

Working Papers

Makers and Takers: The Economics of the Kalshi Prediction Market

Constantin Bürgi, Wanying Deng, Karl Whelan

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

Makers and Takers: The Economics of the Kalshi Prediction Market

Constantin Bürgi* Wanying Deng[†] Karl Whelan[‡]

University College Dublin

September 2025

Abstract

Since 2021, Kalshi has operated as the only federally licensed prediction market in the United States. Using transaction-level data on over 300,000 contracts, we provide the first systematic evidence on its pricing. Kalshi’s contract prices are informative and improve in accuracy as markets approach closing, but they display a clear favorite–longshot bias: low-price contracts win far less often than required to break even, while high-price contracts win more often and yield small positive returns. We interpret these patterns with a simple framework that reflects Kalshi’s quote-driven microstructure. Makers—relatively well-informed traders who post offers—seek positive expected returns but may be slightly over-optimistic, while Takers accept these offers. The model predicts distinct patterns of favorite–longshot bias for Makers and Takers, and the data confirm these predictions.

Keywords: Prediction Markets, Kalshi, Favorite-Longshot Bias, Market Microstructure, Bid-Ask Spreads

JEL Classification: G14, G23, G40

*constantin.burgi@ucd.ie

[†]wanying.deng@ucdconnect.ie

[‡]karl.whelan@ucd.ie. We would like to thank Oana Peia for comments and Prachi Srivastava and Jessica Whelan for help with obtaining the data used in the paper.

1. Introduction

Prediction markets match people who place wagers backing events to happen with people who wager that the event will not happen. Promoted in popular books by Surowiecki (2004) and Sunstein (2006), these markets have long been considered useful for their ability to summarize the public’s information but their development was limited by regulatory restrictions in the United States. A 2008 article by 22 high-profile social scientists including four economics Nobel prize winners (Arrow et al. 2008) recommended easing these restrictions and argued for “The Promise of Prediction Markets.”

In November 2020, the Commodity Futures Trading Commission (CFTC) gave the startup firm Kalshi a license for a market of this type (technically, a designated contract market). Importantly, unlike previous prediction markets such as PredictIt or the Iowa Electronic Markets (IEM), Kalshi’s license allows it to operate without stake limits. Kalshi (Arabic for “everything”) operates markets across a huge range of topics such as culture, weather, sports, company announcements, financial markets and politics. Notably, Kalshi called the 2024 US Presidential election correctly, predicting a high likelihood of a Trump win when most polling suggested the race was a toss-up.

Contracts on Kalshi are quoted with prices ranging from 1c to 99c. If you buy a contract for 70c, you can win the 30c put up by the counterparty. Previous work on prediction markets has described how to break even on average, contracts costing X cents must win X percent of the time. However, Kalshi charges fees on contracts so breaking even on average requires contracts to win more often than the percentage implied by their price.

Kalshi is a decentralized quote-driven market. Traders participate by either making offers to others to back Yes or No for an event or by accepting these offers. Kalshi describes the market structure as follows: *You can think of every trade on the platform as a deal taking place between two participants: a “maker” and a “taker”. The maker is the first one to the table: They declare a side they’re willing to buy (Yes or No), how much they’re willing to pay, and how many contracts they are looking to buy at that price. Takers can see all available offers and match with the most generous one.*¹ In stock exchange terminology, Makers make limit orders and Takers make market orders. This makes Kalshi’s pricing mechanism different from the IEM, which took limit orders on both sides of events and had a centralized mechanism to match up order requests from opposite sides when they were mutually compatible. A distinctive feature of Kalshi’s data is that it records directly which side of each trade was initiated by the Maker and which by the Taker, so we can compare Makers and Takers without relying on noisy trade-classification algorithms.

This paper is the first to present systematic evidence on prices in the Kalshi market and also provides a simple theoretical model to explain our findings. We present evidence on the outcomes of traded contracts on Kalshi since the market opened in 2021. We have transaction-level data on 46,282 different Kalshi contracts on specific outcomes (e.g. Will Margot Robbie win the Oscar for best

¹<https://help.kalshi.com/trading/who-are-you-trading-with>

actress?) from 12,403 individual events (e.g. Who will win the Oscar for best actress?) each of which was open for at least 24 hours. We collect the final traded price as the market closed and also, where available, previous prices from 24 hour intervals up to 10 days before markets closed. This gives us over 300,000 prices when we factor in both sides of each traded contract.

We show that Kalshi's contract prices are relatively accurate predictors of outcomes and their accuracy increases as markets come closer to closing. However, we also find that Kalshi's prices display a systematic favorite-longshot bias. Low-price contracts win systematically less often than required to break even, with the opposite for high-price contracts. In particular, investors who buy contracts costing less than 10c lose over 60 percent of their money. In contrast, there is statistically significant evidence that contracts with prices above 50c earn a small positive rate of return. These findings hold across a wide range of different categories of Kalshi's markets, though there is some evidence that the bias in prices is getting smaller over time.

Prior to fees being charged, the average rate of return on Kalshi contracts is minus 20%. Even though total return before fees sums to zero by definition, losses on cheap contracts are proportionally larger than gains on expensive ones. Factoring in Kalshi's fees, the average return across all contracts is minus 22%.

We provide a simple model to explain the observed pattern of favorite-longshot bias that features two key aspects of Kalshi's market structure: The existence of fees and the asymmetry in the market micro-structure in which people are either Makers or Takers. In our model, Makers put forward offers seeking positive expected returns after accounting for fees but may be slightly optimistic about their chances of success. When trades occur, it is because Takers have accepted these offers. The model predicts that loss rates for Takers will increase as the price of contracts fall, with the incidence of the commission and the profits earned by Makers falling harder on those who have bought cheap contracts. The model also predicts Makers will lose money offering cheap contracts if they are too optimistic about their chances of winning. This can be considered as a form of "winner's curse" where those who make the most generous offers turn out to have slightly over-paid.

The data support the model's key predictions. Kalshi's transaction-level data records the Taker's side of each traded contract. Using this information, we find patterns of favorite-longshot bias for contracts bought by both Makers and Takers but the pattern is more pronounced for prices accepted by Takers. We show how these results are consistent with a version of our model in which Makers seek a modest rate of return but are also slightly too optimistic about their chances of winning.

The paper is organized as follows. Section 2 begins with a brief review of previous literature on prediction markets and describes how the Kalshi market operates. Section 3 presents our initial evidence on the performance of contracts. Section 4 presents our theoretical framework and Section 5 provides our results on differences in outcomes for those who offer versus those who accept contracts. Section 6 discusses explanations for the apparent pricing anomalies we have found.

2. Previous Literature and How Kalshi Works

Here we provide a brief description of some of the relevant previous research on prediction markets and then discuss how the Kalshi prediction market works.

2.1. Previous Literature

The first prediction markets pre-dated the Internet. Forsythe et al. (1992) reported that the Iowa Presidential Stock Market, which operated on a computer network at the University of Iowa, yielded predictions of the vote shares in the 1988 US presidential election that outperformed opinion polls. Subsequent research on what became known as the Iowa Electronic Markets (IEM), such as Berg, Nelson and Rietz (2008) and Berg and Rietz (2019) further reported impressive findings on its predictive accuracy. Wolfers and Zitzewitz (2004) reported similarly strong predictive power for a number of other lower-profile prediction markets. To our knowledge, only one other academic paper has used data from Kalshi: Swanson, Wang and Wu (2025) analyze the impact of Fed announcements on prices on Kalshi's markets on macroeconomic data releases.

Of relevance for the results we present below, Berg and Rietz (2019) reported that IEM prices did not exhibit any favorite-longshot bias. Worth noting, though, is the evidence presented by Page and Clemen (2013) that prices from the InTrade prediction market tended to be too high for longshots and too low for favorites. However, this finding related to contracts that were traded well over 10 days in advance of the closing of the market. Page and Clemen described how the discounting of the potential future payout could explain this phenomenon and showed that there were no significant deviations between prices and win probabilities when contracts were close to closing. Our focus will be on contracts with no more than 10 days before closing. Of relevance, also, is Page's (2012) finding of the so-called Yogi Berra bias ("it ain't over till it's over") using InTrade, in which prices for losing teams in sports events in the final 15 minutes were too high relative to the fraction of times they won.

The theoretical literature on prediction markets has largely focused on whether it is reasonable to interpret a prediction market price as an unbiased estimate of the public's average belief about the probability of an event happening. Manski (2006) presented a simple model with risk-neutral investors who invested all their wealth in their preferred contract and showed that when observing a price p in a prediction market, the average belief could be as low as p^2 or as high as $2p - p^2$. Key contributions by Gjerstad (2004) and Wolfers and Zitzewitz (2006) showed that if agents set the size of their trading positions by maximizing log utility and the distribution of beliefs about a probability in the population were symmetric, then the prediction market price that equated demand and supply on either side of the contract was the average belief of participants. He and Treich (2017) showed that this result does not apply for any other CRRA utility functions.

2.2. How Kalshi Works

Before explaining how the Kalshi market works, we highlight two ways it differs from both the IEM and existing theoretical models of prediction markets.

First, Kalshi is a commercial operation and charges fees. IEM was allowed to operate by the CFTC because it was a research project. Stake sizes on IEM were limited to \$500 and no fees were charged. In contrast, Kalshi charges a per-contract fee of $\$0.07p(1 - p)$ where p is the price in dollars, rounding the total up to the nearest cent. It also has no legally binding stake size limit.

Fees did not feature in the theoretical papers just cited. Whelan (2023) used the log utility model of stake sizing to incorporate various different possible fee structures, including Kalshi's, showing that the addition of fees implied that people who purchased contracts backing low probability outcomes would experience higher percentage losses on their investment than those buying contracts backing more likely outcomes. While the model we introduce here is different, it also features the incidence of fees falling more heavily on those who buy low-price contracts. One factor driving this is that Kalshi's per-contract fee of $\$0.07p(1 - p)$ means the participants on both sides of each contract pay the same fee, so fees are higher as a proportion of the total cost of the investment for low-price contracts.

Second, there is the market price mechanism. IEM operated a continuous double-auction system where people on both sides of the market would submit limit orders indicating their willingness to buy at various prices. IEM maintained a centralized order book and would match compatible offers from opposing sides. If the highest price offer to buy a Yes contract was at least 60c and the highest price offer for a No contract was at least 40c, the IEM would match the two sides and record the market price for a Yes contract at 60c. With IEM as the principal source of evidence, the literature understandably focused on there being a single market-clearing price at any point in time. In contrast, Kalshi does not run a centralized double-auction system. Instead, users post offers that are visible on the site and people can choose to either accept or reject these offers. This means there are always two different potential prices at which a Yes contract will be traded: One based on a Maker offering a Yes contract and one based on one of them offering a No contract.

To illustrate how it works, consider a specific Kalshi event market: "What will monthly inflation be in the April 2025 US Consumer Price Index?" Figure 1 shows what the market looked like on May 7, 2025, six days prior to the CPI release, with the upper panel showing the historical traded prices. The figures quoted as percentages on the left-hand-side are the last traded Yes prices for each option, while the blue and purple boxes show the current best offers available to back Yes and No. For example, at a cost of 32c you could buy a contract that will pay out \$1 if CPI inflation in April was over 0.3% and at a cost of 72c you could buy a contract that will pay out \$1 if CPI inflation in April was 0.3% or under. Note that all trades occur at integer-valued number of cents, so it is not possible to buy a contract for 32.5c.

Figure 1: A sample Kalshi market: CPI for April 2025




Figure 2 shows the options available when you click on Yes for Above 0.3% on the 32c button in blue and provides information on volumes in the order book.

At this point in time, people had four different actions they could have taken with the “Yes/No Above 0.3%” market.

1. **Take the best available Yes offer of 32c.** The order book in this case shows you can buy \$33.60 worth of Yes contracts at this price. If you bought more than this, you would have to pay 33c for the rest of your investment, with \$3,336 available at that price. The 33c Yes price (and other more expensive offers to back Yes) shows up under “Asks” on the order book.
2. **Make an offer for a Yes contract for a price below 32c.** If you consider 32c too high, you can make an offer to buy a Yes contract at a lower price. For example, if you post a 30c offer for Yes, it becomes the best available price for someone seeking to back No (i.e. they would pay 70c for No) so 70c would show in the purple box in Figure 1. A key difference with this strategy is that accepting Yes at 32c means you definitely get your wager placed while seeking a cheaper price may mean your offer doesn’t get matched.
3. **Take the best available No offer at 72c.** When you click on the purple 72c button, a second order book shows up. This has the same information as the Yes order book but the prices are quoted from the No perspective. We can see that \$16.80 is available to buy No contracts at a price of 72c (because this is the amount made available by people seeking a Yes contract at 28c).
4. **Make an offer for a No contract for less than 72c.** For example, if you decided to seek a No contract at 69c, then you would have the best offer under “Yes” and it would be your 69c offer that would show up as 31c in the blue box (because $\$1 - 69c = 31c$). If that offer was accepted, you would have bought a No contract for 69c, which is cheaper than you could get from clicking on the No button and taking the 72c offer.

It is worth noting here that the last Yes contract price for Above 0.3% is 28c and this matches the most expensive bid price. This suggests the last trade likely involved someone accepting an offer from a Maker seeking a Yes contract for 28c. As we will describe in Section 5, the transactions-level data provided by Kalshi explicitly identify which side of a trade was the Maker and which was the Taker.

Figure 2: Options for betting for or against above 0.3%

	Price	Contracts	Total
	37¢	61	\$6,859
	35¢	10,000	\$6,836
	33¢	10,008	\$3,336
Asks	32¢	105	\$33.60
Yes	Last 28¢		
Bids	28¢	60	\$16.80
	27¢	10,000	\$2,717
	26¢	10,008	\$5,319
⚙️ ÷ ⓘ	25¢	20	\$5,324

2.3. Bid-Ask Spreads and Limit versus Market Orders

Kalshi's order book system has some similarities to financial markets. However, there are no institutionalized market makers in the Kalshi market that are driving the bid/ask spreads by "buying low and selling high." Instead, the bid/ask spread—for example the 4c gap between the 28c bid and the 32c ask for Above 0.3%—is driven by the process of Makers looking to get the best possible price for their preferred option and their accounting for the fees charged by Kalshi. This behavior means the sum of the best available Yes and No offers on each option is always greater than \$1. Similarly, the total cost of all the best available Yes contracts exceeds \$1, so buying them all guarantees losing money and the sum of all best available No contracts similarly is more expensive than the guaranteed payout.

In contrast, getting a 32c Yes contract offer accepted gives someone a 68c No contract, while getting a 72c No offer accepted gives someone a 28c Yes contract. In this case, getting both the Yes and No offers accepted would cost 96c, which would guarantee a profit of 4c before fees. In this example, the combined fees would equal about 3c, indicating an opportunity for post-fee arbitrage profits for Makers. But this is not a risk-free strategy since Makers would have no guarantee of getting either or both their Yes and No offers accepted and prices may move against them after the first of these offers is accepted. If arbitrage-seeking behavior by Makers is sufficiently strong so that they only broke even after fees, then the bid-ask spreads would be driven solely by fees. Takers pay the same fees as Makers on each trade so these calculations point to a greater chance of making profits by making offers rather than taking them up.

These considerations suggest that, at any point in time, there are perhaps two different kinds of people that participate in these markets. On the one hand, there are people who take up offers rather than seeking a lower price. As we discuss later, these people may be impatient, or so confident that this is a winning trade at the offered price that they don't want to risk looking for a better price but not getting matched. On the other hand, there are people who believe making a profit after commission likely requires acting as a Maker and seeking better prices.

A useful comparison from research on financial markets is the evidence in Barber et al. (2009) from anonymized trading records from the Taiwan Stock Exchange (TSE). The TSE only accepts limit orders but Barber et al. characterize these orders as either aggressive or passive depending on how high a price the bidder is willing to pay.² They reported that at short horizons (up to 10 days), individual investors earned profits on their passive trades but lost money on aggressive trades.

There is evidence, however, that this pattern does not hold at all times. Linnaimaa (2010) examines the performance of market orders and limit orders on the Helsinki Stock Exchange around the time of earnings announcements and finds that market orders perform better than limit orders. His explanation for this effect is relatively simple: When there is good news about a company, market

²Specifically, orders with an order price below the most recent unfilled buy limit order were categorized as passive.

orders to buy get triggered and generate good returns for buyers and bad returns for those with limit orders to sell. The opposite pattern occurs when bad news is announced with those with limit orders to buy getting bad returns. Whether this pattern holds in the Kalshi data likely depends on how “passive” the Makers are: Makers who are aware of good or bad news that affects the outcome of a Kalshi contract are just as free to cancel their unmatched offers as Takers are to take them up. But this argument shows it is an open question whether Makers or Takers perform better on a market like Kalshi that revolves around reacting to news.

3. Initial Evidence

Here, we introduce our dataset and provide evidence on the predictive power of Kalshi’s prices and returns on these investments. We start with summary statistics and then present regression evidence formally testing for a favorite-longshot bias.

3.1. Data and Summary Statistics

To obtain the data, we registered with Kalshi to get access to their Application Programming Interface (API). We used Python scripts to get transaction-level data from the API for contracts from Kalshi’s inception in 2021 through April 2025. We focus only on contracts that have reached a total trading volume upon market closure of at least \$1,000 to ensure there was a meaningful level of market activity. In addition, we drop contracts where the final bid-ask spread is larger than 20c.³ Kalshi runs some markets that re-set every hour, generally markets that speculate on the value of cryptocurrencies or stock market indices at the end of the hour. We do not analyze these markets, limiting ourselves to markets that are open at least 24 hours. The resulting dataset has the last traded Yes prices before the market closed for 12,403 events (e.g. Who will win the Super Bowl?) and 46,282 different Yes contracts (e.g. Will the Chicago Bears win the Super Bowl?)

Going from final trades on the day a market closed back in 24-hour intervals to ten days before closing, we obtained a total of 156,986 prices for traded Yes contracts, so we have a maximum of 11 daily observations for each Yes contract and the average number of daily observations across Yes contracts is 3.5.⁴ Associated with each Yes price is the corresponding No price that was taken by the other side, which means we have 313,972 prices for purchased contracts. We also obtained data on whether contracts were successful or not, which allows us to check the predictive accuracy of prices and the pre- and post-fee returns for all contracts.

As detailed in Table 1, the most common length of time a market is open in our dataset is 24 hours.

³Kalshi also report final traded prices and volumes separately from their transactions-level data. We omitted 63 Yes contracts from the transactions-level data that did not match these figures.

⁴We collected these as follows. If, for example, the last trade on a market was at 5.57pm on the day the market closes, we also record the last trade before 5.57pm for each of the previous days that trades were available.

Examples of these short-lived markets include Kalshi's markets for the top temperature in various cities tomorrow and markets like "Which app will be the most downloaded to iPhones tomorrow?" Other markets are available for longer. So, for example, we have 6,754 observations on Yes contracts that were traded every day back to 10 days before the market closed and 12,861 markets where we record the final traded price and the price from 24 hours earlier.

Some Kalshi events have simple Yes/No outcomes on a single event, so they have one Yes and one No: For example, Kalshi run markets like "Will there be a government shutdown this year" and you can pick either Yes or No. Other markets have multiple mutually exclusive winners and you can back Yes or No for any of them: For example, "Who will be the most popular artist on Spotify this year?" Some, like the CPI market discussed earlier, offer multiple different Yes contracts available with their success dependent on a single published number or else markets like "Which phrases will be mentioned by the Fed chair during the FOMC press conference?" where there are many different phrases and no limit on how many can win. The average number of traded Yes contracts on the final day per event in our data was 3.7.

Table 2 shows the number of combined Yes and No contracts in our data for each of the 10 unit price intervals going from 1c to 10c up to 90c to 99c. It shows that our data contain far more contracts with extremely high or low prices than with mid-range prices. Of the 313,972 contract prices recorded, about two-thirds have prices either less than 10c or greater than 90c. There are only 8,351 contracts with price between 50 and 59.

Table 3 reports the average total final volumes traded by decile. Typical total volumes on Kalshi are low relative to most financial markets with the median amount of money staked on each contract (including money placed on both Yes and No) being \$8,982. The distribution of volumes is highly skewed, with lots of markets having low volumes but those in the highest decile often having total volumes of over \$1 million. Our data also report the size of individual transactions. The mean transaction size is \$100 while the median is \$35.

Table 1: Yes trade observations by number of days available before closing

Data Availability	Observations
Closing day only	19,338
Closing day and day before closing	12,861
Closing day and 2 days before closing	1,723
Closing day and 3 days before closing	1,083
Closing day and 4 days before closing	1,311
Closing day and 5 days before closing	1,288
Closing day and 6 days before closing	1,461
Closing day and 7 days before closing	159
Closing day and 8 days before closing	150
Closing day and 9 days before closing	154
Closing day and 10 days before closing	6,754
Total	46,282

Table 2: Total observations for all
Yes and No contracts in each price
range

Price Range	Number	Percent
1c-10c	106,209	33.8
11c-20c	20,395	6.5
21c-30c	12,558	4.0
31c-40c	10,049	3.2
41c-50c	7,199	2.3
50c-59c	8,351	2.7
60c-69c	10,049	3.2
70c-79c	12,558	4.0
80c-89c	20,395	6.5
90c-99c	106,209	33.8
Total	313,972	100

Table 3: Average final total volume
deciles

Decile	Average Volume	Freq
1	\$1,199	4,632
2	\$1,735	4,634
3	\$2,455	4,621
4	\$3,593	4,626
5	\$5,454	4,629
6	\$8,481	4,628
7	\$13,387	4,628
8	\$21,286	4,628
9	\$35,960	4,628
10	\$526,245	4,628
Total	\$61,977	46,282

3.2. Winning Fractions and Forecast Accuracy

As you can see from Figure 1, Kalshi's website reports its last traded Yes price in the form of percentages with the clear implication that, for example, if we examined a large sample of Yes contracts with prices equal to 60c, then we would expect the events being backed to happen 60% of the time. A simple way to assess this is to sort the data by price and record the fraction of wins for contracts at each price level.

Specifically, consider a contract with price p_{ij} in dollars which returns a payout of \$1 if event i ends in outcome j . Define

$$y_{ij} = \begin{cases} 1 & \text{if outcome } j \text{ occurs} \\ 0 & \text{if outcome } j \text{ does not occur} \end{cases} \quad (1)$$

Figure 3 plots the empirical average win rate $\bar{y}_{ij}$ for each possible value of p_{ij} , as well as a standard 95% confidence interval. We also show a 45 degree line. If this line lies inside the confidence intervals, then a standard t -test can't reject the hypothesis that the winning fraction was consistent with the prediction from the traded price.⁵ One quick conclusion is that Kalshi's prices are pretty good predictors of outcomes. The winning fraction fluctuates around the 45 degree line, particularly in the middle price ranges where the sample sizes are smaller, but overall we can say that the higher a Kalshi price is, the more likely is the outcome.


Interpreting Kalshi's prices as probabilistic forecasts, we can also assess the accuracy of these forecasts and whether they become more accurate as we get closer to the closing of the market. Evidence that Kalshi forecasts improve over time as more information becomes available is provided in Figure 4. We calculated the absolute error for each contract as $|y_{ij} - p_{ij}|$ and then averaged them to obtain the mean absolute error (MAE) for various samples. The top-left chart in Figure 4 uses the data from all events where contracts were available from 10 days prior to market closing up to the closing price and reports the mean absolute error (MAE) of the forecast implied by the prices for each day. We can see that the MAE declines with each day as the market gets closer to closing. The decline is smooth until the last day, at which point there is a steep drop. This pattern is replicated in the other charts, which show MAE calculations for samples of contracts that were available for smaller numbers of days prior to the market closing.

While this is evidence of Kalshi's ability to forecast accurately, there is also a key shortcoming. Figure 3 shows a systematic pattern of contracts with relatively low prices having win rates (and typically also 95% confidence intervals) that lie below the 45 degree line, while the opposite applies for relatively high-price contracts. This mirrors the well-documented favorite-longshot bias in betting markets, which has not been recorded previously for relatively short-dated prediction markets such as the one we are looking at.⁶

⁵We have also estimated this relationship using bootstrapped kernel regression methods and found the same results.


⁶Snowberg and Wolfers (2008) and Ottaviani and Sørensen (2008) are excellent surveys of the theoretical and empirical

Figure 3: Win percentage sorted by price (shaded areas are 95% confidence intervals)



Notes: The figure shows the fraction of contracts that won for each price for the full sample of 313,972 Yes and No contracts.

Figure 4: Mean Absolute Error by forecast horizon (shaded areas are 95% confidence intervals) for samples of contracts available up to 5 days out, 6 days out etc.



Notes: The figures show the MAE that won for each price for different sub-samples of the 313,972 Yes and No contracts for which prices were available continuously from the number of days prior closing indicated.

3.3. Returns on Contracts

We also calculated ex post returns on contracts. Calculating the return before fees are charged is simple but there is a slight complication when calculating post-fee returns. We could calculate the rate of return on purchasing a single contract using the fee for a single contract. However, it is rare that anyone would buy only one contract. Kalshi calculates their fee as $\$0.07p(1 - p)$ times the number of contracts where p is the price in dollars, rounding the total up to the nearest cent. This rounding means the average fee is generally slightly higher than Kalshi's stated fee rate for a single contract. We calculated an average fee based on the purchase of 100 contracts rounded up to the next cent. The round-up makes the fee on an average contract price of 50c equal to 3.54% of the price, slightly higher than the 3.5% it would be without the round-up.

We calculate the pre-fee rate of return on contracts as

$$r_{ij} = \frac{y_{ij} - p_{ij}}{p_{ij}} \quad (2)$$

and given our imputed commission rate for each contract, c_{ij} , we calculate the post-fee rate of return as

$$r_{ij}^c = \frac{y_{ij} - p_{ij} - c_{ij}}{p_{ij} + c_{ij}} \quad (3)$$

In other words, when calculating the return once fees are considered, we treat the total investment to be the cost of the contract p_{ij} plus the associated commission.


Figure 5 shows the average returns on investment in contracts sorted by 10c-size price bands (1c-10c, 11c-19, ..., 91c-99c) with the upper panel showing returns before fees and the lower panel showing the return after accounting for fees. The results are striking. The pattern of low-price contracts under-performing, which was evident in Figure 3, has a dramatic effect on their rate of return. Average loss rates for contracts costing 10c and under are over 60%. These loss rates fall as prices rise above 10c and there are small positive returns for contracts above 50c. For contracts above 70c, there is evidence of statistically significant, though small, positive post-fee returns.

These results are driven by the fact that, while the shortfall in win rates for low-price contracts shown in Figure 3 may appear small, it is a relatively large fraction of the price. A 5c contract that only wins 3% of time will have a minus 40% average return before fees. In contrast, the over-performance of higher-price contracts represents a relatively small amount of extra money relative to their prices. A 95c contract that wins 98% of the time has a pre-fee average return of 3.1%. Prior to fees, Kalshi participants simply swap money, so by definition the average return prior to fees across the total volume of money invested is zero. However, the asymmetric pattern of returns on cheap and expensive contracts means the average pre-fee return on a Kalshi contract in our data is -20%.

The similarity of the upper and lower panels in Figure 5 shows that fees play a relative minor role in the patterns driving average returns. The average post-fee return on a Kalshi contract in our

data is -22%. However, accounting for fees exacerbates the pattern of cheap contracts having inferior returns because the fee paid on a cheap contract is the same as the fee paid by the counter-party that took the other side. For example, the fee when buying one hundred 5c contracts (so an investment of \$5) is 34c, equivalent to almost 7% of the pre-fee cost. For the counter-parties that spent \$95 on the matching contracts, the fee is 0.03% of the pre-fee cost.

Figure 5: Average rates of returns on investment sorted by 10c-size price bands (shaded areas are 95% confidence intervals)



Notes: The upper figure shows the average rate of return prior to considering fees (as calculated by equation 2) and the lower shows the average rate of return taking fees into account (as calculated by equation 3) for various price range sub-samples of the 313,972 Yes and No contracts.

3.4. Regression Evidence

A formal way to test the accuracy of the forecasts implicit in Kalshi's traded prices is to use the classic Mincer-Zarnowitz (1969) regression

$$y_{ij} - p_{ij} = \alpha + \psi p_{ij} + \epsilon_{ij} \quad (4)$$

The null hypothesis of the price being an unbiased predictor can be tested via the standard F -test of $\alpha = \psi = 0$. This approach has been used often to assess whether forecasts are unbiased, including the well-known contribution of Keane and Runkle (1990).⁷ We report results from applying this approach to various sub-samples of our data in Tables 4 to 9.

A few technical aspects of these regressions are worth mentioning. First, because the variable being forecasted takes the value of either 0 or 1, this is effectively a linear probability model. These models can be objected to because they can predict probabilities outside the $[0, 1]$ interval but, in this case, the null hypothesis is that the expected value of y_{ij} increases one for one with p_{ij} and not the nonlinear relationships implied by Logit or Probit models. Also, any predicted value being outside $[0, 1]$ would itself be a violation of the null hypothesis.

Second, the regressions are restricted to the profits and prices of the 156,986 Yes contracts. For a Yes contract with price p_{ij} and profit $y_{ij} - p_{ij}$, the price and profit for the No contract are just 100c minus the Yes contract values so the No observations are just a linear transformation of the Yes observations. This means a regression specification with both Yes and No contracts would double the apparent sample size without adding any useful information, thus distorting the standard errors downward.

Third, standard errors are clustered at the event-level and at the Yes-contract level, because there are negative correlations within observations on the same event (if one contract wins a mutually exclusive event, then the others did not) and because there are positive correlations between the observations across multiple days for the same Yes contract (if the Yes contract won, then the residual in the regression is positive for all the recorded observations on this contract from the first daily observation that we record up to the final price before the market closes.)

The first column in Table 4 reports the results from regression equation 4 for the full sample of 156,986 Yes contracts. As expected from the figures already provided, the rejection of the null hypothesis of prices as an unbiased predictor of outcomes is highly statistically significant. Indeed, the null hypothesis is firmly rejected for all sub-samples we present here. The form of the rejections fit with the favorite-longshot pattern observed in Figure 3. The intercepts are negative and the coefficient on the price is positive, indicating a pattern in which win rates are lower than predicted for low

⁷A more recent application that relates to this example where probabilistic forecasts are assessed is Hegarty and Whelan (2024) which uses the Mincer-Zarnowitz regression approach to test for the accuracy of probabilities based on bookmakers' odds for soccer and tennis.

prices and then rise to being higher than predicted.

The remaining columns in Table 4 show the regression results for sub-samples for various different types of contracts. As noted above, some Kalshi markets have one contract for a specific event (e.g. is the temperature above X on a specific date and specific place?) Others have a series of related contracts for one event. The outcomes being considered can be mutually exclusive or instead have multiple possible winners as was the case with the CPI inflation market we showed earlier. For mutually exclusive contracts, we further distinguished between contracts where the description includes greater, less or in-between and classify these as numerical (e.g. GDP growth less than 2%, between 2-4% or greater than 4%). The remaining mutually exclusive contracts include events like who wins a particular Oscar. The results here show the hypothesis of Kalshi prices being unbiased predictors of outcomes is strongly rejected for all of the different types of contracts, with the ψ coefficient being largest for single contract markets.

Table 4: Regressions of pre-fee profit on price for different types of contracts

	(1)	(2)	(3)	(4)	(5)	(6)
	All	Single	Non Exclusive	Exclusive Numerical	Exclusive Other	All Exclusive
Price	0.034*** (0.005)	0.045*** (0.009)	0.036*** (0.007)	0.014*** (0.004)	0.020** (0.008)	0.017*** (0.004)
Constant	-1.736*** (0.153)	-0.595 (0.611)	-1.839*** (0.499)	-1.637*** (0.116)	-1.916*** (0.186)	-1.756*** (0.105)
Observations	156,986	16,433	58,602	46,674	35,277	81,951
F -test p -value	0.000	0.000	0.000	0.000	0.000	0.000

Notes: The table reports estimates of regression equation 4 in which the dependent variable is pre-fee profit of the specific contract. Single are single events with one contract, Non-Exclusive are non-mutually-exclusive multiple events (e.g. GDP larger than 2% and GDP larger than 3%). Exclusive Numerical are mutually exclusive events reliant on a numerical outcome and Exclusive Other are the remaining mutually exclusive events (e.g. who wins an Oscar). Standard errors clustered at event and contract level except for single contracts (which are only clustered at the contract level) in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5 shows separate regressions for the final traded Yes prices and the Yes prices drawn from the various one-day intervals prior to this, going back to 10 days before the market closed. While Figure 4 had shown prices becoming more accurate as the number of days until market closing declined, the regressions show the F -test still rejecting the null hypothesis for each of the days.⁸

Table 6 examines whether the amount of traded volume influences the favorite-longshot bias pattern that we have found. One could imagine that the favorite-longshot bias inefficiency might only be observed in markets with limited amounts of trading while arbitrageurs intervene in larger-volume markets to align prices better with winning rates. The table reports separate regressions for closing Yes prices for five quintiles based on the total final volume traded before market closing. The null hypothesis of prices being an unbiased predictor of outcomes is rejected for each of the quintiles, with a favorite-longshot bias pattern evident in the coefficients. The lowest volume quintile has the largest ψ coefficient but, other than that, there is no evidence of prices in higher-volume markets being more efficient predictors.

An alternative possibility could be that, rather than total market volume being what matters, prices could be more accurate when the size of the individual transactions is bigger because people pay more attention to fundamentals when they are placing more money at risk. Table 7 shows that this is not the case. Organizing the data into five quintiles by the average size of the transaction associated with each price, the null hypothesis is rejected for each quintile and, in fact, the quintile with the highest average transaction size has the largest ψ coefficient.

Table 8 reports regressions for separate categories of markets such as financials, crypto, climate and weather. Again, the null is rejected for all categories, though the ψ are smaller for politics and entertainment.

Finally, in Table 9, we report the regressions separately for each calendar year. Volumes on Kalshi have grown since 2021 and the market's users may have become more sophisticated. We find, however, that the null hypothesis of prices being an unbiased forecaster of outcomes is rejected for each of the years. There is some evidence of a weakening in the favorite-longshot bias because the ψ coefficient for the 2025 data is smaller and less statistically significant.

⁸Standard errors here were clustered only at the event level because each Yes contract only shows up once in these regressions.

Table 5: Regressions of pre-fee profit on price for different amounts of days prior to market closing

Days	Price	SE	Constant	SE	Observations	<i>F</i> -test <i>p</i> -value
0 Day	0.036***	(0.001)	-2.025***	(0.049)	46,282	0.000
1 Day	0.017***	(0.006)	-1.453***	(0.176)	26,835	0.000
2 Day	0.041***	(0.007)	-1.964***	(0.219)	13,996	0.000
3 Day	0.036***	(0.008)	-1.448***	(0.271)	12,281	0.000
4 Day	0.032***	(0.008)	-1.217***	(0.293)	11,201	0.000
5 Day	0.029***	(0.009)	-1.598***	(0.322)	9,918	0.000
6 Day	0.021**	(0.009)	-1.324***	(0.348)	8,631	0.000
7 Day	0.037***	(0.010)	-1.569***	(0.363)	7,178	0.000
8 Day	0.035***	(0.010)	-1.644***	(0.375)	7,024	0.000
9 Day	0.040***	(0.010)	-1.821***	(0.379)	6,886	0.000
10 Day	0.038***	(0.010)	-1.680***	(0.389)	6,754	0.000

Notes: The table reports estimates of regression equation 4 in which the dependent variable is the pre-fee profit of the specific contract. Standard errors clustered at event level in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6: Regressions of pre-fee profit on price for different final trading volume quintiles

	(1)	(2)	(3)	(4)	(5)
	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Price	0.045*** (0.003)	0.039*** (0.002)	0.032*** (0.003)	0.036*** (0.002)	0.032*** (0.002)
Constant	-2.422*** (0.120)	-2.242*** (0.108)	-1.903*** (0.096)	-1.870*** (0.093)	-1.668*** (0.098)
Observations	9,266	9,247	9,257	9,256	9,256
<i>F</i> -test <i>p</i> -value	0.000	0.000	0.000	0.000	0.000

Notes: The table reports estimates of regression equation 4 in which the dependent variable is pre-fee profit of the specific contract for the last price of each contract. Quintile 1 are the 20 percent of contracts with the lowest volume, Quintile 5 the ones with the highest volume. Standard errors clustered at the event level in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7: Regressions of pre-fee profit on price for different mean transaction size quintiles

	(1)	(2)	(3)	(4)	(5)
	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Price	0.036*** (0.002)	0.033*** (0.002)	0.038*** (0.002)	0.034*** (0.003)	0.043*** (0.002)
Constant	-2.211*** (0.130)	-1.847*** (0.104)	-1.798*** (0.105)	-1.917*** (0.097)	-2.346*** (0.089)
Observations	9,257	9,256	9,257	9,256	9,256
<i>F</i> -test <i>p</i> -value	0.000	0.000	0.000	0.000	0.000

Notes: The table reports estimates of regression equation 4 in which the dependent variable is pre-fee profit of the specific contract for the last price of each contract. Quintile 1 represents the 20 percent of contracts with the smallest mean transaction size, Quintile 5 the ones with the largest mean transaction size. Standard errors clustered at the event level in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 8: Regressions of pre-fee profit on price by different categories of contracts

	(1)	(2)	(3)	(4)
	All	Financials	Climate & Weather	Crypto
Price	0.034*** (0.005)	0.032*** (0.004)	0.031*** (0.005)	0.058*** (0.014)
Constant	-1.736*** (0.153)	-1.431*** (0.164)	-0.997*** (0.243)	-1.944** (0.756)
Observations	156,986	27,123	29,924	8,150
<i>F</i> -test <i>p</i> -value	0.000	0.000	0.000	0.000

	(5)	(6)	(7)	(8)
	Politics	Entertainment	Economics	Other
Price	0.022 (0.021)	0.020 (0.012)	0.034*** (0.010)	0.053*** (0.009)
Constant	-1.912*** (0.340)	-2.809*** (0.364)	-0.978 (0.972)	-2.392*** (0.467)
Observations	26,819	25,541	24,405	15,192
<i>F</i> -test <i>p</i> -value	0.000	0.000	0.000	0.000

Notes: The table reports estimates of regression equation 4 where the dependent variable is pre-fee profit of the specific contract. Columns split the sample by topical category (Financials, Climate & Weather, Crypto, Politics, Entertainment, Economics, Other). Standard errors clustered at the event and contract level in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 9: Regressions of pre-fee profit on price for different years

	(1)	(2)	(3)	(4)	(5)
	2021	2022	2023	2024	2025
Price	0.041*** (0.015)	0.023** (0.011)	0.036*** (0.009)	0.048*** (0.006)	0.021* (0.011)
Constant	-1.649 (1.224)	-1.589*** (0.366)	-1.531*** (0.357)	-1.793*** (0.289)	-1.851*** (0.253)
Observations	3,855	24,913	23,559	53,338	51,321
<i>F</i> -test <i>p</i> -value	0.026	0.000	0.000	0.000	0.000

Notes: The table reports estimates of regression equation 4 in which the dependent variable is profit of the specific contract. Standard errors clustered at the event and contract level in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

4. Makers and Takers: A Simple Model

Here we present a simple model inspired by Kalshi's market structure in which people can act as either Makers or Takers. As we described earlier, Makers do not accept offers, instead seeking better prices while risking not being matched. This suggests that Makers may take a more strategic approach to trading than Takers. Here, we will assume that Makers seek a specific positive expected rate of return on their investment. We describe the implications of this for both types of participants. We do not provide a theoretical explanation of the behavior of Takers. For now, they can be considered "noise traders" who accept the best available trade offers. We will discuss potential behavioral explanations for the conduct of real-world Takers later.

4.1. The Model

There is an event that will happen with probability π . We consider the prices that will be offered on this event by Makers who seek to make a return on their investment, taking Kalshi's commission into account. We allow for the possibility that Makers may differ slightly in their evaluation of the true probability π and that the marginal Maker that makes the best available offer—and thus sets the traded price—may be too optimistic in their assessment of their chances of winning. So, for example, when the marginal Maker offers someone else a contract that pays out if the event happens, they believe the event's probability of happening is actually $\pi - \delta$. Intuitively, δ can be interpreted as the degree to which a Maker is overly-optimistic. This leads to them making a more generous offer than the other Makers, who choose not to match their price because it does not meet their rate of return requirement. This can be considered a form of the "winner's curse" whereby the person who makes the best offer for something has likely paid too much.⁹

If a Maker offers a contract with price p , then they will lose $1 - p$ if the event happens and win p if the event does not happen. In addition, both sides will pay a commission of $\theta p(1 - p)$ to Kalshi. The Maker can be considered to have made an investment of $1 - p + \theta p(1 - p)$, the price of their contract plus the commission. We assume they require a subjective expected return of γ on this investment. From this, the following relationship equating their subjective expected profit with the required return must hold

$$(1 - \pi + \delta)p - (\pi - \delta)(1 - p) - \theta p(1 - p) = \gamma[(1 - p) + \theta p(1 - p)] \quad (5)$$

Rearranging we obtain the following quadratic equation

$$\theta(1 + \gamma)p^2 + (1 - \theta(1 + \gamma) + \gamma)p + \delta - \pi - \gamma = 0 \quad (6)$$

This has two solutions but for sufficiently small δ the combined term $\delta - \pi - \gamma$ is negative. This means

⁹See Thaler (1988) and Charness and Levin (2009).

one of the solutions is positive and one is negative, so only one of the two solutions is economically meaningful and corresponds to the Taker's price. From this price, we can calculate the post-fee average returns on contracts.

The positive solutions to equation 6 imply that Takers do systematically worse than Makers and that the loss rates on their contracts get bigger as the price of the accepted contracts fall. Even if Makers only seek to break even after the commission, the existence of a commission can cause cheap contracts to have significant loss rates. The quadratic equation defining the prices sought by Makers is complex but the intuition for this result is simple. Consider a Maker who decides to offer a Yes contract on an event with a 1% likelihood of happening. The commission on contracts priced close to either $1c$ or $99c$ is about $0.07c$. For the Maker to break even, they will need their No contract price to be lowered from $99c$ by this amount and thus the Yes contract price to be that much higher than $1c$. In addition to paying approximately $1.07c$, the Taker also incurs the $0.07c$ commission. The Taker's combined average loss rate on their total investment of approximately $\$1.14$ (the Yes price plus commission) will be 13%.

More generally, consider the case where the Maker has $\delta = 0$, seeks to break even after commission fees ($\gamma = 0$) and the commission on both sides of a contract is c . This requires the price of the contract accepted by the Taker to be $p = \pi + c$ and their total investment inclusive of commission is $\pi + 2c$. The rate of return on the Taker's investment in this case would be $\frac{-c}{\pi+2c}$, which clearly worsens as π declines.

4.2. Illustrating the Results

To illustrate the model's results, we start with the assumption that the Maker is correct in their belief about the event's probability of happening ($\delta = 0$). Figure 6 shows the prices Takers will pay for contracts with varying probabilities of success for three different values of the required rate of return: $\gamma = 0$, $\gamma = 0.02$ and $\gamma = 0.04$. The purple line is a 45 degree line showing what prices would be if $p = \pi$. The other lines show that prices for Takers always lie above this line, implying they incur losses on average before fees.

For $\gamma = 0$, when Makers are only expecting to break even, the magnitude of the gap between p and π is smallest at the extremes and largest in the middle. This is because this gap depends only on fees and the size of fees rises as we go up from $1c$, reaching a maximum at $50c$ and then falling again in a symmetric fashion. However, as a fraction of the contract price paid by the Taker, these fees fall as the price rises. For higher values of γ , the overpayment by Takers becomes more pronounced. In addition to their over-payment covering the fees paid by Makers, it must now also provide a rate of return for them. For a low-price contract bought by a Taker, say $5c$, to provide a 4% post-fee return on the Maker's $95c$ contract, this must add almost $4c$ to the cost already required to cover fees. The result is a much larger gap between p and π .

Figure 7 shows the implications of these prices for returns on investment for Takers. For $\gamma = 0$, loss rates decline linearly with π (the lowest value of 0.13 is consistent with the calculation just described above). For the higher values of γ , loss rates fall nonlinearly as π falls with particularly big loss rates for events with very low values of π . The reason for this is simple. If a Maker buys a contract with $\pi = 0.9$ expecting a 4% return, then even without factoring in fees, they will require a price of about 86c, implying the counter-party Taker is buying a contract with a 10% chance of winning for a price of about 14c, hence losing about 30% on average before fees. As the rate of return required by the Maker goes up, loss rates for Takers rise more than proportionately.

The pattern of losses for Takers looks a bit similar to the pattern seen for all contracts in Figure 5 but since here we have half the contracts earning returns of either 2% or 4%, the average returns here across all contracts are a bit too high relative to what is seen in the data and would require higher values of γ and enormous loss rates for Takers for low π contracts.

A mechanism in the model that would allow for the possibility of Makers also losing money on cheap contracts is over-optimism about their investment. Figure 8 shows the return on the Maker's contracts when $\delta = 0.005$. This implies a very small amount of over-optimism on the part of the Makers and leads to them offering contracts that are too cheap. As the figure shows, for most values of π , this over-optimism does not prove to be too costly. However, for events that are unlikely to occur, this can be very damaging for Makers. For example, leaving fees aside, if I believe the probability of an event occurring is 0.015 and it is actually 0.02, then I would believe that offering someone else a contract at 1.5c would break even, when in fact it would have a positive return of 25% for the Taker and a loss rate of 33% for the Maker.

Figure 9 shows the returns for Takers with over-optimism of $\delta = 0.005$ and different rates of return for the Maker. If the Maker is just seeking to break even ($\gamma = 0$), the over-optimism leads to profits for Takers, with particularly big returns for cheap contracts, in line with the calculation just reported. For positive values of γ , the over-optimism implies a slightly tempered form of the nonlinear loss pattern shown previously.

Figure 10 shows the average returns for Makers and Takers under the assumptions that $\delta = 0.005$ and $\gamma = 0.04$, so Makers are over-optimistic but also seek subjective expected returns of 4%. Rates of return for both Makers and Takers show a nonlinear pattern with large loss rates for contracts with low values of π but the loss rates for Takers are considerably larger than for Makers. We will show below that this pattern fits the data well.

Figure 6: Prices for contracts accepted by Takers for $\theta = 0.07$, $\delta = 0$ and various values of π and γ




Figure 7: Rates of return including commission for Takers for $\theta = 0.07$, $\delta = 0$ and various values of π and γ




Figure 8: Rates of return including commission for Makers for $\theta = 0.07$, $\delta = 0.005$ and various values of π and γ




Figure 9: Rates of return including commission for Takers for $\theta = 0.07$, $\delta = 0.005$ and various values of π and γ





Figure 10: Post-commission return rates for Makers and Takers for $\theta = 0.07$, $\delta = 0.005$, $\gamma = 0.04$ and various values of π



5. Evidence on Makers and Takers

Our data allow us to test the model's predictions about outcomes for Makers and Takers because, for every transaction, the data available from Kalshi's API records which side of a contract (either Yes or No) was the Taker's. This level of detail is unusual in empirical work on market microstructure. The standard Lee and Ready (1991) algorithm infers trade direction from quotes and transactions, but misclassifies 10–20% of trades in many applications (Odders-White 2000; Theissen 2001). Kalshi's direct identification of Makers and Takers therefore eliminates a major source of measurement error and provides a degree of precision that is rarely available in this literature.

Table 10 shows the number of observations for contracts bought by Makers for each price range as well as the total number of contracts. It shows that while Makers do buy many lower-priced contracts, they are more likely to buy higher-priced ones. This means the contracts that cost 10c or less, which perform particularly poorly, are mainly bought by Takers.

Table 10: Total observations for all Yes and No contracts in each price range and for contracts bought by Makers

Price Range	Total	Makers	Makers Share
1c-10c	106,209	46,185	43.5%
11c-20c	20,395	9,527	46.7%
21c-30c	12,558	6,136	48.9%
31c-40c	10,049	4,799	47.8%
41c-50c	7,199	3,565	49.5%
50c-59c	8,351	4,210	50.4%
60c-69c	10,049	5,250	52.2%
70c-79c	12,558	6,422	51.1%
80c-89c	20,395	10,868	53.3%
90c-99c	106,209	60,024	56.5%
Total	313,972	156,986	50.0%

Mincer-Zarnowitz F -tests, applied separately to Makers-only and Takers-only contracts, both firmly reject the null hypothesis of the traded price being an unbiased forecaster of outcomes. Again, the violation is due to low cost contracts not winning often enough and higher price contracts winning more often than needed to break even.

Figure 11 repeats the chart for average post-fee returns by price range, previously shown for all contracts in Figure 5b, but this time separately with returns for Takers on the left and for Makers on the right. The first clear pattern is that average returns for Makers clearly exceed those for Takers: The average return on contracts for Makers was -11.99% while for Takers it was -31.46%. The F -statistic for the test of the null hypothesis that the average return for Makers equals the average return for Takers rejects the hypothesis with an extremely high level of significance. This finding is perhaps not surprising. By definition, Makers seek better prices than Takers, so we might expect them to get better returns but the evidence from Linnaimaa (2010) discussed earlier is worth keeping in mind. Makers need to not simply passively place orders but be willing to cancel those orders if new evidence emerges that influences the likely success of their preferred contract.

Beyond average returns for the two sides, our model also makes more specific predictions, particularly that loss rates for Takers will increase as the price of the contract falls and that Makers will also make losses on low-price contracts if they are too optimistic in their assessment of probabilities ($\delta > 0$). Both charts in Figure 11 show nonlinear patterns with increasingly negative returns as contract prices fall but the pattern is more pronounced for Takers than for Makers. There is some evidence that Makers who buy contracts priced above 50c earn statistically significant positive returns but these returns are small relative to the huge loss rates for Takers who buy cheap contracts. On average, Makers who buy contracts costing 50c and over make a 2.6% rate of return before commission and a 1.9% rate of return including commission.


Because we include the same contract at different points during its lifetime, we want to make sure that our results are not driven by the small minority of contracts that are in the sample up to 11 times. At the same time, we want to assess whether this pattern identified in Figure 11 also holds during the various different days before closing. To this end, Figures 12 and 13 reproduce the charts for different times during the lifetimes of contracts. For example, the top-middle sub-figures show the returns using final outcomes with confidence intervals for all prices available 8 days before closing. For each of the days relative to closing, we see the nonlinear pattern of worsening returns for Takers as prices fall, with large and statistically significant loss rates for the prices of 10c and below and typically also for prices in the 11-19c range.

Returns for Makers also continue to show the pattern of returns generally worsening as the price of their contract falls, though most of these returns are not statistically significantly different from zero. The statistical significance of positive returns for Makers for higher-price contracts is generally not replicated here, most likely due to the smaller sample sizes and thus larger standard errors. Still,

contracts that Makers buy for 10c or less have statistically significant negative rates of return for 5 of the 6 days shown here. Interestingly, the pattern of losses for Makers on closing prices (bottom left) is similar to those for Takers. This may point to over-optimism from Makers being a bigger issue as the market comes closer to closing than in previous days. Indeed, the larger losses for Makers and Takers on low-price contracts in the final day suggests a version of Page's (2012) "Yogi Berra effect" may be displayed in these markets.


Overall, these results are broadly consistent with the final calibration of the model illustrated above in Figure 10, in which Makers seek a return of about $\gamma = 0.04$ but are slightly too optimistic about their probability of winning and so offer prices consistent with them thinking their chance of winning is a half percentage point higher than it actually is ($\delta = 0.005$). This calibration of the model captures both the shape and sizes of the loss rates for Takers as well as the pattern for Makers of earning small profits for high-price contracts and significant losses (though smaller than for Takers) on low-price contracts.

Figure 11: Rates of return including commission by price (by 10c-size price bands) for Makers (right) and Takers (left) (shaded areas are 95% confidence intervals)




Notes: The left figure shows the average rate of return taking fees into account (as calculated by equation 3) for various price range sub-samples of the 156,986 contracts bought by Takers. The right figure shows the same for the 156,986 contracts bought by Makers.

Figure 12: Rates of return including commission by price (by 10c-size price bands) for Takers (shaded areas are 95% confidence intervals)



Notes: The figures show the average rate of return taking fees into account (as calculated by equation 3) for various price range sub-samples of the 156,986 contracts bought by Takers and for various numbers of days prior to the market closing.

Figure 13: Rates of return including commission by price (by 10c-size price bands) for Makers (shaded areas are 95% confidence intervals)



Notes: The figures show the average rate of return taking fees into account (as calculated by equation 3) for various price range sub-samples of the 156,986 contracts bought by Makers and for various numbers of days prior to the market closing.

6. Discussion

Here, we provide some further discussion of potential explanations for our results.

6.1. Behavioral Biases and Platform Design

Given the high loss rates recorded, some participants in Kalshi's markets may be influenced by different possible behavioral biases.

The first bias, **over-optimism**, is something that we have explicitly modeled for Makers. Our baseline approach was that Makers sought a subjective expected rate of return on their investments and had accurate beliefs about event probabilities ($\delta = 0$) but this was inconsistent with the patterns shown in Figures 11 and 13. The model only matched the relatively large losses experienced by Makers on cheap contracts once we assumed they were overly optimistic about the option they were backing. We did not design this over-optimism to be specific to low probabilities, modeling it as an additive error relative to the true probability, but the implications were that Makers incurred significant average loss rates on contracts with particularly low probabilities of winning.

The more extreme losses on low-price contracts for Takers also suggest that these traders may **overweight the likelihood of events with small probabilities** in a more extreme way than defined by our additive assumption for Makers. There has been a lot of experimental evidence for the idea that people overestimate the likelihood of low-probability events, which was formalized as part of Prospect Theory by Kahneman and Tversky (1979), and has regularly been cited as a potential explanation for the favorite-longshot bias exhibited in sports betting.¹⁰ We can note that some of the other explanations for favorite-longshot bias in sports betting settings, such as manipulation of odds by bookmakers with superior information (as described by Shin, 1991, Levitt, 2004 and Hegarty and Whelan, 2005) or the combination of disagreement and the specific market structure of pari-mutuel betting markets (as discussed by Ali, 1977, Ottaviani and Sørensen, 2010 and Gandhi and Ricardo Serrano-Padial, 2015) do not apply here and yet the bias is still observed.¹¹ It is also interesting that this bias is evident in Kalshi's prediction market and not in the previously reported results for the IEM, which used a single price double auction to match trades.

The **design of Kalshi's platform** is likely to particularly encourage those who are highly over-optimistic to make bad financial decisions. A review of previous research on prediction markets would tell you that their prices tend to be relatively accurate summaries of the probability of winning.

¹⁰See Ottaviani and Sørensen (2008) and Snowberg and Wolfers (2008) for two excellent summaries of the theoretical and empirical literature on the favorite-longshot bias in sports betting markets.

¹¹Bias in pari-mutuel markets due to disagreement can be explained as follows. Suppose there are two competitors, one of which has a probability of 0.7 of winning and the pari-mutuel operator does not take any money out. Decimal odds of $\frac{1}{0.7}$ on the favorite and $\frac{1}{0.3}$ on the longshot will equate expected returns on the two bets. Suppose the public is correct in its median assessment of the probability of the favorite winning. At these prices, half the public would back the favorite and half would back the longshot. But the pari-mutuel system means this implies equal odds for the two competitors, which is a contradiction. Equilibrium requires better odds for the favorite than the longshot.

Someone who believes this would likely conclude that they would lose money on average once you factor in Kalshi's fee. And indeed, on average this is correct. Kalshi earns money from commissions and its participants on net lose money. This illustrates a paradox of commercial prediction markets previously discussed by Whelan (2023): Those with the most accurate views will most likely not participate.

Those who think their assessment of event probabilities gives them a small "edge" (such as the over-confident Makers that we modeled above) may think they can earn a positive expected return but perhaps realize that they are more likely to do so as a Maker rather than a Taker. However, those who have a high level of over-optimism, and thus believe they have a big edge, may not want to risk missing out on their expected profits by seeking a better price but then possibly not getting matched. The market structure thus encourages those who are overly optimistic in their beliefs to lose to those who are more realistic and the evidence is consistent with there being plenty of people willing to accept the offers from Makers of cheap contracts at prices that produce terrible returns.

Our discussion so far could be interpreted as suggesting that Makers and Takers are two different groups of people with different behavioral traits but this isn't necessarily the case. A person could have accurate beliefs about one event and believe they can only earn profits as a Maker while also having less accurate beliefs about another event and believing that it is worth being a Taker rather than risk not getting matched. However, there may be some behavioral traits that make people more likely to be Makers than Takers. For example, differences in **impatience** may explain why some people want to get their trade placed immediately rather than waiting to see if it gets matched later at a better price. There may also be differences in people's views on the **complexity** of the two roles, with getting a trade immediately completed seen as having a lower "mental cost" than placing a limit order that may or may not be filled.

Finally, it is possible that the behavior of Takers (and that of Makers towards low contract prices) could reflect a **risk-loving preference for high-variance bets**. This idea has often been suggested as an explanation of favorite-longshot bias, for example by Quandt (1986) in the context of pari-mutuel betting markets. More recently, Moskowitz and Vasudevan (2022) present evidence of a favorite-longshot bias in "moneyline" bets with US sportsbooks and the absence of this bias in bets based on points spreads. They point to several pieces of evidence suggesting an appetite among some bettors for high-variance "lottery" type bets may explain this difference.

Unfortunately, we have no information on the identity of the participants in transactions (such as anonymized ID numbers), so we can't find out whether some people tend to mainly act as Makers or mainly act as Takers or whether certain people like to purchase cheap contracts. Given this limitation, there is little we can do with our data to distinguish between the various interpretations put forward here.

6.2. Why Aren't Biases Competed Away?

An interesting question about our results is why investors have not come in to undercut the profits made on higher-priced contracts by Makers and thus eliminate the favorite-longshot bias. We have noted that the rate of return after commission for Makers on contracts costing 50c and over is 1.9%. This may seem modest when compared with annualized returns on most financial instrument but, of course, it is not an annualized return. Someone could buy a Kalshi contract a few days before closing, make profits and then re-invest those profit. So there could be potential for large cumulative returns. We can point to three potential reasons why this hasn't happened up to now.

Small Volumes: The first potential reason is the small volumes on Kalshi's markets. Even its largest-volume markets are small relative to the kinds of markets that professional investors will typically be willing to participate in. Table 3 showed average final trading volume in the top decile of Kalshi markets was only \$526,245. And at any point in time, the amount of liquidity available is far smaller: Notice the relatively small amounts available in the order book for the April CPI shown in Figure 2. This likely means that someone who wanted to invest lots of money as a Maker seeking to buy high-price contracts may have to post prices that are less advantageous to Makers than the typical trades that we recorded here. This would further reduce the potential size of investments that could be made to exploit the bias in prices documented here. Furthermore, some of the attempts to trade as a Maker would not be matched, again reducing the overall amount actually invested.

Riskiness: The second potential reason is the riskiness of investments in Kalshi contracts. The asymmetric payoff structure of these contracts means that the standard deviations of returns dwarf the averages. The standard deviation of the rate of return on contracts bought by Makers costing 50c and over is 33%. This huge variance may deter many investors from buying these contracts. A counter-argument is that a strategy of making lots of these investments could rely on the Law of Large Numbers to essentially guarantee profits with relatively little risk. However, Paul Samuelson (1963) famously argued that this idea ran counter to expected utility theory: If you rejected one gamble, you should also reject two independent versions of the same gamble. Pratt and Zeckhauser (1987) called this property of rejecting multiple independent gambles that would be rejected on a standalone basis "proper risk aversion" and demonstrated that this property applied to all power utility and logarithmic utility functions, i.e. those used in the vast majority of modern financial economic theory. It is thus possible that rational investors that have examined Kalshi's market have concluded that the positive expected returns for Makers on higher-price contracts represent no more than an appropriate return for the risk taken.

Lack of Information: The third potential reason is that perhaps participants in Kalshi's markets have

not been aware of the favorite-longshot bias pattern that we have documented or the evidence for positive returns for Makers when buying higher-priced contracts. We think it will be interesting to see if the biases and return patterns that we have reported persist now that they have been publicly documented or whether they follow the pattern of other financial market anomalies in tending to disappear once they have been publicized (see Zaremba, Umutlub and Maydybura, 2020 and Shanaev and Ghimire, 2021).

7. Conclusion

This is the first academic paper to present a systematic study of prices on the Kalshi prediction market. Originally, we had two separate questions. The first was to check whether Kalshi’s prices could be considered unbiased probabilistic forecasts of events. The second was to check whether its quote-driven “Maker-and-Taker” market structure (which differed from the IEM double auction approach that had previously been the focus of empirical and theoretical research on prediction markets) had implications for the properties of its prices. We believe our results show that these two questions are related. Our model, while simple, illustrates how Kalshi’s market microstructure encourages there to be a favorite-longshot bias and influences the form this takes for both Makers and Takers.

We have also shown that this bias leads to very large loss rates for low-price contracts, particularly for those who buy these contracts as Takers. We have speculated about the reasons why some of the participants in Kalshi are willing to buy these contracts but, given the restricted amount of data available from Kalshi’s API, there are limited options available to explore the various hypotheses. We do believe, however, that the evidence is consistent with participants who are impatient and overly optimistic being more likely to enter the market as Takers rather than Makers.

Our results that Makers do better than Takers are consistent with the evidence of Barber et al. (2009) from Taiwan that investors who make “passive” limit orders at prices below recent trades perform better than those willing to buy at market prices. However, it would perhaps be a mistake to characterize Makers on Kalshi as passive. While Linnaïmaa (2010) notes how stock exchange limit orders can perform badly in response to news as market orders exploit the news and inflict losses on those who placed limit orders, it seems likely that Makers on Kalshi are just as active as Takers in responding to news, canceling and updating their offers based on either good or bad news.

Overall, our analysis shows that Kalshi’s distinctive fee structure and Maker–Taker market design generate systematic departures from unbiased pricing, with losses concentrated among Takers of low-priced contracts. These findings highlight how market microstructure can shape prediction market outcomes and suggest that design features, as much as information aggregation, determine whether prices provide reliable forecasts.

References

- [1] Ali, Mukhtar (1977). "Probability and Utility Estimates for Racetrack Bettors," *Journal of Political Economy*, Volume 85, pages 803-815.
- [2] Arrow, Kenneth J. et al (2008). "The Promise of Prediction Markets," *Science*, Vol 320, Issue 5878, pages 877-878.
- [3] Barber, Bruce, Yi-Tsung Lee, Yu-Jane Liu and Terrance Odean (2009). "Just How Much Do Individual Investors Lose by Trading?" *Review of Financial Studies*, Volume 22, pages 609-632.
- [4] Berg, Joyce, Forrest Nelson and Thomas Rietz (2008). "Prediction Market Accuracy in the Long Run," *International Journal of Forecasting*, Volume 24, pages 285-300.
- [5] Berg, Joyce, and Thomas Rietz (2019). "Longshots, Overconfidence and Efficiency on the Iowa Electronic Market," *International Journal of Forecasting*, Volume 35, pages 271-287.
- [6] Charness, Gary and Dan Levin (2009). "The Origin of the Winner's Curse: A Laboratory Study," *American Economic Journal: Microeconomics*, Volume 1, page 207-236.
- [7] Forsythe, Robert, Forrest Nelson, George Neumann and Jack Wright (1992). "Anatomy of an Experimental Political Stock Market," *American Economic Review*, Volume 82, pages 1143-1161.
- [8] Gandhi, Amit and Ricardo Serrano-Padial (2015). "Does Belief Heterogeneity Explain Asset Prices: The Case of the Longshot Bias," *Review of Economic Studies*, Volume 82, pages 156-186.
- [9] Gjerstad, Steven (2004). Risk Aversion, Beliefs and Prediction Market Equilibrium. University of Arizona Working Paper 04-17.
- [10] He, Xue-Zhong and Nicolas Treich (2017). "Prediction Market Prices under Risk Aversion and Heterogeneous Beliefs," *Journal of Mathematical Economics*, Volume 70, pages 105-114.
- [11] Hegarty, Tadgh and Karl Whelan (2024). "Comparing Two Methods for Testing the Efficiency of Sports Betting Markets," *Sports Economics Review*, Volume 8, paper 100042.
- [12] Hegarty, Tadgh and Karl Whelan (2025). "Market Structure and Prices in Online Betting Markets: Theory and Evidence," forthcoming, *Oxford Economic Papers*.
- [13] Kahneman, Daniel and Amos Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, Volume 47, pages 263-291.
- [14] Keane, Michael P. and David E. Runkle (1990). "Testing the Rationality of Price Forecasts: New Evidence from Panel Data," *American Economic Review*, Volume 80, pages 714-735.

- [15] Lee, Charles, and Mark Ready (1991). "Inferring Trade Direction from Intraday Data," *Journal of Finance*, Volume 46, pages 733–746.
- [16] Levitt, Steven (2004). "Why Are Gambling Markets Organised So Differently From Financial Markets?," *Economic Journal*, Volume 114, pages 223–246.
- [17] Linnainmaa, Juhani (2010). "Do Limit Orders Alter Inferences about Investor Performance and Behavior?," *Journal of Finance*, Volume 65, pages 1473–1506.
- [18] Manski, Charles (2006). "Interpreting the Predictions of Prediction Markets," *Economic Letters*, Volume 91, pages 425–429.
- [19] Mincer, Jacob and Arnold Zarnowitz (1969). "The Evaluation of Economic Forecasts" in *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*. New York: National Bureau of Economic Research.
- [20] Moskowitz, Tobias and Kaushik Vasudevan (2022). Betting Without Beta, working paper, Yale University.
- [21] Odders-White, Elizabeth R. (2000). "On the Occurrence and Consequences of Inaccurate Trade Classification," *Journal of Financial Markets*, Volume 3, pages 259–286.
- [22] Ottaviani, Marco and Peter Norman Sørensen (2008). "The Favorite-Longshot Bias: An Overview of the Main Explanations," in *Handbook of Sports and Lottery Markets*, Elsevier.
- [23] Ottaviani, Marco and Peter Norman Sørensen (2010). "Noise, Information, and the Favorite-Longshot Bias in Parimutuel Predictions," *American Economic Journal: Microeconomics*, Volume 2, pages 58–85.
- [24] Page, Lionel (2012). "'It Ain't Over Till it's Over' Yogi Berra Bias on Prediction Markets," *Applied Economics*, Volume 44, pages 81–92.
- [25] Page, Lionel and Robert Clemen (2013). "Do Prediction Markets Produce Well-Calibrated Probability Forecasts?," *Economic Journal*, Volume 123, pages 491–513.
- [26] Pratt, John and Richard Zeckhauser (1987). "Proper Risk Aversion," *Econometrica*, Volume 55, pages 143–154.
- [27] Quandt, Richard (1986). "Betting and Equilibrium," *Quarterly Journal of Economics*, Volume 101, pages 201–208.
- [28] Samuelson, Paul (1963). "Risk and Uncertainty: A Fallacy of Large Numbers," *Scientia*, Volume 98, pages 108–113.

- [29] Shanaev, Savva and Binam Ghimire (2021). "Efficient Scholars: Academic Attention and the Disappearance of Anomalies," *The European Journal of Finance*, Volume 27, pages 278-304
- [30] Shin, Hyun Song (1991). "Optimal Betting Odds Against Insider Traders," *Economic Journal*, Volume 101, pages 1179-1185.
- [31] Snowberg, Erik and Justin Wolfers (2008). "Examining Explanations of a Market Anomaly: Preferences or Perceptions?" in *Handbook of Sports and Lottery Markets*, Elsevier.
- [32] Sunstein, Cass (2006). *Infotopia: How Many Minds Produce Knowledge*. Oxford University Press.
- [33] Surowiecki, James (2004) *The Wisdom of Crowds*. Abacus.
- [34] Swanson, Eric, Renxuan, Wang and Yanbin Wu (2025). The Effects of Monetary Policy on Macroeconomic Expectations: High-Frequency Evidence from Traded Event Contracts, working paper, China Europe International Business School, Shanghai.
- [35] Thaler, Richard (1988). "Anomalies: The Winner's Curse," *Journal of Economic Perspectives*, Volume 2, pages 191–202.
- [36] Theissen, Erik (2001). "A Test of the Accuracy of the Lee–Ready Trade Classification Algorithm," *Journal of International Financial Markets, Institutions and Money*, Volume 11, pages 147–165.
- [37] Whelan, Karl (2023). "On Prices and Returns in Commercial Prediction Markets," *Quantitative Finance*, Volume 23, pages 1699-1712.
- [38] Wolfers, Justin and Eric Zitzewitz (2004). "Prediction Markets," *Journal of Economic Perspectives*, Volume 18, pages 107–126.
- [39] Wolfers, Justin and Eric Zitzewitz (2006). Interpreting Prediction Market Prices as Probabilities. NBER working paper 12200.
- [40] Zaremba, Adam, Mehmet Umutlub and Alina Maydybura (2020). "Where Have the Profits Gone? Market Efficiency and the Disappearing Equity Anomalies in Country and Industry Returns," *Journal of Banking and Finance*, Volume 121. Article 105966.