The Zohran Effect: What Moves Prediction Markets?

ECON 4438: Applied Econometrics

Milan Singh

October 25, 2025

1 Introduction

The New York City mayoral race has gotten a lot of attention of late. On June 24, Zohran Mamdani, a state assemblyman and a member of the Democratic Socialists of America, beat former New York governor Andrew Cuomo in the Democratic primary. This was widely seen as an upset: of the 32 polls of the primary included in the *New York Times* database (see References), only two had Mamdani leading Cuomo.[1]

However, Mamdani's odds of becoming mayor began to surge at the start of June, based on data from Polymarket, a popular online prediction market. This raises a question: what did Polymarket traders know, and when did they know it? Were traders adjusting their bets based on polls that showed a tightening race? Or were they looking at neighborhood-level turnout during the early voting period? What about campaign finance reports? In other words, what factors drive the prices on betting markets for elections?

This race provides an especially interesting case study. Mayoral elections—even this one—are much lower-salience events than a presidential or Senate race. There is less polling done and fewer historical comparisons to draw on. This election is also of particular interest to the author because the polling organization that he heads conducted a survey of the primary, released on June 23, which turned out to be quite far off.[2] One goal of this paper, then, is to determine how much of an effect the author's inaccurate poll had on the betting markets.

In this paper, I construct a model that uses changes in polling averages for the Democratic primary, fundraising data for both candidates, and an indicator variable for early voting to predict Mamdani's odds of becoming mayor of New York on Polymarket. I find that movement in polling averages has a miminal effect on Mamdani's odds. Fundraising data for both candidates moves the prediction market in the way one would expexct: increases in Cuomo's (Mamdani's) cumulative funds raised are associated with an decrease (increase) in Mamdani's odds. There is a statistically significant but small effect of an indicator variable for whether a given day featured early voting in the Democratic primary on Mamdani's odds; the sign of the estimated coefficient is negative, implying that markets saw early voting turnout data as negative for Mamdani. Over 18 specifications of model, the polls-only equation explains at most 17.5% of the variation in Mamdani's Polymarket odds; polls and fundraising models explain at most 20.3% of the variation; and polls, fundraising, and early voting models explained up to 97.3 % of the variation.

Section 2 summarizes the literature on what factors affect political betting markets. Section 3 discusses data sources, the paper's hypothesis, and how the hypothesis is tested. Section 4 discusses results, and Section 5 offers some concluding remarks, including some ideas about what other factors might account for the variation in Mamdani's Polymarket odds that is not explained by this model.

[1] New York City uses ranked choice voting for municipal primary elections, while general elections use first past the post voting. Not all polls in the *New York Times* database included full ranked choice tabulations. A survey from Public Policy Polling fielded from June 6-7 had Mamdani leading Cuomo by 5 points in the first round, but did not include full ranked choice tabulations. A survey from Emerson College fielded from June 18-20 had Mamdani leading Cuomo by 4 points in the final round of ranked-choice tabulations.

[2] We had Cuomo beating Mamdani 57-43 in the final round; the actual result was Mamdani beating Cuomo 56-44 in the final round. Our postmortem analysis found that poor weighting assumptions accounted for 10 points out of the total 26-point miss.

2 Literature review

Prediction markets for elections have a long and storied history in the United States, beginning in the late 19th century and experiencing a resurgance in the early 2000s (Rhode & Strumpf 2004; Wolfers & Zitzewitz 2004). In recent years, prediction markets have grown further, particularly after a 2018 Supreme Court decision overturned a federal law which banned sports gambling in most jurisdictions (Singh 2022).

Several studies have looked at the effect of prediction markets or election outcomes on the stock prices of particular firms or industries. Snowberg et al. (2007) analyzes prediction market data for presidential elections from 1880 to 2004 and finds that "electing a Republican president raises equity valuations by 2-3 percent" and that since Reagan, Republican presidents have raised bond yields. Snowberg et al. (2006) looks at prediction markets for control of Congress and finds that while "equity values, oil prices, and Treasury yields are slightly higher with Republican majorities in Congress," a switch in the majority party of a chamber of Congress has only 10-30 percent of the effect of a switch in partisan control of the presidency.

Knight (2006) tests whether the Bush and Gore campaign platforms were factored into the equity prices of firms that stood to gain or lose from each candidates platform, finding that "under a Bush administration, relative to a counterfactual Gore administration, Bush-favored firms are worth 3% more and Gore-favored firms are worth 6% less, implying a statistically significant differential return of 9%" based on data on the probability of a Bush victory from the Iowa Electronic Market. Shon (2006) looks at the level and partisanship (ratio of donations to Democratic vs. Republican presidential candidates) of campaign contributions for the 2000 election and equity returns during the Florida recount period, finding that there is a positive (negative) relationship between both the level and partisanship of contributions to Bush (Gore) and equity prices of Bush (Gore) favored firms.

Wagner et al. (2018) looks at the stock prices of companies that stood to benefit from the 2017 Trump tax cuts as the prospects for tax reform and the contents of the reform bill changed from the 2016 election through the end of 2017. Over the course of the period, firms with high tax burdens saw their stock prices rise relative to the market overall, while firms with heavy foreign exposure saw their stock prices fall behind the market, suggesting that the stock market was incorporating information about the Trump administration's economic policy.

Other papers investigate the factors that drive trading on prediction markets. Brown et al. (2019) looks at prediction markets for American elections and finds that "poll releases stimulate an immediate uptick in trading activity," but that "much of this activity involves relatively inexperienced traders, meaning that the price efficiency declines in the immediate aftermath of a poll release, and does not recover until more experienced traders enter the market in the following hours." The implication is that "information releases do not necessarily improve prediction market forecasts" as the influx of "noise traders" temporarily reduces price efficiency. Rothschild & Sethi (2016) study prediction markets for the 2012 presidential election on a platform called Intrade, finding evidence that one particular trader lost $4 million by trying to manipulate the market price in favor of Mitt Romney. Reade & Williams (2019) develop "an empirical method for converting vote shares into outcome probabilities, then compare the two sources of forecasts by considering both the bias and the precision," finding that "converted opinion polls perform well in terms of bias, while prediction markets are good for precision."

This paper contributes to this literature by investigating what factors—polling, fundraising, turnout data from early voting—move the prediction market for the winner of the 2025 New York City mayoral election. The next section describes data sources and the paper's hypothesis.

3 Data and hypothesis

Betting market data was downloaded from Polymarket on October 10, 2025 at around 4:30pm. I obtained both hourly and daily market data. However, the hourly data only goes back to July 10, 2025, which is after the Democratic primary (which took place on June 24). The hourly and daily Polymarket datasets are merged, such that all price data from April 23, 2025 to July 9, 2025 is daily, and all price data from July 10, 2025 to October 10, 2025 is hourly.

Presumably, the switch to data reported hourly reflects increased interest in this election after Mamdani's victory leading to an influx of bettors and a more liquid market. This influx has not had much of an effect on Mamdani's odds of winning the mayoralty: per Figure 1, he has remained the favorite since June 25, with his odds staying within the 75-90% range.

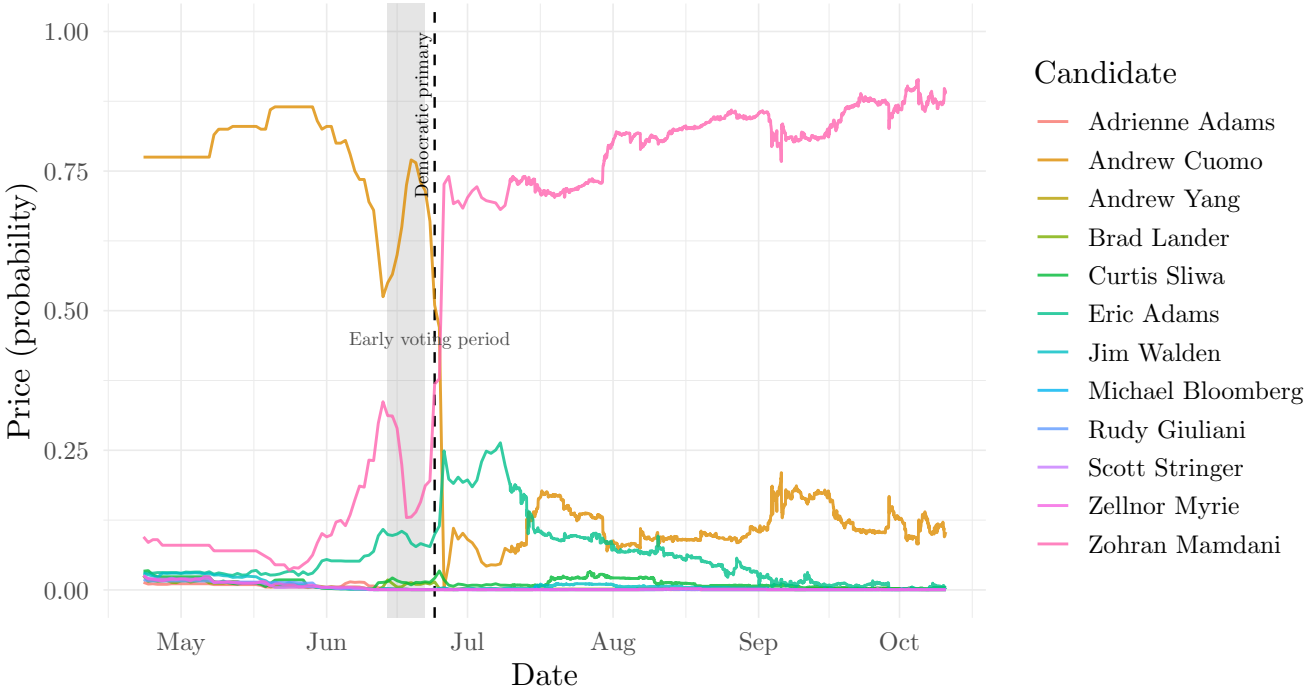

Figure 1: Polymarket odds for New York City mayoral race

A couple interesting patterns can be seen in the betting odds. First, the initial surge in Mamdani's odds of becoming mayor came in early June, well before early voting started. This is when the first poll showing Mamdani leading in the first round of the Democratic primary was released (on June 12). Once early voting was underway, Mamdani's odds fell (and Cuomo's odds rose). On June 24th, the day of the Democratic primary, Mamdani's odds jumped to around 75%, as election returns showed him leading Cuomo in the first round, making it all but certain that he would win the primary. Since the primary, Mamdani's odds have steadily but non-monotonically increased; as of October 10th, Mamdani has a roughly 80% chance of becoming mayor.

What accounts for the movements in this prediction market? There are a couple of possible options:

- **New polls releasing.** As new polls of the Democratic primary were released, Polymarket traders may have updated their beliefs about the state of the race. In particular, the release of a poll showing Mamdani gaining ground—even if he was not leading Cuomo—may have made traders more bullish about his candidacy.
- **Fundraising reports.** The New York City Campaign Finance Board (CFB) maintains an official database of total reciepts and contributions, cash on hand, public matching funds received, and expenditures for all candidates. News stories about CFB fundraising reports—particularly stories that focused on Mamdani's success with public matching funds and Cuomo's success with private donors—may have informed expectations about the race.
- **Turnout data from early voting.** During the early voting period, observers might have looked at turnout rates in particular neighborhoods or boroughs to guess whether certain demographic groups expected to favor Cuomo or Mamdani were turning out at higher or lower than expected rates. Zachary Donnini '25, the author's colleague at the Yale Youth Poll, noted ahead of the primary election that early voting data indicated an electorate that was likely to be favorable to Mamdani, which was confirmed in post-election analyses (see Singh et al.; Beck et al.).

Polling data were obtained from the *New York Times* database ("New York City Mayoral Election 2025: Latest Polls"). There are three types of polls in the database: first-round polls for the Democratic primary, final round polls for the Democratic primary, and general election polls.[3] Figure 2 displays a scatter plot of these polls, for both the primary and the general election.

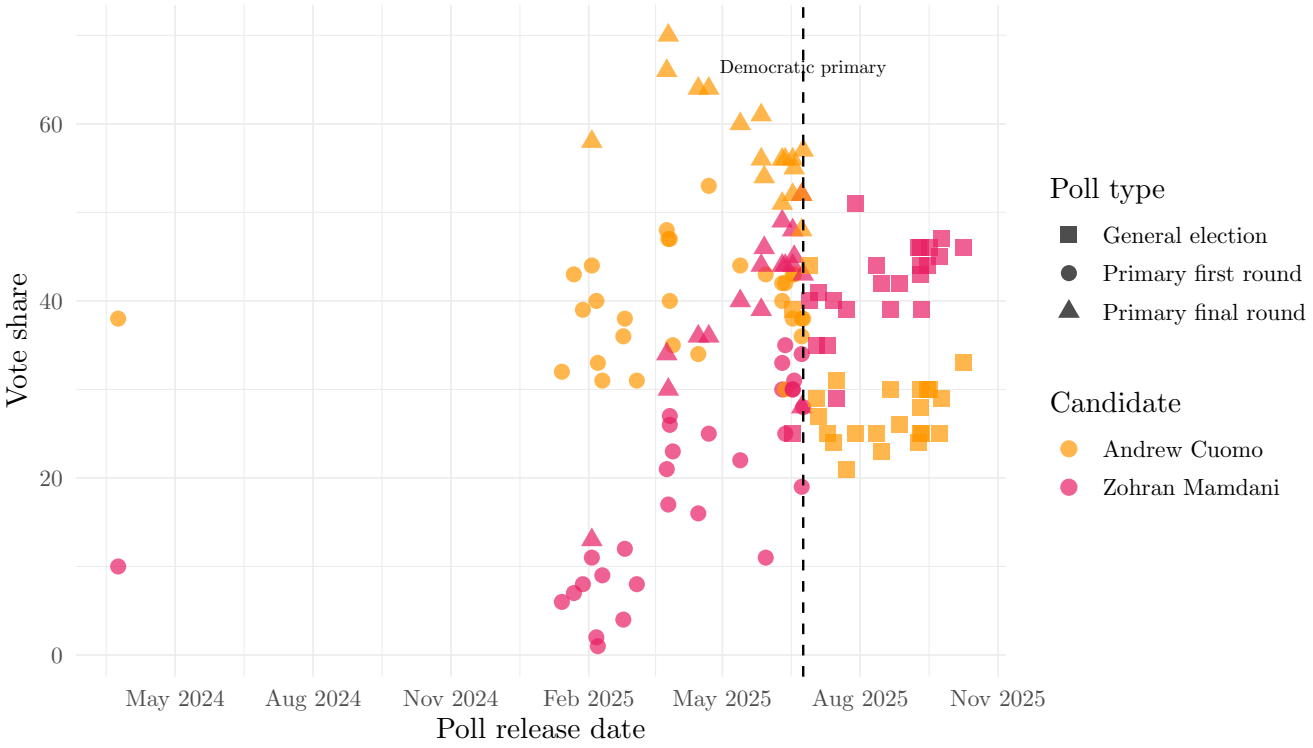

Figure 2: Polls of the New York City mayoral race

At the start of 2025, Cuomo had a fairly large lead in the polls (of all three types). By the late spring/early summer, Cuomo's margin over Mamdani began to tighten significantly, which coincides with the first spike in Mamdani's odds (and corresponding decline in Cuomo's odds) on Polymarket.

To get a more precise picture of how polls shifted, I construct 7-day, 14-day, and 30-day rolling polling averages for the first round of the Democratic primary, the final round of the Democratic primary, and the general election. These averages are displayed in Figure 3.

The polling averages are used to create a daily polls movement variable, defined as:

$$P_{r,w,t} = (\hat{M}_t - \hat{C}_t) - (\hat{M}_{t-1} - \hat{C}_{t-1}) \tag{1}$$

Here, $P_{r,w,t}$ is the change in Mamdani's margin over Cuomo on day t in polling average for election round r (primary first round, primary final round, or general election) with a w-day window. $\hat{M}_t$ and $\hat{C}_t$ represent Mamdani's and Cuomo's average vote shares in a given polling average on day t, respectively. For the first day on any polling average, $P_{r,w,t} = 0$. If polling averages move prediction markets, we should find that a positive (negative) value for $P_{r,w,t}$ is associated with an increase (decrease) in Mamdani's Polymarket odds on a given day.

[3]As of October 11, 2025, when data were collected, the *New York Times* database contained a total of 32 polls of the Democratic primary (of which 18 included final-round tabulations) and 23 polls of the general election. Two first-round primary election polls—one from Slingshot Strategies, added on January 17, 2025; and one from NYT/Siena College, added October 26, 2024—were dropped because they did not include Mamdani.

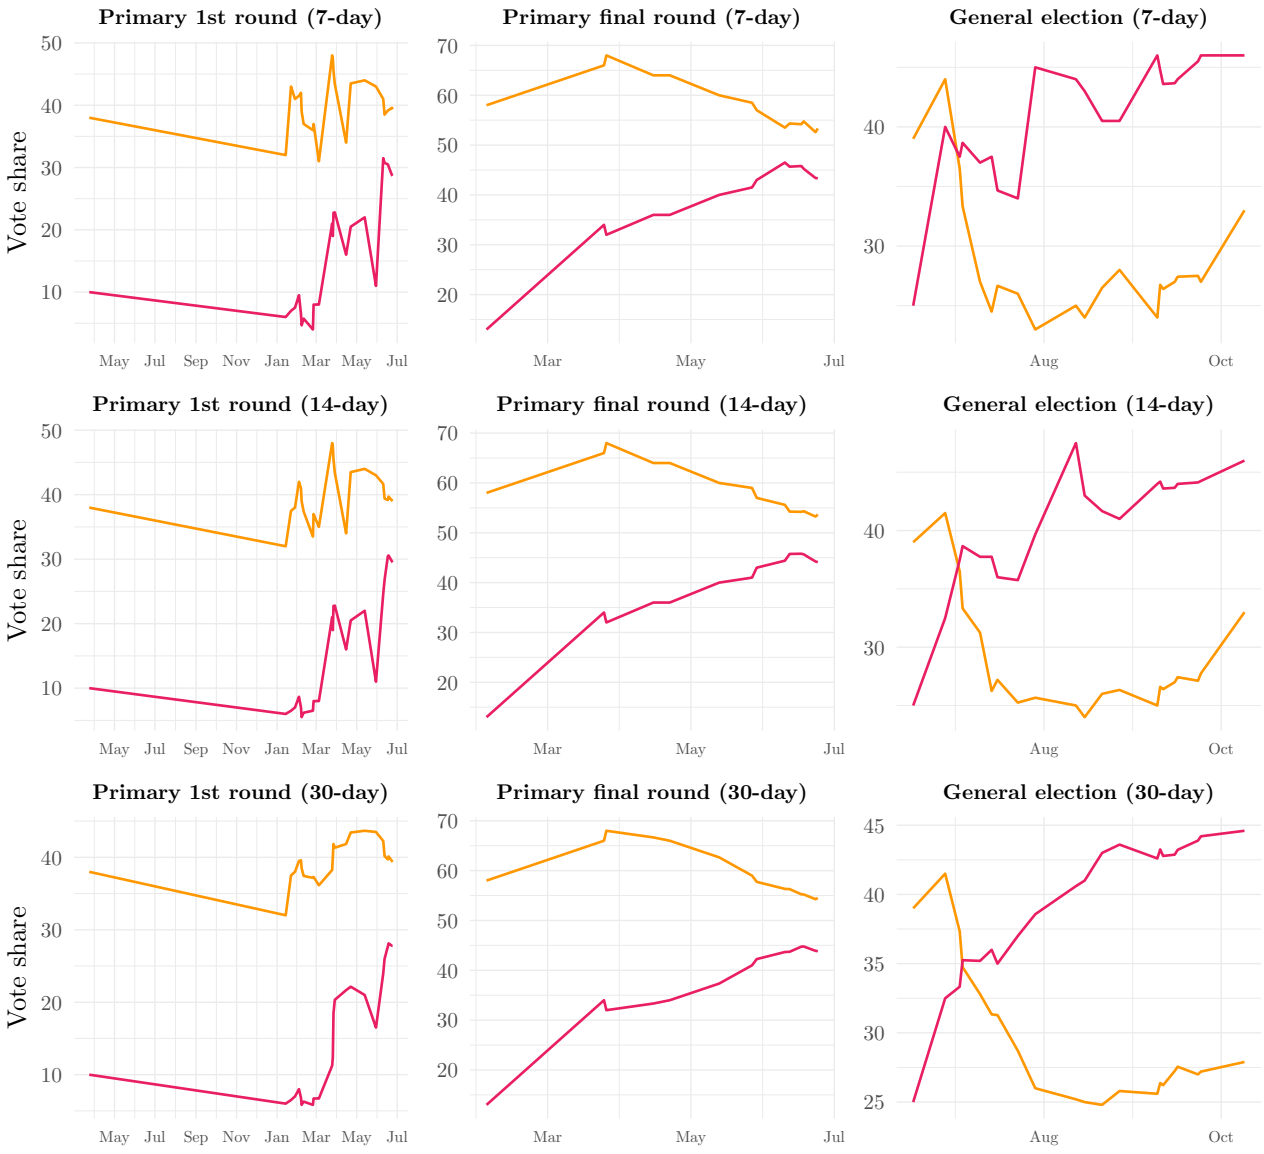

Figure 3: Polling averages for the New York City mayoral election

Campaign finance data were obtained from the CFB's website ("Data Library"). The CFB data reports include contributions to each candidate by date.[4] Figure 4 shows cumulative fundraising totals for Mamdani and Cuomo. Note that Cuomo announced his candidacy for mayor on March 1, 2025, which probably explains the initial surge in Cuomo contributions around that time. Throughout the entire race—even after Mamdani won the Democratic primary—Cuomo has had a sizeable fundraising advantage.

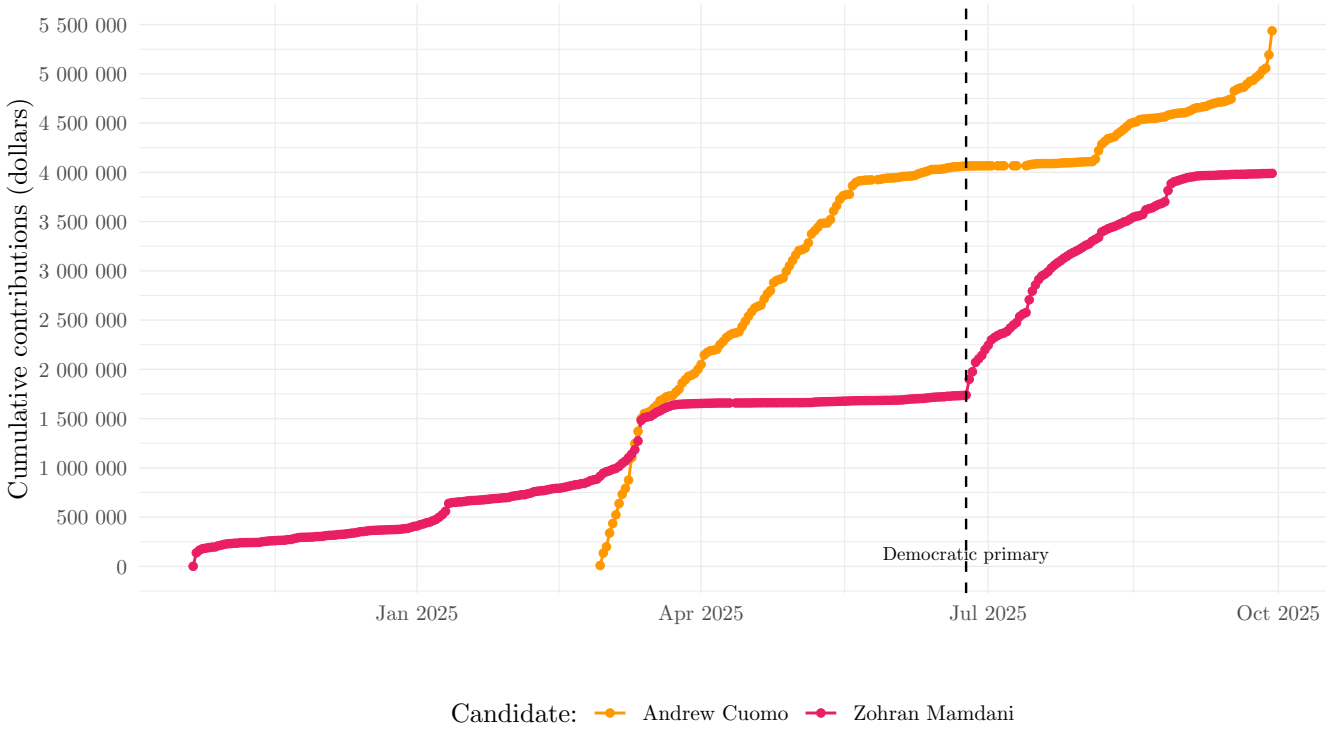

Figure 4: Cumulative fundraising in the 2025 New York City mayoral race

A full description of all variables can be found in Table 1.

Table 1: Variable Definitions

Variable	Definition
M_t^p	Change in Zohran Mamdani's Polymarket odds on day t
$P_{r,w,t}$	Change in Mamdani's margin over Cuomo on day t in a w-day rolling window polling average for election round r ($r = 1$, primary first round; $r = 2$, primary final round; $r = 3$, general election)
F_t^m	Cumulative funds raised by Zohran Mamdani on day t, millions of dollars
F_t^c	Cumulative funds raised by Andrew Cuomo on day t, millions of dollars
E_t	1 if there is early voting for the Democratic prinmary for New York City mayor on day t; 0 otherwise

I propose the following hypothesis: movement in polling averages, fundraising reports, and early voting data

[4]New York City has a program that matches small-dollar donations. However, the CFB reports does not report public matching fund disbursements by date, so I could not include these data in the model. This omission might bias the absolute value of estimated coefficients downard, as Mamdani did quite well with public matching funds.

are all incorporated into prediction markets. This hypothesis is tested with the following regression:

$$M_t^p = \beta_0 + \beta_1 P_{r,w,t} + \beta_2 F_t^m + \beta_3 F_t^c + \beta_4 E_t + \epsilon_t \tag{2}$$

Equation (2) tests whether movements in polling averages, fundraising reports, and early voting are predictive of movements in Mamdani's Polymarket odds. If my hypothesis is true, then $\hat{\beta_1}, \hat{\beta_2} > 0$ and $\hat{\beta_3} < 0$; the sign of $\hat{\beta_4}$ would indicate whether early voting data was seen as favorable for Mamdani or Cuomo.

4 Results

OLS estimates were obtained for 18 different specifications of equation (2). Importantly, estimates were only obtained for the time period of April 23, 2025 to June 24, 2025, the date of the Democratic primary. After June 24, Mamdani's Polymarket odds have not moved all that much (see Figure 1).

OLS estimates are obtained using the 7-day, 14-day, and 30-day versions of $P_{r,w,t}$ for the first round of the Democratic primary. The results of these nine regressions are displayed in Tables 2-4. I also obtained OLS estimates using the 7-day, 14-day, and 30-day versions of $P_{r,w,t}$ for the final round of the Democratic primary. The results of these nine regressions are displayed in Tables 5-7.

I did not regress on the versions of $P_{r,w,t}$ for the general election, since most of the market movement occured before the Democratic primary. Because New York City is a heavily-Democratic jursidiction, it stands to reason that the bettors treat winning the Democratic primary as tantamount to election.

Table 2: Regression Results: 7-day change in primary round 1 polling average

	Dependent variable:		
	Mamdani Polymarket odds movement		
	Baseline	+Fundraising	+Early Voting
	(1)	(2)	(3)
$P_{1,7,t}$ (polling average movement)	0.00250***	0.00179***	0.00002
	(0.00038)	(0.00044)	(0.00008)
F_t^c (Cuomo fundraising)		0.02925	−0.34344***
		(0.14455)	(0.02683)
F_t^m (Zohran fundraising)		−0.93152*	5.75209***
		(0.49018)	(0.12527)
E_t (Early voting)			−0.19649***
			(0.00258)
Constant	−0.00713**	1.47898***	−8.38226***
	(0.00338)	(0.49486)	(0.15779)
Observations	204	204	204
R^2	0.17936	0.21561	0.97401
Adjusted R^2	0.17530	0.20384	0.97349
Residual Std. Error	0.04532 (df = 202)	0.04452 (df = 200)	0.00812 (df = 199)
F Statistic	44.14888*** (df = 1; 202)	18.32502*** (df = 3; 200)	1,864.71000*** (df = 4; 199)

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses. Sample period through June 24.

Table 2 displays coefficients for equation (2)) estimated using the 7-day rolling average of first-round Democratic primary polls. The first column shows results for a model using only poll movement, while the second and third columns add in fundraising data and the early voting indicator variable. In the baseline model, $\hat{\beta_1}$ is estimated at 0.00491 (statistically significant at the $p < 0.01$ level) and 9.7% of the variation in Mamdani's Polymarket odds is explained. In the polls and fundraising model, $\hat{\beta_1}$ drops to 0.00252, while $\hat{\beta_2} = 0.05612$

Table 3: Regression Results: 14-day change in primary round 1 polling average

	Dependent variable:		
	Mamdani Polymarket odds movement		
	Baseline	+Fundraising	+Early Voting
	(1)	(2)	(3)
$P_{1,14,t}$ (polling average movement)	0.00256^{***}	0.00103	-0.00014
	(0.00062)	(0.00069)	(0.00012)
F_t^c (Cuomo fundraising)		0.08728	-0.33845^{***}
		(0.14981)	(0.02689)
F_t^m (Zohran fundraising)		-1.47946^{***}	5.72065^{***}
		(0.50313)	(0.12687)
E_t (Early voting)			-0.19696^{***}
			(0.00249)
Constant	-0.00697^{*}	2.18985^{***}	-8.34757^{***}
	(0.00382)	(0.50040)	(0.15954)
Observations	204	204	204
R^2	0.07891	0.16052	0.97417
Adjusted R^2	0.07435	0.14793	0.97365
Residual Std. Error	0.04801 (df = 202)	0.04606 (df = 200)	0.00810 (df = 199)
F Statistic	17.30537^{***} (df = 1; 202)	12.74776^{***} (df = 3; 200)	$1{,}876.52100^{***}$ (df = 4; 199)

Note: *p<0.1; **p<0.05; ***p<0.01

Standard errors in parentheses. Sample period through June 24.

Table 4: Regression Results: 30-day change in primary round 1 polling average

	Dependent variable:		
	Mamdani Polymarket odds movement		
	Baseline	+Fundraising	+Early Voting
	(1)	(2)	(3)
$P_{1,30,t}$ (polling average movement)	0.00491^{***}	0.00252^{**}	-0.00018
	(0.00106)	(0.00119)	(0.00021)
F_t^c (Cuomo fundraising)		0.05612	-0.33859^{***}
		(0.15033)	(0.02709)
F_t^m (Zohran fundraising)		-1.32247^{***}	5.72873^{***}
		(0.50542)	(0.12670)
E_t (Early voting)			-0.19697^{***}
			(0.00251)
Constant	-0.00830^{**}	2.04394^{***}	-8.36095^{***}
	(0.00384)	(0.49398)	(0.15864)
Observations	204	204	204
R^2	0.09652	0.16997	0.97411
Adjusted R^2	0.09204	0.15752	0.97359
Residual Std. Error	0.04755 (df = 202)	0.04580 (df = 200)	0.00811 (df = 199)
F Statistic	21.57897^{***} (df = 1; 202)	13.65146^{***} (df = 3; 200)	$1{,}871.58800^{***}$ (df = 4; 199)

Note: *p<0.1; **p<0.05; ***p<0.01

Standard errors in parentheses. Sample period through June 24.

(not statistically significant) and $\hat{\beta}_3 = -1.32247$ (statistically significant at the $p < 0.01$ level), and 15.7% of the variation in Mamdani's Polymarket odds is explained.

In the polls, fundraising, and early voting model, the coefficient on poll movement loses statistical significance, while the coefficient on Cuomo fundraising variable is -0.34344 and the coefficient on the Mamdani fundraising variable is 5.75209 (both statistically significant at the $p < 0.01$ level). The estimate for $\hat{\beta}_4$ is -0.19649 (statistically significant at the $p < 0.01$ level) and over 97% of the variation in Mamdani's Polymarket odds is explained.

The results from Table 3 and Table 4 are similar. In both the polls and fundraising and the polls, fundraising, and early voting specifications, the estimates for the fundraising coefficients $\hat{\beta}_2$ and $\hat{\beta}_3$ have the same signs and approximately the same magnitudes as in Table 2, with the same level statistical significance. Adjusted R^2 is larger for the polls, fundraising, and early voting specification than for the polls and fundraising-only specification of the model in Tables 2-4. Notably, the largest jump in adjusted R^2 comes from incorporating fundraising data, with a smaller jump from incorporating the early voting indicator variable.

Table 5: Regression Results: 7-day change in primary final round polling average

	Dependent variable:		
	Mamdani Polymarket odds movement		
	Baseline	+Fundraising	+Early Voting
	(1)	(2)	(3)
$P_{2,7,t}$ (polling average movement)	0.00664***	0.00465***	0.0000003
	(0.00122)	(0.00136)	(0.00035)
F_t^c (Cuomo fundraising)		0.18837	−0.34454***
		(0.14337)	(0.03683)
F_t^m (Zohran fundraising)		−1.28496***	5.11773***
		(0.44641)	(0.15881)
E_t (Early voting)			−0.18500***
			(0.00329)
Constant	−0.00705**	1.44490***	−7.29299***
	(0.00347)	(0.42060)	(0.18709)
Observations	212	212	212
R^2	0.12331	0.17301	0.94928
Adjusted R^2	0.11914	0.16108	0.94830
Residual Std. Error	0.04593 (df = 210)	0.04483 (df = 208)	0.01113 (df = 207)
F Statistic	29.53756*** (df = 1; 210)	14.50461*** (df = 3; 208)	968.51900*** (df = 4; 207)

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses. Sample period through June 24.

Tables 5-7 repeat this exercise but using polling averages for the final round of the Democratic primary. The results are similar to those in Tables 2-4. Movement in polling averages has very little effect on Mamdani's Polymarket odds once the fundraising and early voting variables are incorporated. The magnitude of the coefficients on the Cuomo and Mamdani fundraising variables, as well as the early voting dummy variable, are similar across Tables 5-7. The fully-specified models in Tables 5-7 all predict just under 95% of the movement in Mamdani's Polymarket odds, with the largest jumps in adjusted R^2 coming from adding the early voting dummy variable.

5 Conclusion and discussion

This paper analyzes the effect of movement in the polls, fundraising data, and early voting turnout data on prediction markets for the 2025 New York City mayoral race. Polling movement has only a minimal effect

Table 6: Regression Results: 14-day change in primary final round polling average

	Dependent variable:		
	Mamdani Polymarket odds movement		
	Baseline	+Fundraising	+Early Voting
	(1)	(2)	(3)
$P_{2,14,t}$ (polling average movement)	0.01602***	0.01466***	−0.00208**
	(0.00243)	(0.00315)	(0.00084)
F_t^c (Cuomo fundraising)		0.35624**	−0.38108***
		(0.14643)	(0.03894)
F_t^m (Zohran fundraising)		−1.34755***	5.19056***
		(0.43101)	(0.15918)
E_t (Early voting)			−0.18796***
			(0.00337)
Constant	−0.00850**	0.87375*	−7.26805***
	(0.00337)	(0.45127)	(0.18442)
Observations	212	212	212
R^2	0.17163	0.20884	0.95073
Adjusted R^2	0.16769	0.19743	0.94978
Residual Std. Error	0.04465 (df = 210)	0.04385 (df = 208)	0.01097 (df = 207)
F Statistic	43.51125*** (df = 1; 210)	18.30178*** (df = 3; 208)	998.58600*** (df = 4; 207)

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses. Sample period through June 24.

Table 7: Regression Results: 30-day change in primary final round polling average

	Dependent variable:		
	Mamdani Polymarket odds movement		
	Baseline	+Fundraising	+Early Voting
	(1)	(2)	(3)
$P_{2,30,t}$ (polling average movement)	0.00029	−0.00756***	0.00058
	(0.00225)	(0.00248)	(0.00063)
F_t^c (Cuomo fundraising)		0.05862	−0.33897***
		(0.14731)	(0.03696)
F_t^m (Zohran fundraising)		−1.73746***	5.15421***
		(0.44256)	(0.16321)
E_t (Early voting)			−0.18569***
			(0.00327)
Constant	0.00051	2.75715***	−7.37822***
	(0.00398)	(0.43211)	(0.20793)
Observations	212	212	212
R^2	0.00008	0.16375	0.94949
Adjusted R^2	−0.00469	0.15169	0.94851
Residual Std. Error	0.04906 (df = 210)	0.04508 (df = 208)	0.01111 (df = 207)
F Statistic	0.01606 (df = 1; 210)	13.57686*** (df = 3; 208)	972.74620*** (df = 4; 207)

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses. Sample period through June 24.

on the market; fundraising data has a meaningful impact; and early voting data had the large impact on betting markets. Across specifications, $\hat{\beta}_1$ lacks statistical significance, but $\hat{\beta}_2 > 0$, $\hat{\beta}_3 < 0$ and $\hat{\beta}_4 < 0$. This suggests that betting markets did not react to movement in the polls, but did see more fundraising for Cuomo (Mamdani) as bad (good) for Mamdani and early voting data as bad for Mamdani.

Some existing papers (Snowberg et al. 2007; Snowberg et al. 2006; Knight 2006; Shon 2006; Wagner et al. 2018) seek to analyze the effects of elections on certain business sectors, often finding that firms who stand to benefit from a given candidate's policy platform see their stock prices increase (decrease) when that candidate wins (loses) the election. It would have been interesting to analyze the effect of Mamdani becoming the Democratic nominee on firms that invest heavily in New York City real estate. One of Mamdani's core campaign planks was freezing the rent for rent-stabilized apartments in New York City. Rent-stabilized units make up about 40% of rental units in New York City, and are mostly owned by private landlords (Holliday). The real estate industry appears to have perceived Mamdani as a threat to their bottom lines: landlords contributed $2.5 million to a pro-Cuomo super PAC (Ashford). Unfortunately, there are no publicly-traded firms that focus heavily or exclusively on New York City real estate, which makes it impossible to replicate the analyses used in the literature to gauge at the impact of campaign platforms on specific business sectors.

One interpretation of this paper's results is that polling simply did not move this prediction market very much. This is good for the author, in that it allows him to sleep soundly at night knowing that his firm's poll did not make Polymarket prices less accurate. The results of these regressions suggest that fundraising data had the biggest impact on Polymarket prices, followed by early voting data. But that still leaves more than half of the variation in Mamdani's odds unexplained. What gives?

One explanation is that municipal elections—particularly off-cycle primaries—are simply harder to predict that other kinds of elections. There are far fewer polls released for municipal elections, and it is trickier to make accurate assumptions about what the primary electorate will look like *ex ante* than for other races (see Beck et al.).

Another explanation is that Zohran Mamdani's victory in the primary was genuinely an upset. Throughout the primary, Mamdani lagged Cuomo in the polls (see Figure 2 and Figure 3) and after entering the race Cuomo has comfortably out raised Mamdani (see Figure 4). Perhaps reflecting these data, betting markets had Cuomo as the favorite to win the mayoral election until the day of the Democratic primary.

Studying Figure 1 suggests a possible answer to the question posed in the title of this paper, albeit one that is harder to prove empirically. Mamdani's odds of becoming mayor spiked on the day of the Democratic primary. One of the author's vices is spending time on a particular corner of the internet know as "Election Twitter," where amateur analysts, political junkies, and journalists follow and discuss the results of elections in real-time, as polls close and votes are reported. The author recalls distinctly that on June 24, the overall discourse on Election Twitter concluded within a few hours of polls closing that Mamdani was likely to beat Cuomo in the Democratic primary; the crystallization of this consensus online roughly matches the the window when Mamdani's Polymarket odds spiked on election night.

References

- Ashford, Grace. "Landlords Commit $2.5 Million to Help Cuomo's Mayoral Campaign". The *New York Times*. June 5, 2025, https://www.nytimes.com/2025/06/05/nyregion/cuomo-landlords-donation.html.
- Beck et al. "Postmortem: Why we got the New York City mayoral primary wrong". Yale Youth Poll. August 13, 2025, https://yalepolling.substack.com/p/postmortem-why-we-got-the-new-york.
- Brown et al. "When are prediction market prices most informative?" *International Journal of Forecasting*. January–March 2019, 420-428, https://www.sciencedirect.com/science/article/pii/S0169207018300852.
- "Data Library". New York City Campaign Finance Board. Retrieved October 13, 2025, https://www.nyccfb.info/follow-the-money/data-library/?.
- Holliday Smith, Rachel. "Can Zohran Mamdani Really Freeze the Rent?". The City. September 2, 2025, https://www.thecity.nyc/2025/09/02/rent-freeze-mamdani-affordable-housing-stabilized-apartments/.
- Knight, Brian, "Are Policy Platforms Capitalized into Equity Prices? Evidence from the Bush/Gore 2000 Presidential Election." *Journal of Public Economics*, 2006, 751-773, https://fairmodel.econ.yale.edu/ec438/knight1.pdf.
- "New York City Mayoral Election". Polymarket. Accessed October, 10, 2025, https://polymarket.com/event/new-york-city-mayoral-election?tid=1760132014221.
- "New York City Mayoral Election 2025: Latest Polls". The *New York Times*. Updated October 10, 2025. Retrieved October 10, 2025, https://www.nytimes.com/interactive/polls/nyc-mayoral-election-polls-2025.html
- "Ranked Choice Voting Explained". New York City Public Advocate. May 30, 2025, https://advocate.nyc.gov/blog/ranked-choice-voting-explained.
- Reade, James J. and Williams, Leighton Vaughan. "Polls to probabilities: Comparing prediction markets and opinion polls." *International Journal of Forecasting*. January–March 2019, 336-350, https://www.sciencedirect.com/science/article/pii/S0169207018300633.
- Rhode, Paul W., and Koleman S. Strumpf. "Historical Presidential Betting Markets." *Journal of Economic Perspectives*, Spring 2004, 127-142, https://fairmodel.econ.yale.edu/ec438/rhode1.pdf.
- Rothschild, David and Sethi, Rajiv. "Trading Strategies and Market Microstructure: Evidence from a Prediction Market." *The Journal of Prediction Markets*, 2016, https://www.ubplj.org/index.php/jpm/article/view/1179.
- Shon, John J. "Do Stock Returns Vary With Campaign Contributions? The Bush versus Gore 2000 Presidential Elections." 2006, https://fairmodel.econ.yale.edu/ec438/shon.pdf.
- Singh, Milan. "The case against legal gambling." Slow Boring. March 19, 2022, https://www.slowboring.com/p/gambling.
- Singh et al. "We polled New York City. Here's what we found". Yale Youth Poll. June 23, 2025, https://yalepolling.substack.com/p/new-yoll-city.
- Snowberg et al. "Party Influence in Congress and the Economy." November 30, 2006, https://fairmodel.econ.yale.edu/ec438/swz.pdf.
- Snowberg et al. "Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections." *The Quarterly Journal of Economics*, May 2007, 807-829, https://fairmodel.econ.yale.edu/ec438/snow1.pdf.
- Wagner et al. "Unequal Rewards to Firms: Stock market Responses to the Trump Election and the 2017 Corporate Tax Reform." *AER Papers and Proceedings*, 2018, https://fairmodel.econ.yale.edu/ec438/wagner1.pdf.
- Wolfers, Justin, and Eric Zitzewitz. "Prediction Markets." *Journal of Economic Perspectives*, Spring 2004, 107-126, https://fairmodel.econ.yale.edu/ec438/wolf1.pdf.