Liquidity Risk in Prediction Markets: A Polymarket Case Study

CHLOE SHAO[1], KATHERINE LIANG[2]

[1]: High School Intern from Camellia Venture Capital High School Summer Intern Program, Stony Brook, NY
[2]: Corresponding Author

ABSTRACT

This paper examines liquidity risk and market dynamics in online prediction markets through a case study of Polymarket, the world's largest prediction market. We analyze two high-volume events—"Israel x Hamas Ceasefire Before August?" and "Which Company Has Best AI Model End of July?"—to investigate correlations, liquidity, and external event impacts. Using automated hourly data collection over a two-week period, we track key metrics including yes/no share prices, bid–ask spread, liquidity, and trading volume. Our findings reveal that high-probability markets exhibit greater volatility, while low-probability markets remain more stable. Price fluctuations strongly correlate with bid–ask spread and short-term trading volume spikes, often triggered by external event developments. These results show that prediction markets can track public opinion in real time and help forecast events for financial and policy decisions.

Keywords: *Prediction markets, Polymarket, Liquidity risk, Market volatility, Geopolitical forecasting, AI competition.*

Table of Contents

I. Introduction 2
II. Background 2
2.1 Differences Between Traditional Betting and Online Prediction Markets 2
2.2 Key Statistics 3
2.3 Potential 3
III. Methodology 4
IV. Results 4
4.1 Israel x Hamas Ceasefire Market 4
4.2 Best AI Model Market 7
V. Discussion 10
5.1 General Patterns in the Markets 10
5.2 Correlations 12
5.3 Comparisons Between Markets 12
5.4 Patterns in the Israel x Hamas Ceasefire Market 13
5.5 Patterns in the Best AI Model Market 13
VI. Conclusion 14

I. Introduction

As society becomes increasingly digital, there is a quickly growing phenomenon among investors: online prediction markets. These markets are taking the place of traditional betting, with several key differences between the two. This paper analyzes the effectiveness of prediction markets through a case study of two high-volume macro-political markets on Polymarket, the largest prediction market platform. By connecting these two high-volume markets to external events, we aim to provide insight on how prediction markets change and adapt in real time to guide investors and businesses that may be able to use prediction markets for financial decisions. The review of the two Polymarket events covers two broad categories: geopolitical events and technology competition.

II. Background

2.1 Differences Between Traditional Betting and Online Prediction Markets

With the shift from traditional betting to online prediction markets, there are a few key differences to keep in mind.

Traditional betting is typically done for entertainment, usually in the forms of wagers or bookmakers. They are often more unofficial, with odds set by either the bettors or the bookmaker. Oftentimes, the wagers are exactly 1:1, and don't always reflect the odds of winning. They also are sometimes less secure and upfront with rules and procedures, unlike prediction markets.

Prediction markets, by contrast, are structured to function like financial markets. Participants buy and sell shares in possible outcomes. A "Yes" share priced at $0.65 suggests a 65% likelihood that the event will occur. These markets are dynamic, with prices changing based on supply and demand, thereby aggregating diverse perspectives into a probability-based forecast. As users purchase more and more of a share, the odds of that rise along with the price. Additionally, prediction markets are run on software like blockchain technology, making them much more secure as users' data is saved into the larger chain, ensuring that the data won't be lost.

The purposes also diverge, as traditional betting is typically recreational, while prediction markets aim to collect and reflect the most accurate public sentiment or data about an event. For this reason, prediction markets are often used in research, forecasting, and policy analysis. They are considered efficient because participants are financially incentivized to be accurate, not just lucky.

Another difference is legal and regulatory differences. Traditional betting is strictly regulated and often limited to certain jurisdictions. Prediction markets, particularly decentralized ones like Polymarket, operate on blockchain technology and exist in a more ambiguous

regulatory space. They may be framed as informational tools rather than gambling platforms, but the legal space is much more ambiguous.

In summary, while both systems involve financial speculation on uncertain events, traditional betting centers on fixed odds and entertainment, whereas prediction markets harness collective intelligence to forecast outcomes more accurately.

2.2 Key Statistics

In prediction markets like Polymarket, investors should pay close attention to key statistics, especially liquidity, which is crucial for smooth trading and accurate price discovery. Liquidity can be measured through several metrics that show market depth, activity, and stability. One important measure is the bid–ask spread—the gap between the highest price a buyer will pay and the lowest price a seller will accept. A smaller spread means better liquidity and lower trading costs. On Polymarket, the "Buy Yes" and "Buy No" prices represent the ask and bid, and a narrower gap makes trading easier and more cost-effective.

Market depth refers to the available volume at different price levels, while liquidity pool size (in AMM-based markets) shows how much capital is locked in a market to facilitate trades. On Polymarket, this is often displayed in the "Liquidity" section of a market page, showing how much USDC is in the pool. A deeper pool means greater capacity to absorb large trades with minimal slippage.

Volume indicates how much has been traded within a given period (often 24 hours). This reflects market activity and interest. High volume suggests better price discovery and more reliable probabilities. On Polymarket, volume can be found on the market page under "Volume (24h)" or "All-Time Volume."

Slippage is the difference between the expected price of a trade and the price actually received due to insufficient liquidity. Polymarket provides estimated slippage when you input a hypothetical trade amount—this is shown in real-time before confirming a trade. High slippage means a trade could move the market significantly, making it more expensive or less predictable to execute.

2.3 Potential

Polymarket's structure as a decentralized prediction market with trading-esque mechanics makes it especially appealing to users interested in both prediction and finance. Dynamic price discovery is enabled first by the platform, with users capable of profiting from buying low and selling high as market sentiment shifts. This constitutes a more interactive trading experience than fixed-odds betting, where users are not simply betting on outcomes but instead speculating on perceived changes in probability. Second, Polymarket financially incentivizes early and accurate intuition, attracting informed participants. Early movers on new data or breaking news

can place positions at favorable prices, increasing the payout for researched predictions. This arrangement incentivizes users to stay knowledgeable and think critically, rather than rely on luck. Third, the site offers a portfolio-like experience, allowing users to monitor multiple markets at once. They can hedge positions, balance risk, and make decisions similar to financial investing. This level of strategic engagement adds depth and longevity to the user experience.

Polymarket also holds excellent potential for machine data analysis and harvesting. Because market data is openly accessible via blockchain records or APIs, developers and researchers can track user sentiment, prediction precision, and market trends in real time. This opens up possibilities for journalism, scholarly research, and even public policy tracking using the market shifts as an indicator of collective belief or concern.

III. Methodology

After conducting background research on Polymarket, we selected two markets to track and analyze. The first market selected was in the Geopolitics category, where we tracked “Israel x Hamas Ceasefire Before August?” and the second was in the Technology category, where we tracked “Which Company Has Best AI Model End of July?”. Both markets were selected for current relevance, as well as the high volume of trades (around $4 million each at the time of selection, which was July 18th). The geopolitics market contained all trades in just one market, while the tech market spanned thirteen markets, each corresponding to the likelihood of a specific company—like Google—having the leading AI model by the end of July. This leading AI model was defined by the company with the highest arena score based off the Chatbot Arena LLM Leaderboard on July 31st, 12:00 p.m. EST.

Then, we developed an automated data collection script in Python to gather market data at hourly intervals over a two-week period ending in late July, when both markets concluded. The collected variables included the yes share price, no share price, bid–ask spread, total trading volume, 24-hour trading volume, and liquidity. Data processing was performed in Python to generate time-series plots and calculate key descriptive statistics—mean, median, standard deviation, and range—which serve as the basis for analyzing market fluctuations.

IV. Results

4.1 Israel x Hamas Ceasefire Market

We collected data on the event “Israel x Hamas Ceasefire Before August?” for a two-week period, starting on July 19th and ending on July 31st, which resulted in 152 data points. This data spanned six statistics: yes share price, no share price, bid–ask spread, total trading volume, 24-hour trading volume, and liquidity. Below is a table displaying several measurements for each statistic.

Statistic	**Yes Price**	**No Price**	**Bid-Ask Spread**	**Volume**	**24-Hour Volume**	**Liquidity**
Mean	0.124	0.876	0.0064	9323143.72	554913.89	106765.62
Median	0.0365	0.9635	0.003	9842073.29	295644.19	105712.47
Standard Deviation	0.1491	0.1491	0.0062	1261455.97	400983.39	35422.85
Minimum	0.001	0.53	0.001	4647863	131542.4	43960.84
Maximum	0.47	0.999	0.03	11150551.57	1307176.44	215074.14

Figure 1. "Israel x Hamas Ceasefire Before August?" Statistics

Additionally, we plotted time graphs of each statistic to visualize trends across the data collection period.

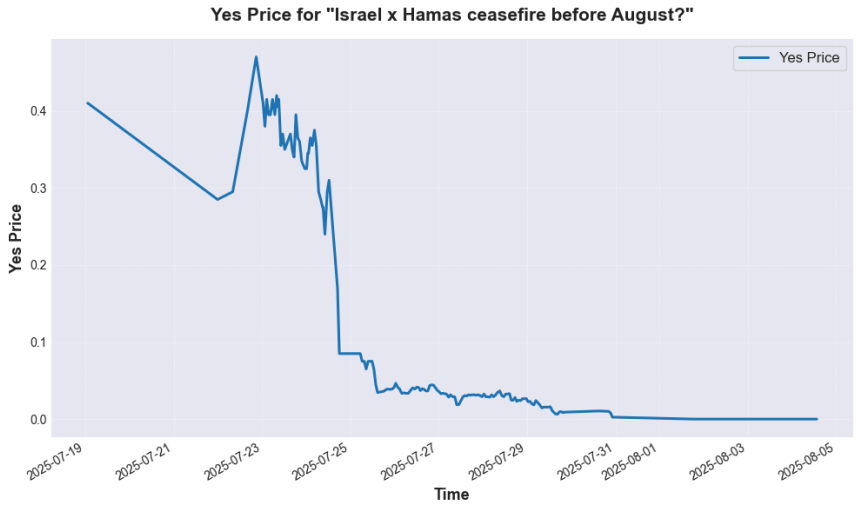

Figure 2. "Israel x Hamas Ceasefire Before August?" Yes Price Timeplot

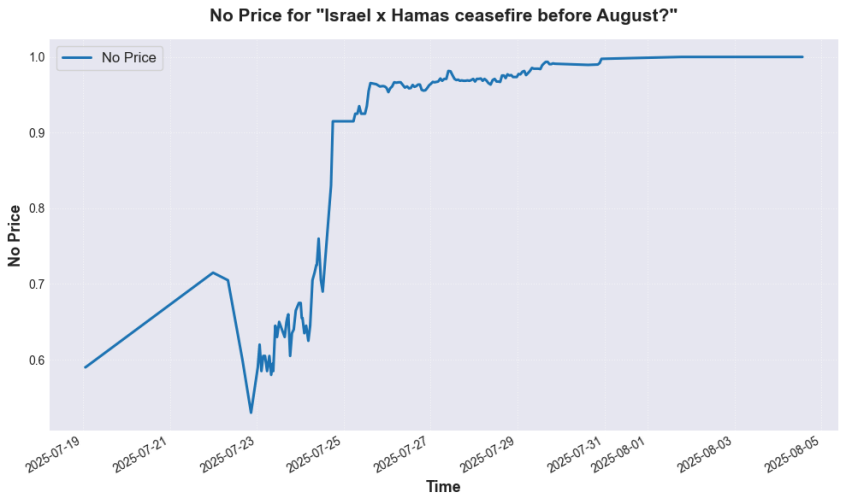

Figure 3. "Israel x Hamas Ceasefire Before August?" No Price Timeplot

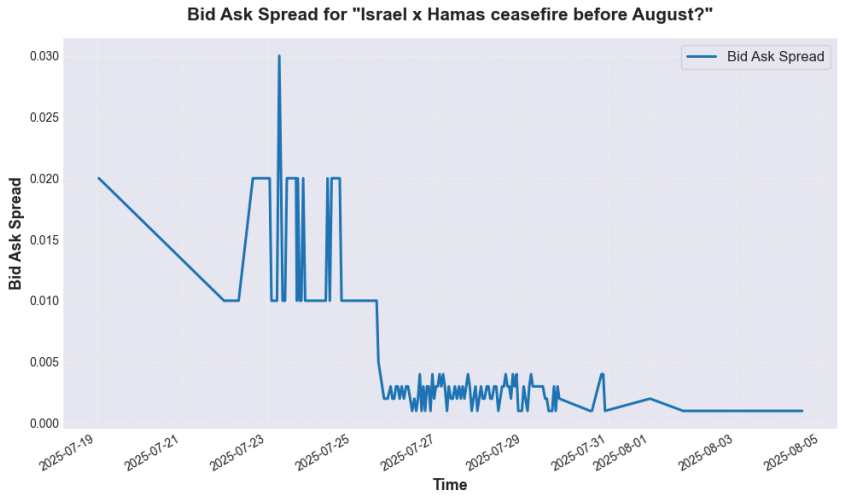

Figure 4. "Israel x Hamas Ceasefire Before August?" Bid-Ask Spread Timeplot

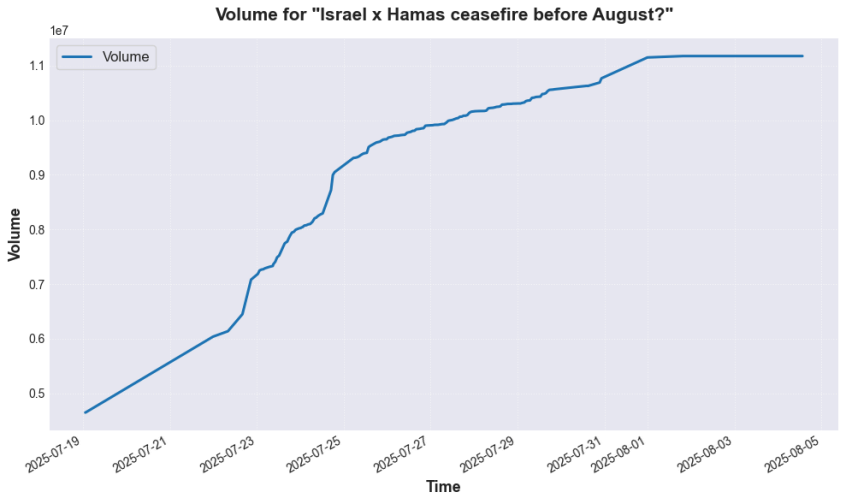

Figure 5. "Israel x Hamas Ceasefire Before August?" Volume Timeplot

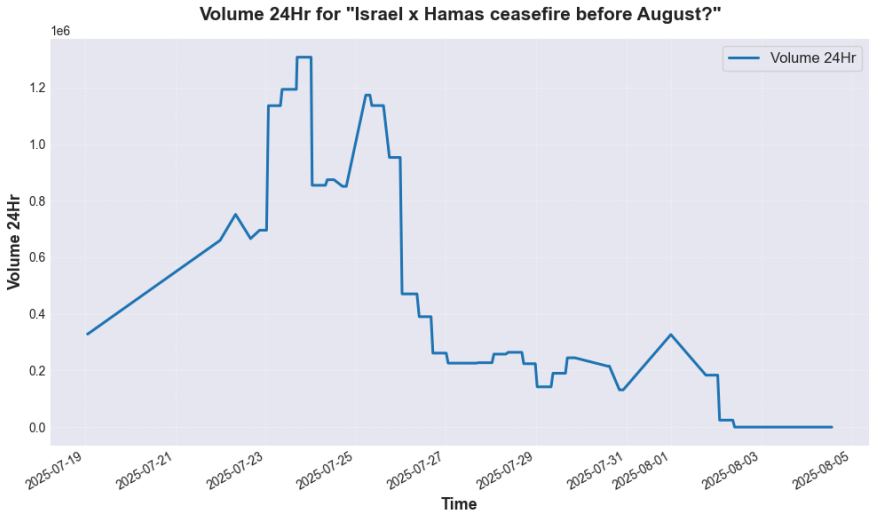

Figure 6. "Israel x Hamas Ceasefire Before August?" 24-Hour Volume Timeplot

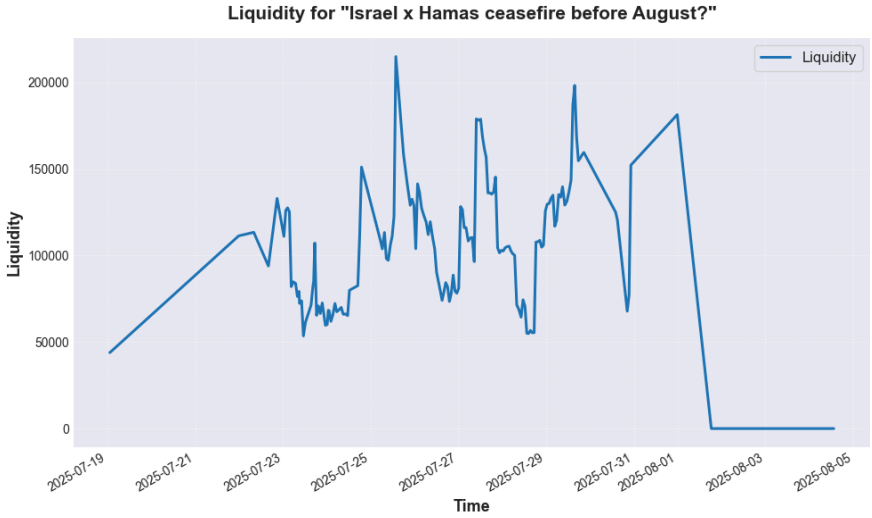

Figure 7. "Israel x Hamas Ceasefire Before August?" Liquidity Timeplot

4.2 Best AI Model Market

For our second event, "Which Company Has Best AI Model End of July?", we collected data for the same period of time, July 19th-31st. We are only displaying seven of the thirteen markets under this event due to six markets being placeholders with statistics of zero for every category. Below is a table comparing the six different metrics for the markets under the event.

Metric	*Statistic*	Google	Anthropic	xAI	Open AI	DeepSeek	Alibaba	Meta
Yes Price	*Mean*	0.9572	0.0011	0.0071	0.0326	0.001	0.0014	0.0008
	Standard Deviation	0.018	0.0008	0.0056	0.0144	0.0005	0.0006	0.0005
No Price	*Mean*	0.0428	0.9989	0.9929	0.9674	0.999	0.9986	0.9992
	Standard Deviation	0.018	0.0008	0.0056	0.0144	0.0005	0.0006	0.0005
Bid-Ask Spread	*Mean*	0.0055	0.0012	0.002	0.0031	0.001	0.0013	0.001
	Standard Deviation	0.0048	0.0005	0.0009	0.0016	0.0002	0.0005	0.0002
Volume	*Mean*	672173.26	674985.52	1186307.63	889734.68	525111.16	739746.28	544687.49
	Standard Deviation	75675.56	35424.9	76817.51	98244.54	46031.34	86438.77	24607.4
24-Hour Volume	*Mean*	29387.94	14218.29	30658.64	40239.85	18404.01	30386.22	8284.83
	Standard Deviation	13391.29	9496.65	11932.5	16984.56	13204.86	28040.39	7631.83
Liquidity	*Mean*	19614.24	78127.56	34859.93	29274	90537.05	71876.64	66628.06
	Standard Deviation	4745.71	32517.25	8489.06	5603.31	14208.67	23035.25	13446.92

Figure 8. "Which Company Has the Best AI Model by the End of July?" Statistics

We also graphed a timeplot of the data to visualize trends across the data collection period. Below are a few graphs that compare the different statistics (yes price, no price, bid-ask spread, volume, 24-hour volume, and liquidity) between the seven markets over time.

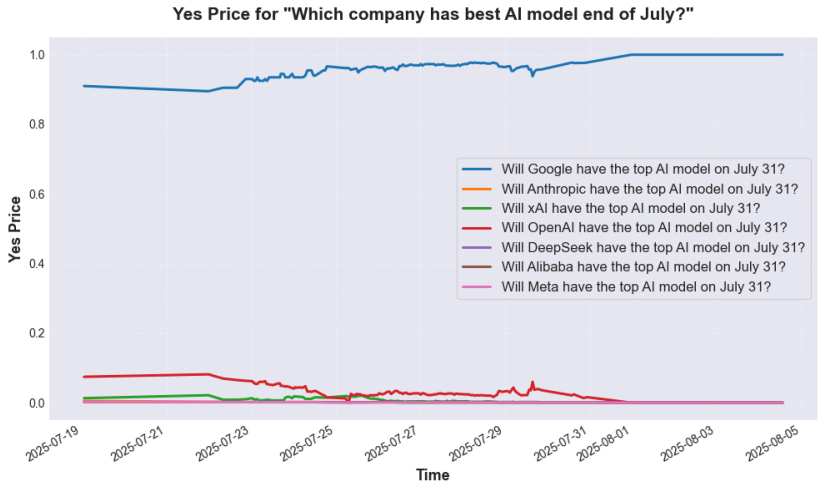

Figure 9. "Which Company Has the Best AI Model by the End of July?" Yes Price Timeplot

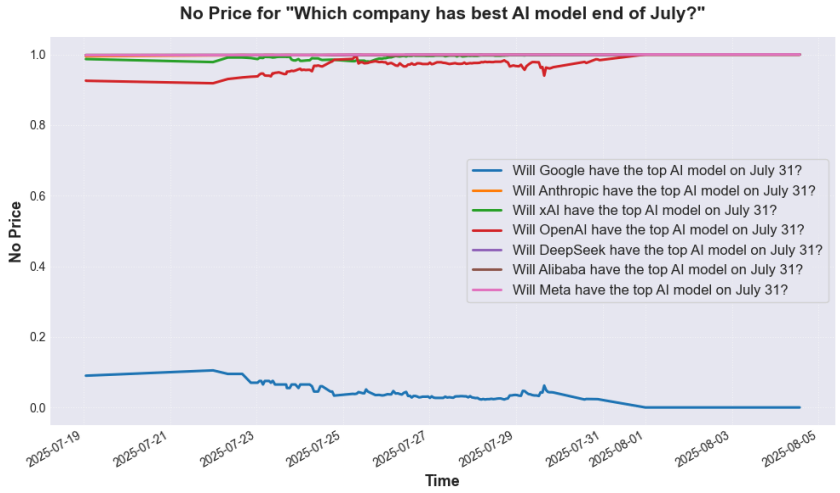

Figure 10. "Which Company Has the Best AI Model by the End of July?" No Price Timeplot

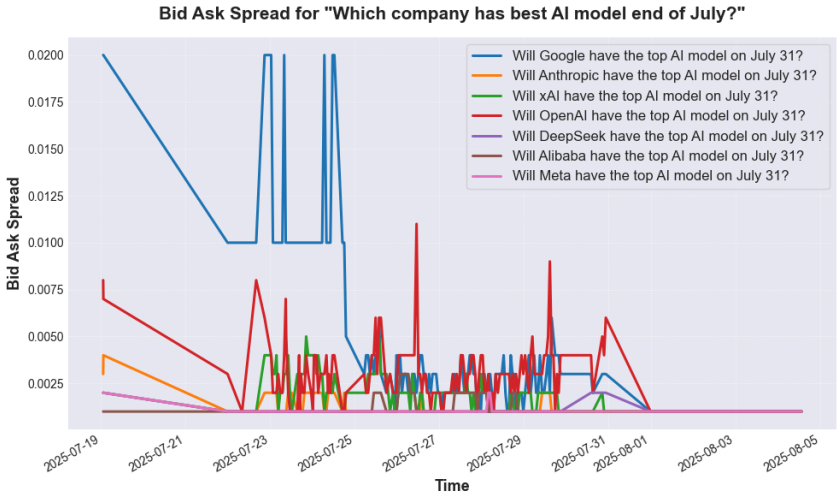

Figure 11. "Which Company Has the Best AI Model by the End of July?" Bid-Ask Spread Timeplot

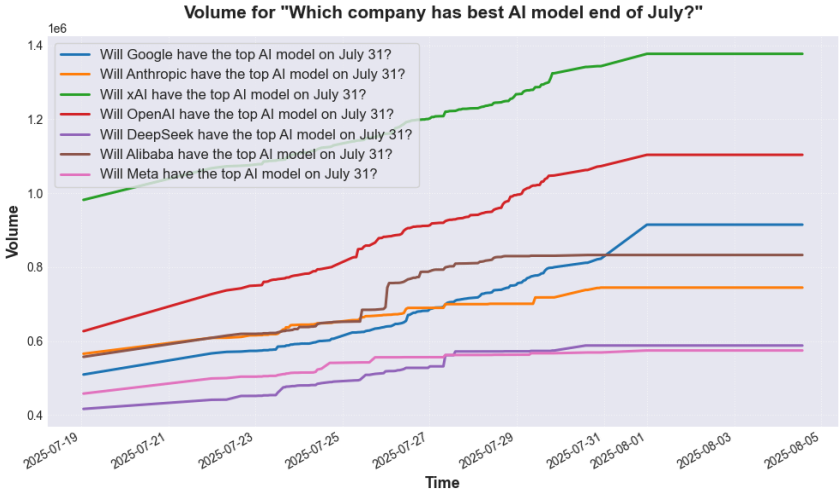

Figure 12. "Which Company Has the Best AI Model by the End of July?" Volume Timeplot

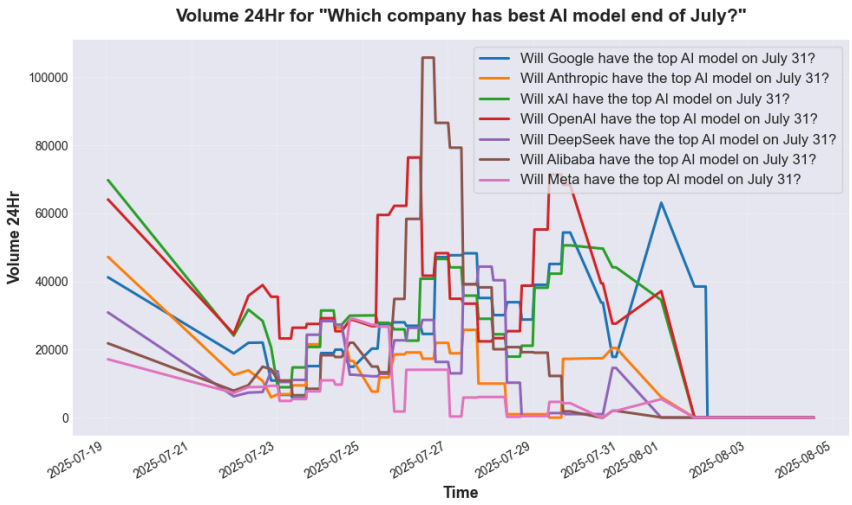

Figure 13. "Which Company Has the Best AI Model by the End of July?" 24-Hour Volume Timeplot

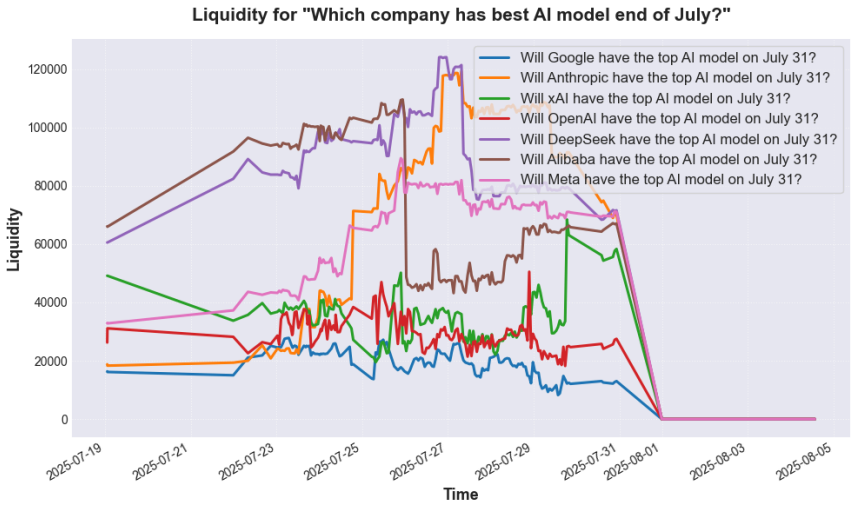

Figure 14. "Which Company Has the Best AI Model by the End of July?" Liquidity Timeplot

V. Discussion

5.1 General Patterns in the Markets

In all markets, there are times when several statistics may dramatically increase or decrease. We will refer to these points as "spikes" for the rest of the paper. These spikes seem to typically happen around moments of change, which will be discussed later on. Most prominently, they occur in markets with large volumes of trades, demonstrated by the two graphs below.

Figure 15. "Will xAI have the top AI model on July 31st?" Yes Price Timeplot

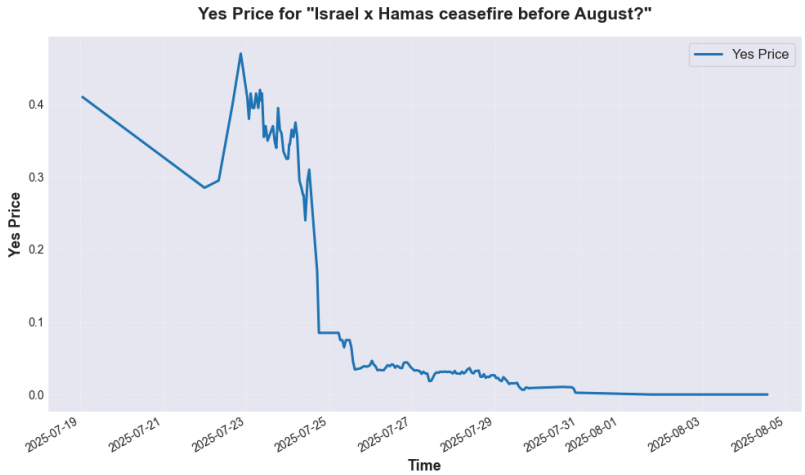

Figure 16. "Israel x Hamas Ceasefire Before August?" Yes Price Timeplot

Both have extremely large volumes of trades, as compared below.

Statistic	**Israel x Hamas Ceasefire**	**xAI Best AI Model**
Mean	9323143.7230	1186307.63
Median	9842073.29	1198608.97
Standard Deviation	1261455.97	76817.51
Minimum	4647863	981931.46
Maximum	11150551.57	1343571.09

Figure 17. Israel x Hamas Ceasefire Market vs xAI Best Model Market Volume Comparison

Additionally, volume appears to be the least variable statistic, as it is the only one that steadily increases across all markets. Variability is defined slightly unusually here, as volume may appear to vary quite a bit from the statistics, but generally increases at the most constant rate with very little sharp increase or decreases. In contrast, the most variable statistics seem to be the prices, liquidity, and bid ask spread, meaning they are more sensitive to sharp changes in the market.

5.2 Correlations

Multiple variables seem to be correlated. Most notably, yes share price and no share price are directly inversely correlated, which can be seen in both graphs and tables. However, large changes in both also seem to correlate to changes in bid-ask spread as well as volume 24hr. This can be seen in multiple graphs, most notably in the Israel x Hamas Ceasefire Market, where sharp spikes in yes share and no share prices directly correlated to spikes in bid-ask spread around July 23rd and 25th. This is likely because changes in yes share and no share prices affect the bid and ask prices, resulting in a larger gap between the two and a larger bid-ask spread. The same is shown as prices level out around July 27th, with a smaller bid-ask spread representing less of a disconnect between sellers and buyers.

Additionally, changes in yes share and no share prices are correlated with spikes in volume. This is likely because consumers are making larger amounts of trades as the climate changes, which results in a change in probability.

Between markets under the same event, prices appear to be relatively correlated, as there is only $1.00 allocated across all markets. In the AI Technology markets, most of this was in the Google market, as traders believed that market had the highest probability and Google was most likely to have the best AI model at the end of July. For the most part, volume, 24-hour volume, and liquidity also appear to be relatively positively correlated. This is possibly because many traders will invest in multiple markets - ie, a trader may believe that Google will have the best AI model and therefore buy “no” shares for the other companies.

5.3 Comparisons Between Markets

Overall, low probability events appeared to be more stable, with smaller ranges and standard deviations across all statistics except volume, 24-hour volume, and liquidity. High probability markets appeared to fluctuate more. This can be seen when comparing any of the other AI Technology markets to the Google market.

Between the two events, the Israel x Hamas Ceasefire market fluctuated much more, with a significantly sharper decline compared to the AI tech markets, specifically the Google market. It also had a higher volume because the trades are not split across multiple sub markets, and both had extremely variable liquidity. However, the Israel x Palestine War market has significantly higher liquidity, so large shares were overall less likely to affect predictions.

There was no indication of these two unrelated markets affecting each other, likely due to the drastically different topics of the two.

5.4 Patterns in the Israel x Hamas Ceasefire Market

In the "Israel x Hamas Ceasefire Before August?" market, there are some patterns that we can observe. The sharp increase in yes share price around July 23rd can be attributed to the ceasefire proposal submitted by Hamas. Likewise, the decrease around July 25th can be attributed to the United States and Israel withdrawal from the ceasefire negotiations.

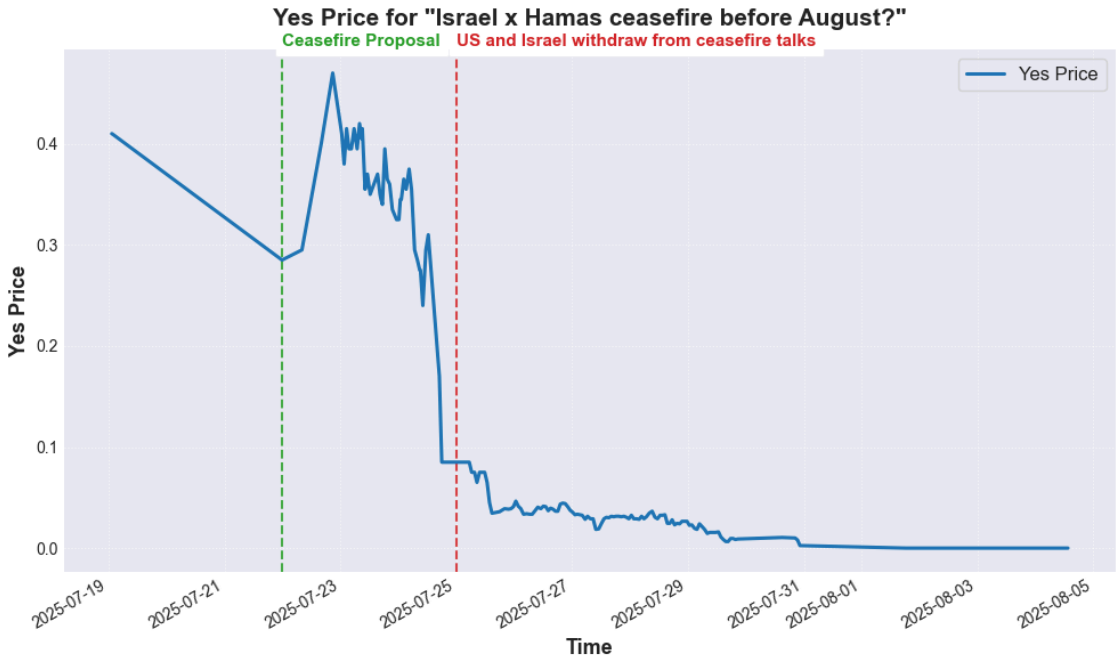

Figure 18. "Israel x Hamas Ceasefire Before August?" Yes Price External Events Timeplot

Additionally, all statistics seem to level out and become less variable as the month nears its end, as the market only accounts for a ceasefire before August. As pictured, the yes share price returned to zero after July 31st for the same reason, since there was no ceasefire in July. Prices also become more extreme (either high or low) as the month goes on; this will be observed in the AI Technology markets as well later on.

5.5 Patterns in the Best AI Model Market

For the AI Model event, we will primarily be analyzing Google, as it is the market with the largest standard deviation, range, and generally more variable.

The graph of yes price for Google is relatively consistent, with very little extreme spikes in either direction. The increase on July 23rd can be attributed to the announcement of Google's strong Q2 earnings, which they attributed to their focus on AI. This may have driven investors to

invest more in "yes" for Google, foreseeing further improvement. Additionally, the price also increased more and more near the end of the month as traders became more confident in the market. After July 31st, this market price increased to $1.00, because Google had the top AI model on July 31st.

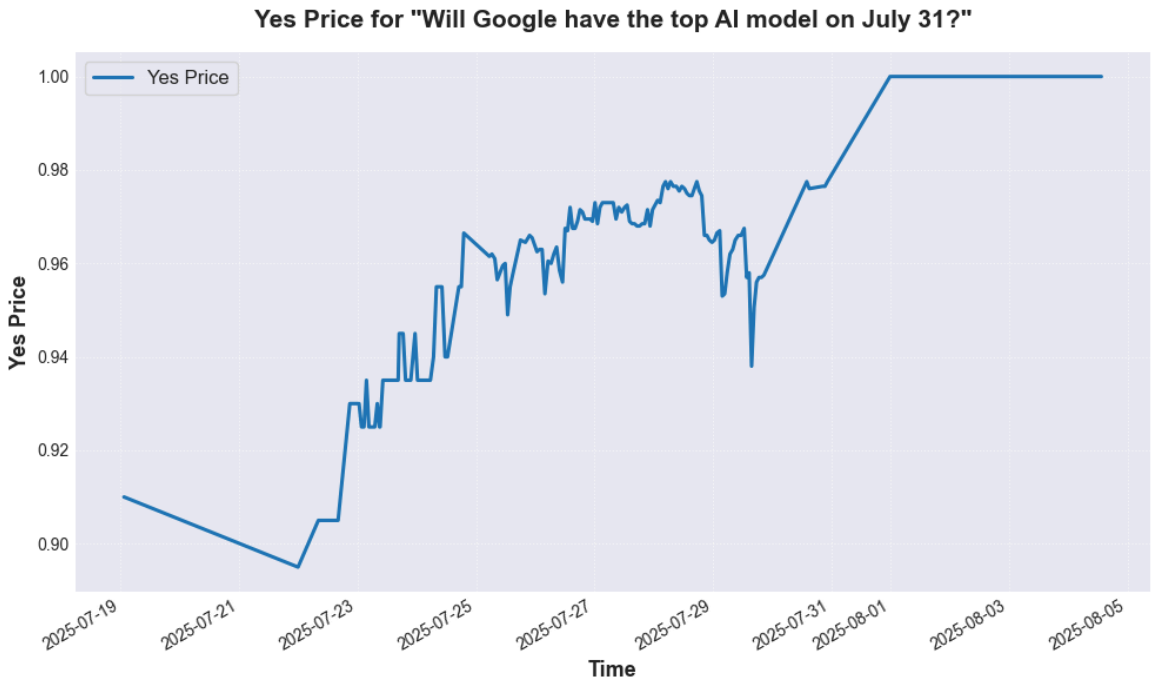

Figure 19. "Will Google Have the Top AI Model on July 31st?" Yes Price Timeplot

Since Google's yes share price remained above $0.90 for almost the entire data collection period, every other company was therefore relegated to an extremely low probability. In the other graphs, the same decline near the end of the month can be observed as the market becomes increasingly certain of which AI model is the best at the end of the month.

VI. Conclusion

Prediction markets are a powerful tool for data analysis. Tracking these different statistics across several different markets gives us a good comparison of how varying different markets and statistics can be, and how each measurement responds to changes in markets. This is useful for those looking to invest in these markets and also investors and government officials, as well as anyone looking to stay informed of current news. For example, tracking spikes in the Israel x Hamas Ceasefire market can make large moments of change extremely obvious. This is the same for any investment related markets, such as the AI Technology markets. Because investors have a financial incentive to stay well-informed, prediction markets can be extremely accurate and have their place in both investment technologies as well as academia as sources of data.

REFERENCES

1. [fill out with references]