How manipulable are prediction markets?

Itzhak Rasooly[1] Roberto Rozzi[2]

[1]City St George's, University of London
[2]University of Siena

27th August, 2025 - EEA Congress 2025

The magic of prediction markets

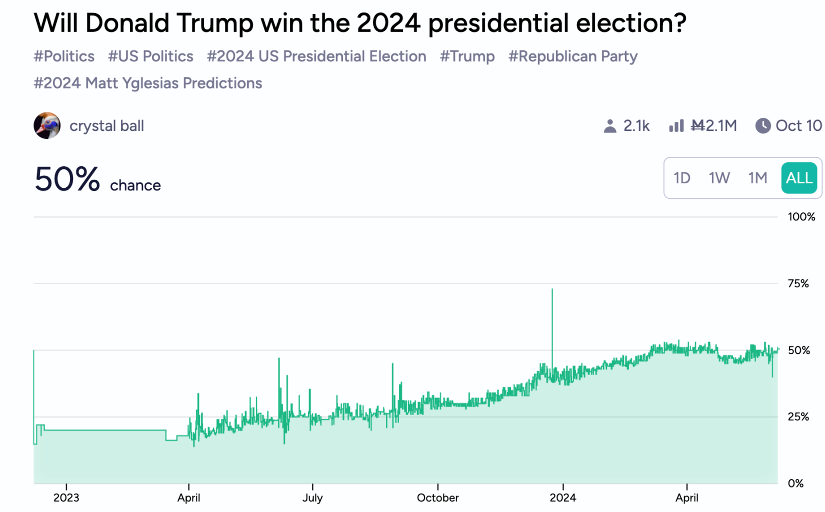

The magic of prediction markets

Prediction markets are remarkable information aggregators:

- Can work as well or better than alternative forecasting methods (Figlewski, 1979; Pennock et al., 2001; Wolfers and Leigh, 2002; Berg et al., 2008; Dreber et al., 2015)
- Largely self-financing
- Perhaps most importantly, often the *only* source of probability estimates on important questions

The perils of manipulation

Despite this promise, prediction markets are hampered by long-standing concerns about manipulability:

- Plenty of manipulation attempts in historical prediction markets (Rhode and Strumpf, 2004)
- Concerns about manipulation were used to justify the cancellation of PAM (Hanson et al., 2006)
  - Stiglitz: '[trading] could be subject to manipulation, particularly if the market has few participants — providing a false sense of security or an equally false sense of alarm'
- Concerns about manipulability also prominent in more recent media coverage (FT, 2023; NYT, 2023; Vox, 2024; WSJ, 2024; NYT, 2024)

THE WALL STREET JOURNAL.

English Edition ▾ | Print Edition | Video | Audio | Latest Headlines | More ▾

MARKETS & FINANCE

A Mystery $30 Million Wave of Pro-Trump Bets Has Moved a Popular Prediction Market

Introduction

TIME

IDEAS • ECONOMY

Don’t Trust the Political Prediction Markets

8 MINUTE READ

Introduction

The New York Times Prediction Markets Tell a Different Story From the Polls

Skeptics wonder if someone has been trying to move these markets. Miller, who predicted a Harris victory a month ago, wrote on his website that he had "identified no single event that explains" the sudden shift in Trump's odds.

Some observers have pointed to a number of recent big pro-Trump bets, including by one pseudonymous user on Polymarket. Others have cited posts by Elon Musk on Oct. 7 promoting a Trump lead in prediction markets, suggesting that he may have encouraged other Trump supporters to jump in. Election Betting Odds, a market results aggregator, shows that Trump's lead jumped around then.

This paper

- First large-scale field experiment on the manipulability of prediction markets ($n = 817$ markets)
- We randomly place *yes bets* (+5 p.p.), *no bets* (-5 p.p.) or *do nothing* (the 'control')
- We collect hourly price data over a 30 day period ($\sim$620k price observations in total) along with rich data on market features (historic trading volume, close date, etc.)
- To help interpret our results, we also build a theoretical model of the impact of price manipulation
- In a follow-up experiment, we replicate our main results using a rather different set of markets ($n = 110$)

Preview of findings

- Prediction markets can be manipulated: the effects of our bets are visible even 60 days after our trades
- However, as predicted by our model, the effect of manipulation decays over time: on average, prices have reverted by about 25% after 1 week
- Markets with more traders, greater trading volume, and an 'external' source of probability estimates are harder to manipulate

Institutional background

Manifold markets

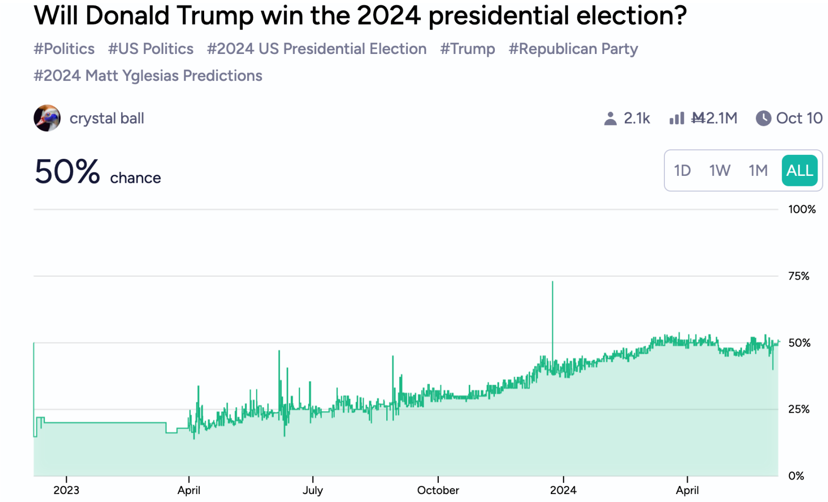

Manifold markets

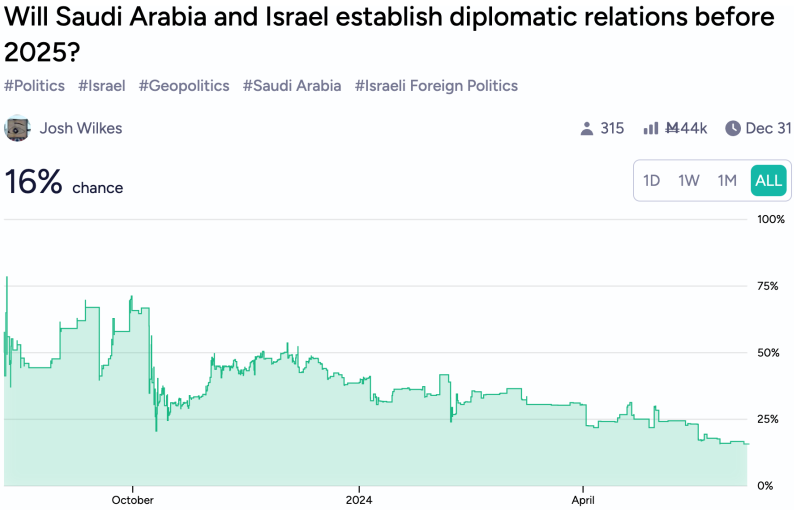

Manifold markets

Will Harvard be found liable for damages to Gino, conditional on a trial verdict being reached by 2026?

#Science #Francesca Gino

Natália 🔍 14 Ṁ1.0k 2026

33% chance

1D 1W 1M ALL

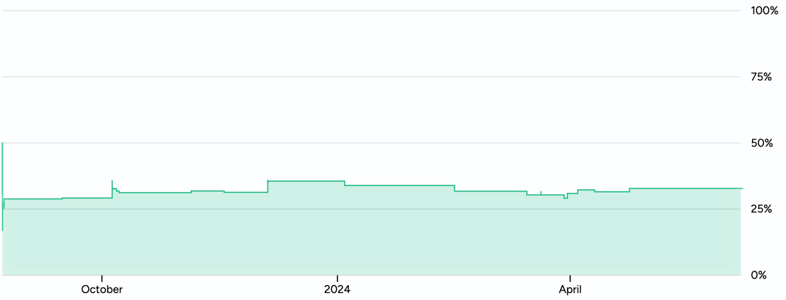

Manifold markets

In some respects, Manifold is an unusual platform:

- Markets are user created and resolved
- Unusually large number of markets
- A large portion of trade is conducted by bots
- The markets run on Maniswap (a generalisation of the constant product rule)
- Markets are run on a platform specific currency (‘Mana’)

Incentives

Despite running on Mana, traders have various incentives to make profitable trades

- Financial incentives: Mana can be converted to charitable donations ($316k raised by Manifold users as of 16 May)
- Social-image incentives (enhanced by leaderboards)
- Self-image incentives (enhanced by personalised Brier scores and calibration charts)

Predictive performance

Given these incentives, it is not surprising that the predictive performance of Manifold is comparable to that of more traditional platforms:

- The markets are generally well-calibrated
- In a study of the 2022 US midterm elections, Manifold outperformed the more traditional prediction markets in the sample (Sigma, 2024)
- Manifold achieves Brier scores that are comparable but slightly worse than Metaculus (EA Forum, 2024)
- See also Servan-Schreiber et al. (2004)

Calibration

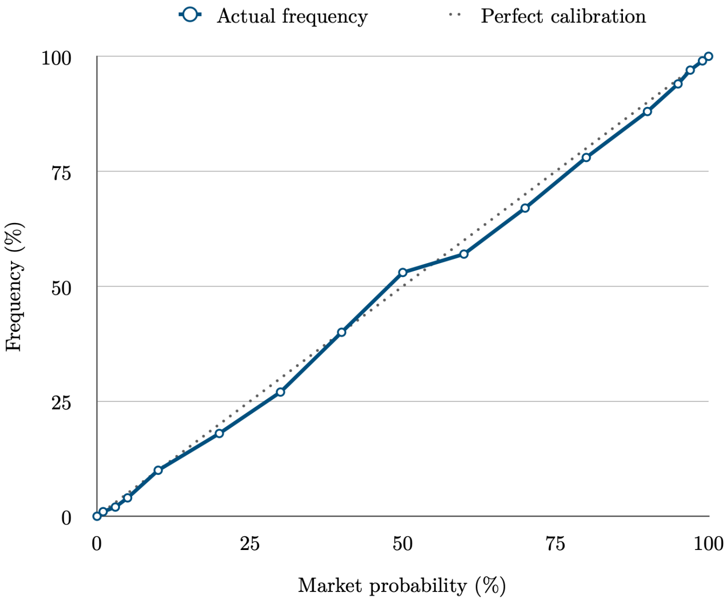

Experimental design

The basic idea

PREREGISTERED

- We conducted a large-scale and 'market level' field experiment ($n = 817$)
- We randomly place *yes bets* (+5 p.p.), *no bets* (-5 p.p.) or *do nothing* (the 'control').
- To see if manipulation yields persistent effects, one can check if the gap in prices between the yes and no groups disappears over time

Exclusion criteria

PREREGISTERED

We excluded markets that

- Resolve after 2025 or within 30 days, or started within the last 7 days
- Had fewer than 10 traders (at the time of our trade)
- Were closely related to another market in our sample
- Cost more than 200M to manipulate in either direction by 5 percentage points

Data

We collected

- Hourly price data, starting 24 hours before the bet and continuing for 30 days ($24 \times 31 \times 822 \approx 610$k prices in total)
- One price 60 days after our bet
- Activity measures: total volume of trade, number of traders, number of comments, etc.
- Whether each market's question was also on Metaculus
- Other information, including each market's question, opening date and closing date

Experimental results

The markets in the sample

Descriptive statistics

Average prices over time (8 days)

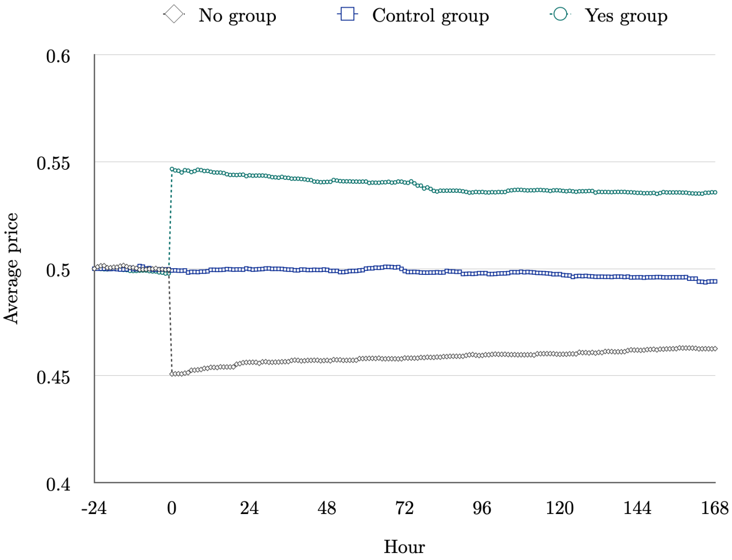

Estimation

To study this formally, we estimate regressions of the form

$$p_{t,i} = \beta_0 + \beta_1 \mathbb{1}_i(\text{`Yes'}) + \beta_2 \mathbb{1}_i(\text{`Control'}) + \beta_3 p_{-1,i} + u_i$$

where

- $p_{t,i}$ is the price in market i at time $t \geq 0$
- $\mathbb{1}_i(\text{`Yes'})$ is a dummy variable that equals 1 if market i is in 'Yes' group
- $\mathbb{1}_i(\text{`Control'})$ is defined analogously
- $p_{-1,i}$ is the price in market i just before the bet

168 regression coefficients (*yes* vs *no*)

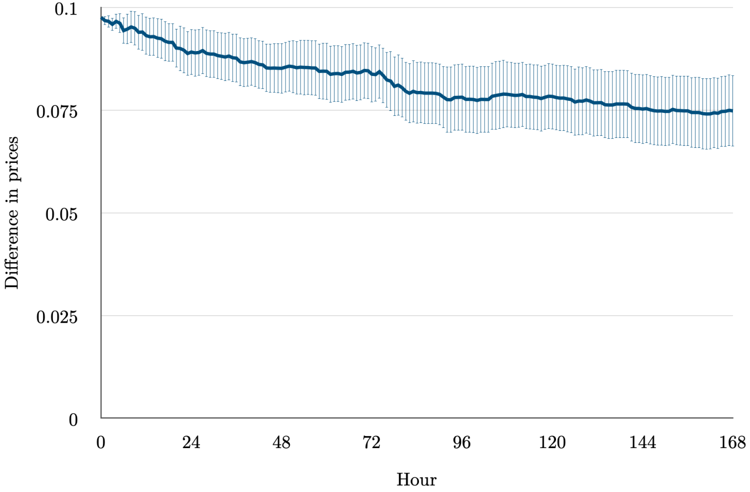

168 regression coefficients (*no* vs *control*)

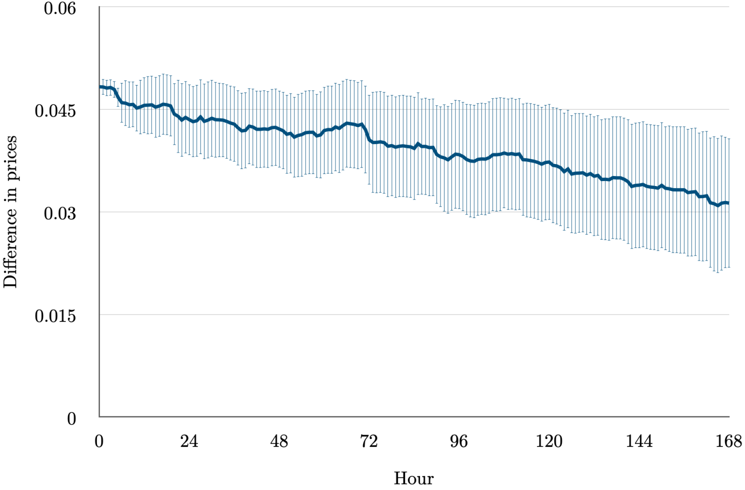

Longer term results (2)

- As we have seen, prices revert by about 25% on average after 7 days
- After 30 days, they have reverted by about 32% on average (a reduction in decay speed, as predicted by our model)
  ▸ graph ▸ all reg
- Despite the expected inflation of standard errors over time, effects are still significant ($p < 0.01$)
- Even after 60 days, effects remain significant (41% reversion in total) ▸ all reg

Heterogeneity in 7 day effects

Variable	Above median	Below median	Difference
Metaculus	0.053***	0.077***	0.024*
24 hour volume	0.049***	0.081***	0.032**
Total volume	0.069***	0.081***	0.012
Comments	0.069***	0.084***	0.015*
Total traders	0.067***	0.083***	0.017*

See also

A follow-up experiment

Sweepcash markets

- In September 2024, Manifold announced that it would introduce 'Sweepcash' markets
- 1 unit of Sweepcash $\approx$ 1 USD
- Useful opportunity to learn whether 'Mana matters'
- *But* fewer Sweepcash markets $\implies$ only 110 markets in the sample $\implies$ effects less precisely estimated
- This time, we shock prices by 4 p.p. (either $\uparrow$ or $\downarrow$) and omit the control

Follow-up Results

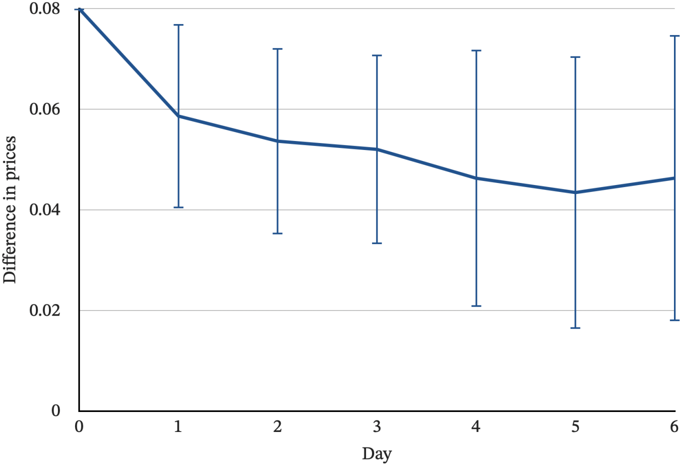

all controls

Conclusions

- In their review of the existing evidence, Wolfers and Zitzewitz (2004) state that manipulation attempts do not have ‘much of a discernible effect on prices, except during a short transition phase’.
- Our large-scale field experiment challenges this conclusion: in the main experiment, we can detect the effects of our manipulations even 60 days after they were made
- However, as predicted by our model, we also find substantial reversion and important heterogeneities in the expected directions

Futre directions

Our experiments also open the door to a lot of future work, e.g.

- *Manipulation via buzz* (e.g. by leaving appropriately chosen comments)
- *Optimal manipulation* (here, one anticipates a ‘U-shape’)
- *Efficiency of manipulated markets* (does the presence of a manipulator affect predictability of markets?)

Thanks for listening!

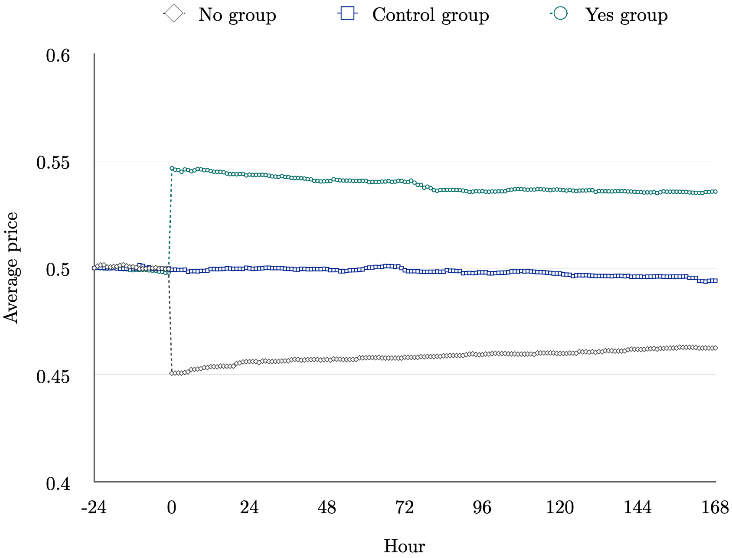

References I

Berg, J., Forsythe, R., Nelson, F., and Rietz, T. (2008). Results from a dozen years of election futures markets research. *Handbook of experimental economics results*, 1:742–751.

Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., Nosek, B. A., and Johannesson, M. (2015). Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences*, 112(50):15343–15347.

EA Forum (2024). Predictive Performance on Metaculus vs. Manifold Markets.

Figlewski, S. (1979). Subjective information and market efficiency in a betting market. *Journal of Political Economy*, 87(1):75–88.

FT (2023). If prediction markets can tell the future, why is the US so afraid of them? *Financial Times Ltd.*

References II

Hanson, R., Oprea, R., and Porter, D. (2006). Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60(4):449–459.

NYT (2023). The Wager That Betting Can Change the World. *The New York Times*.

NYT (2024). Prediction markets tell a different story from the polls. *The New York Times*.

Pennock, D. M., Lawrence, S., Giles, C. L., Nielsen, F. A., et al. (2001). The real power of artificial markets. *Science*, 291(5506):987–988.

Rhode, P. W. and Strumpf, K. S. (2004). Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2):127–142.

References III

Servan-Schreiber, E., Wolfers, J., Pennock, D. M., and Galebach, B. (2004). Prediction markets: Does money matter? *Electronic markets*, 14(3):243–251.

Sigma, F. (2024). Comparing election forecast accuracy.

Vox (2024). Why prediction markets are bad at predicting who'll be president. *Vox*.

Wolfers, J. and Leigh, A. (2002). Three tools for forecasting federal elections: Lessons from 2001. *Australian Journal of Political Science*, 37(2):223–240.

Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126.

WSJ (2024). A mystery $30 million wave of pro-Trump bets has moved a popular prediction market. *The Wall Street Journal*.

Descriptive statistics

Variable	Mean	Std. dev.	Min	Max
24 hour volume	24.7	119.9	0	1,701
Total volume	1,879	4,279	89	56,485
Metaculus	0.093	0.291	0	1
Comments	3.53	7.26	0	83
Traders	27.0	25.0	10	300
Final traders	47.1	44.9	10	321
Final trades	136.0	189.7	13	1,400

back

Short term results with all baseline controls

back

	(1)	(2)	(3)	(4)
Variable	1 week effects	1 week effects	30 day effects	30 day effects
Yes	0.0748***	0.0759***	0.0686***	0.0697***
	[0.00446]	[0.00431]	[0.00908]	[0.00912]
Control	0.0313***	0.0305***	0.0239***	0.0235***
	[0.00484]	[0.00490]	[0.00853]	[0.00854]
$p_{-1,i}$	1.000***	1.003***	1.028***	1.030***
	[0.00781]	[0.00781]	[0.0161]	[0.0155]
Constant	-0.0371***	-0.0297***	-0.0525***	-0.0644***
	[0.00547]	[0.0112]	[0.00910]	[0.0198]
n	817	817	817	817
R^2	0.933	0.937	0.798	0.803

Longer term results

back

	(1)	(2)	(3)
Variable	60 day effects	60 day effects	60 day effects
Yes	0.0559***	0.0855	0.0586***
	[0.0135]	[0.183]	[0.0114]
Control	0.0122	-0.155	0.0258**
	[0.0135]	[0.138]	[0.0103]
$p_{-1,i}$	1.020***	0.860***	1.035***
	[0.0258]	[0.277]	[0.0209]
Constant	-0.0474***	0.0835	-0.0596***
	[0.0141]	[0.135]	[0.0110]
n	817	46	771
R^2	0.617	0.156	0.723

Longer term results (yes vs no)

back

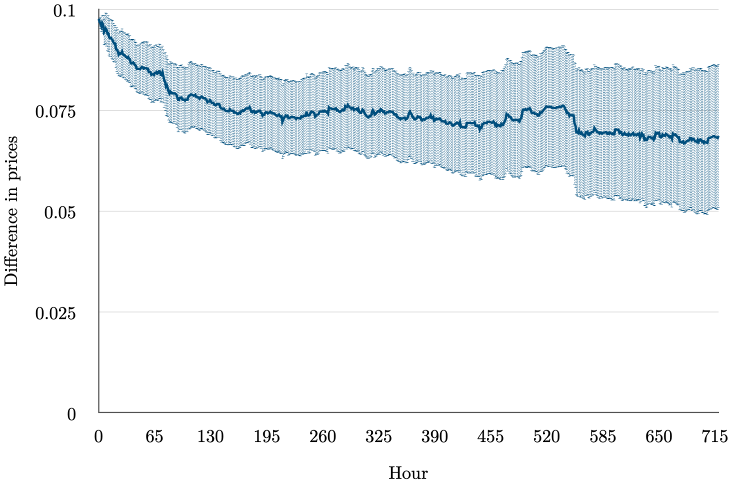

Heterogeneity by number of traders

back

	(1)	(2)	(3)	(4)
Variable	Full sample	25% split	50% split	75% split
Yes	0.0748***	0.0693***	0.0666***	0.0625***
	[0.00446]	[0.00565]	[0.00602]	[0.00999]
Control	0.0313***	0.0268***	0.0264***	0.0259**
	[0.00484]	[0.00594]	[0.00673]	[0.0108]
$p_{-1,i}$	1.000***	0.997***	1.009***	1.003***
	[0.00781]	[0.0103]	[0.0157]	[0.0283]
Constant	-0.0371***	-0.0325***	-0.0406***	-0.0378**
	[0.00547]	[0.00719]	[0.00824]	[0.0149]
n	817	640	419	207
R^2	0.933	0.915	0.908	0.838

Heterogeneity by total volume

back

	(1)	(2)	(3)	(4)
Variable	Full sample	25% split	50% split	75% split
Yes	0.0748***	0.0706***	0.0692***	0.0539***
	[0.00446]	[0.00485]	[0.00623]	[0.0106]
Control	0.0313***	0.0284***	0.0374***	0.0195*
	[0.00484]	[0.00575]	[0.00661]	[0.0112]
$p_{-1,i}$	1.000***	0.999***	1.012***	0.998***
	[0.00781]	[0.0103]	[0.0124]	[0.0245]
Constant	-0.0371***	-0.0359***	-0.0479***	-0.0315**
	[0.00547]	[0.00604]	[0.00699]	[0.0130]
n	817	613	410	205
R^2	0.933	0.926	0.918	0.858

Follow-up Robustness

back

Day	Full sample	Full sample with controls	Reduced sample
0	0.0800***	0.0800***	0.0800***
	[0.000]	[0.000]	[0.000]
1	0.0587***	0.0575***	0.0561***
	[0.00934]	[0.0104]	[0.0111]
2	0.0537***	0.0537***	0.0481***
	[0.00945]	[0.0103]	[0.0112]
3	0.0520***	0.0512***	0.0483***
	[0.00960]	[0.0103]	[0.0112]
4	0.0463***	0.0447***	0.0401***
	[0.0130]	[0.0137]	[0.0151]
5	0.0435***	0.0422***	0.0372**
	[0.0138]	[0.0147]	[0.0160]
6	0.0463***	0.0455***	0.0392**
	[0.0145]	[0.0153]	[0.0165]
7	-	-	0.0348**
	-	-	[0.0169]
n	110	110	89