How Manipulable Are prediction markets?

Itzhak Rasooly¹ Roberto Rozzi²

 $^1 {\rm Sciences}$ Po & Paris School of Economics $^2 {\rm University}$ of Siena

The magic of prediction markets



Prediction markets are remarkable information aggregators:

- Can work as well or better than alternative forecasting methods (Figlewski, 1979; Roll, 1984; Pennock et al., 2001; Wolfers and Leigh, 2002; Berg et al., 2008)
- Largely self-financing
- Perhaps most importantly, often the *only* source of probability estimates on important questions

Partly for this reason, prediction markets are currently undergoing something of a renaissance:

- Polymarket (2020-): \$380 million traded in July 2023
- Kalshi (2021-): ~\$22 million traded in July 2023
- Manifold (2021-): largest prediction market platform as measured by total number of markets

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)

Some questions

This all raises the questions:

- Are these concerns about manipulability justified?
- If so, which markets are most manipulable?

Answering such questions is also an indirect test of the efficient market hypothesis (Fama, 1970):

- If market prices just reflect 'the fundamentals', then the effects of random trades should be transient.
- If markets are inefficient, the effect of random trades could be more persistent.

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 (~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

- Prediction markets can be manipulated: the effects of our bets are visible even 30 days after our trades
- However, as predicted by our model, the effect of manipulation decays over time: on average, prices have reverted by about 24% after 1 week
- Markets with more traders, greater trading volume, and an 'external' source of probability estimates are harder to manipulate

(1) The original inspiration: Camerer (1998) Comment: very different environment, so not surprising that we obtain very different results

(2) Analysis of historical manipulation attempts (Rhode and Strumpf, 2004, 2006; Hansen et al., 2004; Rothschild and Sethi, 2016)

Comment: hard to know the counterfactual price path!

(3) Lab experiments on manipulation (Plott and Sunder, 1982; Hanson et al., 2006; Oprea et al., 2008; Veiga and Vorsatz, 2009; Buckley and O'Brien, 2017; Choo et al., 2022) *Comment*: only study a small number of markets, which are in any case very different from real prediction markets (4) An experiment on the IEM: Rhode and Strumpf (2006) Comment: just 15 bets in total on 2 (inter-related) markets, so only powered to detect immediate effects

(5) Models of prediction markets (Gjerstad, 2005; Manski, 2006; Wolfers and Zitzewitz, 2006; Ottaviani and Sørensen, 2007; Hanson and Oprea, 2009; Chen et al., 2015)

Comment: we study manipulation within a Gjerstad (2005) style model altered to allow for disequilibrium prices and non-price taking behaviour

Manipulation in theory

A model of manipulation

- We consider a single (binary) market.
- A yes share pays out €1 iff the event takes place; a no share pays out €1 iff the event does not take place
- A trader who buys (e.g.) q yes shares has expected utility

$$\pi_i u(w + q - C(q)) + (1 - \pi_i) u(w - C(q))$$

where π_i is their belief about the chance that the event will happen, w is their wealth, C(q) is the cost of the shares

- We assume u' > 0, u'' < 0, $\lim_{w_s \to 0} u' = \infty$ and decreasing -u''/u' (DARA)
- The cost C(q) is determined by an AMM that implements the constant product rule.

To illustrate, suppose that

- The AMM's reserves are (y, n) = (10, 10). Note: $10^2 = 100$.
- If I decide to spend €1 on yes shares, the AMM converts this into 1 yes share and 1 no share.
- Its reserves become (11, 11). But $11^2 = 121 \neq 100!$
- It thus gives me q yes shares, where $(11 q) \times 11 = 100$, i.e. $q \approx 1.9$.

Costs under the constant product rule

Lemma 1

Under the constant product rule,

- MC(0) = n/(n+y)
- MC'(q) > 0 for all $q \ge 0$
- $\lim_{q\to\infty} MC(q) = 1$

Similarly,

- $\lim_{q\to 0^+} AC(q) = n/(n+y)$
- AC'(q) > 0 for all q > 0
- $\lim_{q\to\infty} AC(q) = 1$

Illustration with n = y = 10



Lemma 2

Define $p = \frac{n}{n+y}$. Then

- If $\pi_i > p$, the trader will buy a positive quantity of yes shares.
- If $\pi_i = p$, the trader will not hold any shares.
- If $\pi_i < p$, the trader will buy a positive quantity of no shares.

Lemma 3

Suppose that the price increases from p to $p + \Delta$. Then

- Traders with $\pi_i \ge p + \Delta$ will decrease their holdings of yes shares.
- Traders with $\pi_i \in (p, p + \Delta)$ will switch from holding yes shares to holding no shares.
- Traders with $\pi_i \leq p$ will increase their holdings of no shares.

Simulations

We use simulations to study the price adjustment path:

- The market is initialised and given t periods to reach a stable state; a manipulator then increases the price by 0.05
- The market is then given t' periods to adjust
- At each time, one trader is randomly selected to re-adjust her holdings; thus, we run each simulation 5,000 times
- As an extension, we allow for learning:

$$\pi_i' = \lambda \pi_i + (1 - \lambda)p$$

• In the baseline case, n = y = w = t = t' = 100, s = 10, $\lambda = 0$; we also assume that beliefs are uniform

Results (baseline case)



19/49

Varying the number of traders: m = 10 vs m = 60



Varying the learning rate: $\lambda = 0$ vs $\lambda = 1$



Varying agreement: uniform vs common beliefs



Summary

- The model predicts that manipulation can have persistent effects even without price learning
- However, the model also predicts the effect of manipulation should be somewhat 'undone' by future traders
- The model predicts that markets with more traders, more 'activity', and less learning (e.g due to the existence of external information) should revert faster

Institutional background

Manifold markets

Will Donald Trump win the 2024 presidential election?

#Politics #US Politics #2024 US Presidential Election #Trump #Republican Party #2024 Matt Yalesias Predictions





Manifold markets

Will Saudi Arabia and Israel establish diplomatic relations before 2025?

#Politics #Israel #Geopolitics #Saudi Arabia #Israeli Foreign Politics





Manifold markets

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

#Science #Francesca Gino



In some respects, Manifold is an unusual platform:

- Markets are user created and resolved
- 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')

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)

One highly ranked trader:

'In the unusual world in which I find myself, for better or worse, doing well on a prediction markets website is somewhat of a badge of honour . . . I wish I had more noble motivations but, alas, I think that's a good chunk of it. Another important motivation for me using Manifold relates to charitable giving.'

Effect of the devaluation on GiveWell donations



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

- 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

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 $(24 \times 31 \times 822 \approx 610 \text{k prices in total})$
- 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

- We pre-registered our experiment (with an analysis plan) in December 2023
- We started making bets in December 2023 and finished in April 2024
- We finished the main data collection in May 2024 (and collected some follow-up data in June/July).

Experimental results

The markets in the sample



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

168 regression coefficients (yes vs no)



43/49

168 regression coefficients (no vs control)



Longer term results

- 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)
- 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)

	Above median	Below median
Metaculus	0.053	0.077
24 hour volume	0.049	0.081
Total volume	0.069	0.081
Total traders	0.067	0.083
Comments	0.069	0.084

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: we can detect the effects of our manipulations even 30 days after they made
- However, as predicted by our model, we also find substantial reversion ($\sim 25\%$ after a week) and important heterogeneities in the expected directions

Conclusions

- Our findings somewhat confirm the concerns raised by prediction markets' critics
- However, they do *not* mean that prediction markets are useless: even if manipulable, their prices can still be somewhat informative (Hanson, 2006)
- Although non-causal, our heterogeneity results suggest that making prediction markets more 'active' (higher volume, more traders, etc.) can make them more robust to manipulation attempts

Our experiment also opens 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')

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.
- Buckley, P. and O'Brien, F. (2017). The effect of malicious manipulations on prediction market accuracy. *Information* Systems Frontiers, 19:611–623.
- Camerer, C. F. (1998). Can asset markets be manipulated? A field experiment with racetrack betting. Journal of Political Economy, 106(3):457–482.
- Chen, Y., Gao, X. A., Goldstein, R., and Kash, I. A. (2015). Market manipulation with outside incentives. *Autonomous Agents and Multi-Agent Systems*, 29:230–265.
- Choo, L., Kaplan, T. R., and Zultan, R. (2022). Manipulation and (mis) trust in prediction markets. *Management Science*, 68(9):6716–6732.

References II

- EA Forum (2024). Predictive Performance on Metaculus vs. Manifold Markets.
- Fama, E. F. (1970). Efficient capital markets. Journal of Finance, 25(2):383–417.
- 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.*
- Gjerstad, S. (2005). Risk aversion, beliefs, and prediction market equilibrium. *Economic Science Laboratory, University* of Arizona.
- Hansen, J., Schmidt, C., and Strobel, M. (2004). Manipulation in political stock markets-preconditions and evidence. *Applied Economics Letters*, 11(7):459–463.

References III

- Hanson, R. (2006). Foul play in information markets. In Hahn,
 R. W. and Tetlock, P. C., editors, *Information Markets: A* New Way of Making Decisions in the Public and Private Sectors, pages 126–141. AEI Press, Washington, DC.
- Hanson, R. and Oprea, R. (2009). A manipulator can aid prediction market accuracy. *Economica*, 76(302):304–314.
- 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.
- Manski, C. F. (2006). Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429.
- NYT (2023). The Wager That Betting Can Change the World. The New York Times.

References IV

- Oprea, R., Porter, D., Hibbert, C., Hanson, R., and Tila, D. (2008). Can manipulators mislead prediction market observers? Working Paper 08-01, Chapman University, Economic Science Institute.
- Ottaviani, M. and Sørensen, P. N. (2007). Outcome manipulation in corporate prediction markets. *Journal of the European Economic Association*, 5(2-3):554–563.
- 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.
- Plott, C. R. and Sunder, S. (1982). Efficiency of experimental security markets with insider information: An application of rational-expectations models. *Journal of Political Economy*, 90(4):663–698.

References V

- Rhode, P. W. and Strumpf, K. S. (2004). Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2):127–142.
- Rhode, P. W. and Strumpf, K. S. (2006). Manipulating political stock markets: A field experiment and a century of observational data. *University of Arizona, mimeo*.
- Roll, R. (1984). Orange juice and weather. The American Economic Review, 74(5):861–880.
- Rothschild, D. M. and Sethi, R. (2016). Trading strategies and market microstructure: Evidence from a prediction market. *The Journal of Prediction Markets*, 10(1):1–29.
- Servan-Schreiber, E., Wolfers, J., Pennock, D. M., and Galebach, B. (2004). Prediction markets: Does money matter? *Electronic markets*, 14(3):243–251.

References VI

- Sigma, F. (2024). Comparing election forecast accuracy.
- Veiga, H. and Vorsatz, M. (2009). Price manipulation in an experimental asset market. *European Economic Review*, 53(3):327–342.
- 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.
- Wolfers, J. and Zitzewitz, E. (2006). Interpreting prediction market prices as probabilities. *National Bureau of Economic Research*.