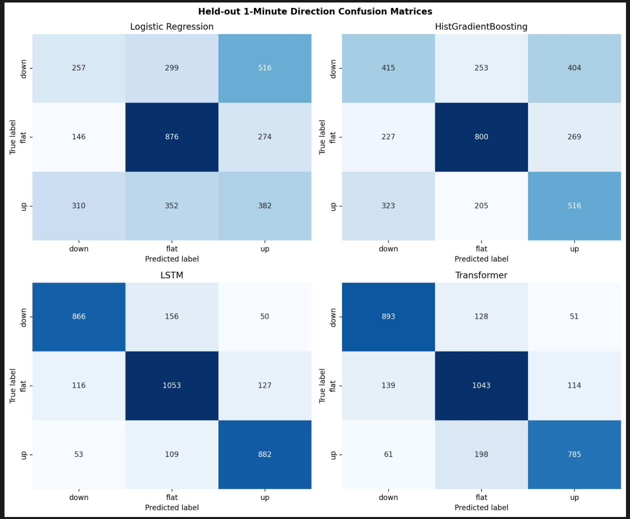

Results and Interpretation: Time Series Regression

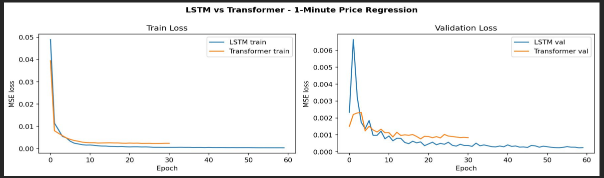

LSTM vs Transformer - 1-Minute Price Regression

The regression models were evaluated on their ability to predict the next minute's price based on 20-step historical sequences.

- The **LSTM model** shows steady convergence in both training and validation loss, suggesting robust learning of temporal dependencies.
- The **Transformer model** achieves a slightly lower R^2 of 0.9802, indicating diminished performance in capturing complex temporal relationships.
- Low MAE and RMSE across both models confirm high precision in price forecasting.

Case Study Analysis

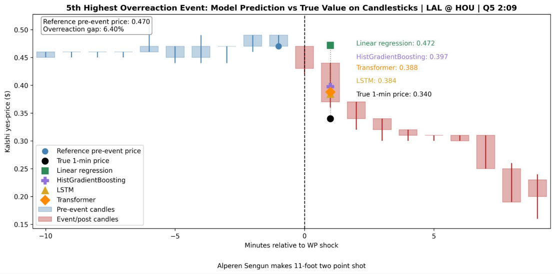

Model Metrics Summary

Model	MAE	MSE	RMSE	R^2
Linear Regression	0.0227	0.0012	0.0340	0.9703
Hist Gradient Boosting	0.0278	0.0014	0.0379	0.9630
LSTM	0.0111	0.0002	0.0154	0.9939
Transformer	0.0214	0.0008	0.0277	0.9802

Results and Interpretation: Time Series Regression

Case Study Analysis

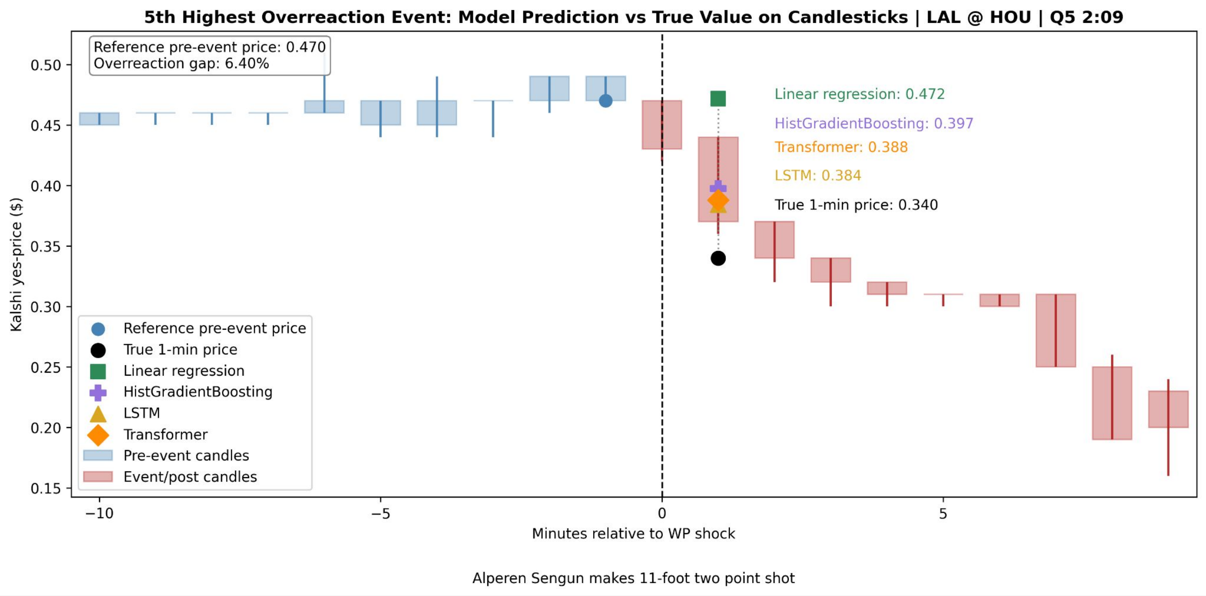

Conclusions and Future Work

Classification

Hist Gradient Boosting outperforms logistic regression in classifying whether the Kalshi market will react appropriately.

Forecasting

However, both LSTM and Transformer models perform better and can very accurately predict the future direction of the Kalshi market based on influential events.

Other Applications

Similar modeling could potentially hold applications to financial markets, forecasting stock prices using news as market shocks.

Thank you