How manipulable are prediction markets?

Itzhak Rasooly* Roberto Rozzi†

1 Overview

The goal of this experiment is to learn whether prediction markets can be manipulated through random trades. To investigate this, we plan to make a series of random bets across hundreds of prediction markets on the Manifold platform. While these bets will automatically shift prices in the short run, our main goal is to learn whether these price effects persist in the long run or are instead 'undone' by the reactions of other traders. We will also study which kinds of market can most easily be manipulated.

2 Experimental design

Manifold. Manifold is a prediction market platform that allows users to bet on whether various events will occur. Prices are determined using an algorithm called 'Maniswap' and can be used to estimate the probability that the event will happen (see Manski, 2006, Gjerstad and Hall, 2005 and Wolfers and Zitzewitz, 2006 for discussion).

Manifold differs in several respects from more traditional prediction markets. First, markets are user-created and user-resolved. As a result, Manifold has a very large number of markets, which makes the kind of large scale field experiment that we will conduct possible. Second, users trade using a virtual currency called 'Mana' (M), which can be converted into charitable donations. This means that market prices on Manifold are not necessarily tied to prices of similar markets on other platforms (via a no-arbitrage condition), which makes manipulation much more feasible. Despite these differences, the evidence suggests that Manifold markets are reasonably well calibrated (Manifold Markets, 2023) and perform comparably to more traditional prediction markets (EA Forum, 2023).

Exclusion criteria. We will only bet on binary markets, i.e. markets that must resolve as 'Yes' or 'No' (or N/A). We will exclude:

1. Markets that do not resolve based on an external event by the end of 2025 (e.g. self-referential markets). *Explanation*: self-referential markets (e.g. "will this market resolve at above 50%?") may be trivial to manipulate but may not be representative of more normal markets.
2. Markets with fewer than 10 traders (at the time of the trade). *Explanation*: while it may be easy to manipulate very small markets, we want to see if we can manipulate larger markets; particularly since prediction markets of interest to the public (e.g. on elections) will generally have many traders.
3. Markets which cost at least 200M to move in either direction by 5 percentage points. *Explanation*: including such markets would reduce the number of markets that we can manipulate holding fixed our experimental budget.
4. Markets that (i) started within the last 7 days or (ii) end within the next 30 days. *Explanation*: markets can be very volatile within the first 7 days of their life cycle, which reduces statistical power. To examine 30 day effects, we need the second restriction.

*Sciences Po & Paris School of Economics.
†University of Padova.

5. Markets that are closely related to other markets in our sample (e.g. "Will Trump win?" vs "Will Trump lose?"). *Explanation*: the prices of such markets are linked by no-arbitrage conditions. As a result, betting on one market can in theory alter the price in the other market, which can bias our estimates due to a "spillover" effect.

Intervention. For each market, we will first randomly determine (with equal probabilities) whether we will bet on 'Yes', bet on 'No', or do nothing (the 'control' group). In the 'Yes' and 'No' groups, we will then buy however many shares are necessary to either increase or (respectively) decrease the market price by 5 percentage points. Note that, although the instantaneous effect of the intervention is obvious, our main interest is in estimating the persistence of the effect on prices.

In the analysis (see below), we will focus on comparing the 'Yes' and 'No' groups. If 'Yes' bets make prices higher than they would have been, and 'No' bets make prices lower than they would have been, then the markets in the 'Yes' group should end up with higher prices than the markets in the 'No' group. Conversely, if neither 'Yes' nor 'No' bets have a persistent effect on market prices, then no persistent differences between the 'Yes' and 'No' groups will be observed.

Data. For each market, we will collect data on:

1. The market question/topic
2. The treatment ('Yes' or 'No')
3. The size of our bet
4. The time of the intervention
5. The market's opening time
6. The market's closing time
7. The total volume of trade in the last 24 hours
8. The total volume of trade
9. The number of traders
10. The 'elasticity' of the price with respect to trades
11. The amount of liquidity
12. The size of the 'pool'
13. Whether the market is publicly listed
14. Whether the market is ranked
15. Whether the market is subsidised
16. Whether the market was created by a moderator
17. The number of user comments on the market
18. Whether there is a similar market on Metaculus

Note that variables 3-15 are Manifold created and are defined on their website.

In addition to these variables, we also record the market prices, starting from 24 hours before our trade and running until 30 days after our trade. Although Manifold stores trading times in Unix time, we will convert this to hourly data to simplify the analysis.

Power. Immediately after our intervention, the 'Yes' and 'No' markets will differ in prices by 10 units (here, we are measuring prices on a 0-100 scale as on Manifold). While it is easy to detect this difference statistically, it becomes more difficult to detect differences as time goes on due to the properties of random walks. We conducted power calculations in order to assess whether we would be able to detect differences across the groups after 1 week.

We conjectured that, on average, the difference in prices between the two 'treated' groups will be about 3 units after 1 week. Using data from a small pilot, we estimated that the variance of prices in our treatment groups would be about 11 units after 1 week. In order to have 90% power to detect this hypothesised difference with $\alpha = 0.05$, one needs a sample size of 566. Since the two treated groups will comprise around 2/3 of the sample (recall that the remaining 1/3 of markets are allocated to the control), this leads to a total sample of $n = 849$. Note that, while this sample size is our target, we are unsure whether there are enough markets on Manifold that meet our inclusion criteria to allow us to achieve this.

3 Analysis plan

Our main regressions will take the form

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

where $\mathbb{1}_i$('Yes') is a dummy variable that equals 1 if market i is in 'Yes' group, $\mathbb{1}_i$('Control') is a dummy variable that equals 1 if market i is in the control group, and $p_{-1,i}$ is the price in the market in the hour before our bet. Thus, the 'No' group is the omitted category. We include the $p_{-1,i}$ variable to increase statistical power, and will also report regressions that include the full suite of controls that are available to us (see above).

We will focus on the difference in mean prices between the 'Yes' and 'No' groups (β_1) since this our best powered comparison. We briefly note three points. First, since $\mathbb{1}_i$('Yes') is randomly assigned, β_1 has a causal interpretation. Second, given that 'Yes' bets drive up prices and 'No' bets drive them down, it seems extremely unlikely that β_1 would be negative; accordingly, we will conduct 1-sided tests. Third, although the regression is cross-sectional, it can be run for various values of $t > 0$ to assess short run, medium run and long run effects. We will focus on effects within one week and run regressions for all $t \in \{1, 2, ..., 7 \times 24\}$. We anticipate that treatment effects will partially, although not entirely, dissipate over the 1 week period.

Although comparing the 'Yes' and 'No' groups allows one to assess the markets' vulnerability to manipulation, it does not allow one to separately estimate the effect of 'Yes' and 'No' bets on prices (unless one makes a symmetry assumption, or instead assumes that prices would counterfactually have followed a random walk). To estimate these effects, we will also compare the 'Yes' and 'No' groups to the control. Specifically, we will use an F-test to evaluate the hypothesis that 'Yes' and 'No' bets have symmetric effects, which is equivalent to the linear restriction that $\beta_1 = 2\beta_2$.

We will also conduct heterogeneity analysis to investigate what kinds of markets can most easily be manipulated. For a each binary pre-treatment variable x_i, we will run regressions of the form

$$p_{t,i} = \beta_0 + \beta_1 \mathbb{1}_i(\text{`Yes'}) + \beta_2 \mathbb{1}_i(\text{`Control'}) + \beta_3 p_{-1,i} + \beta_4 \mathbb{1}_i(\text{`Yes'}) x_i + v_i \tag{2}$$

Here, $\mathbb{1}_i$('Yes')x_i is the interaction term, so β_4 (as usual) captures the difference in treatment effects between the $x_i = 1$ and $x_i = 0$ subgroups. For non-binary pre-treatment variables y_i, we will follow the same procedure after first converting y_i to a binary variable depending on whether it exceeds its median value. Our main expectation is that markets that are more 'active' (e.g. as measured by the number of comments) and markets with a closer 'close date' will be harder to manipulate.

4 Changes to the plan (added: 20 August 2024)

In general, we conformed closely to the design and analysis plan that we pre-registered. (The original plan was submitted in December 2023 and can be viewed in the trial history.) However, we ultimately deviated from the plan in a few (generally minor) ways:

- In the original analysis plan, we planned to conduct heterogeneity analyses using median splits. Although we subsequently conducted these splits, we realised that there was an error in the regression that assesses whether effects are significantly different across the two sub-groups. The correct regression (which we report in the paper) is

$$\begin{aligned} p_{168,i} =& \beta_0 + \beta_1 \mathbb{1}_i(\text{`Yes'}) + \beta_2 \mathbb{1}_i(\text{`Control'}) + \beta_3 p_{-1,i} + \beta_4 \mathbb{1}_i(\text{`Yes'}) x_i \\ &+ \beta_5 \mathbb{1}_i(\text{`Control'}) x_i + \beta_6 p_{-1,i} x_i + v_i \end{aligned}$$

- In the original analysis plan, we planned to conduct one-sided tests to assess if effects are significantly different from zero. However, much of the analysis reported in the paper is graphical and relies on (two-sided) confidence intervals. In the end, we decided to report standard two-sided p-values. However, it is straightforward for the reader to divide these numbers by two if they wish to do so (and this would not change any of our results).
- When analysing the data, we realised that some of the 849 markets on which we had bet inadvertently violated at least one of our exclusion criteria. After discarding these markets, we ended up with $n = 817$ markets in our sample. We should emphasise that our main results are entirely unaffected by the inclusion (or lack of inclusion) of these markets.
- Although we had originally planned to only collect data for 30 days after the shocks, the strength of the results after 30 days suggested that very long-run effects might be observable. As a result, we decided to additionally record the price in every market 60 days after the bets. Note that, although our dataset contains hourly data over the course of the first 30 days, we only recorded a single 'price snapshot' to measure 60 day effects.

References

EA Forum. Predictive performance on metaculus vs. manifold markets, 2023. URL https://forum.effectivealtruism.org/posts/PGqu4MD3AKHun7kaF/predictive-performance-on-metaculus-vs-manifold-markets.

Steven Gjerstad and McClelland Hall. Risk aversion, beliefs, and prediction market equilibrium. *Economic Science Laboratory, University of Arizona*, 2005.

Manifold Markets. Track record and accuracy, 2023. URL https://manifold.markets/calibration.

Charles F Manski. Interpreting the predictions of prediction markets. *Economics Letters*, 91(3): 425–429, 2006.

Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities, 2006.